Mapping the Gravitational-Wave Background

Arianna Renzini

Department of Physics

Imperial College London

A thesis submitted for the degree of

Doctor of Philosophy

September 2020

Abstract

Given the recent detection of gravitational waves (GWs) from individual sources it is almost a certainty that some form of background of gravitational waves will be detected in future. The most promising candidates for such a detection are backgrounds made up of the incoherent superposition of the signal of unresolved astrophysical or, possibly, earlier primordial sources. Such backgrounds will also be anisotropic as they would trace the distribution of the underlying sources and the anisotropy may also be detected. To this end, this thesis presents Gravitational Wave Background (GWB) map-making algorithms that use the cross-correlation of multiple data streams. The resulting maps are maximum likelihood representations of the GWB intensity on the sky. Two different methods are developed, depending on whether the detector timestreams may be considered uncorrelated or not. These are then tailored to two GW detector classes, ground-based interferometers and the Laser Interferometer Space Antenna (LISA), and tested on mock data as proof of concept. The mapping method optimised for the Laser Interferometer Gravitational-wave Observatories (LIGOs) is applied to LIGO open data sets, providing the most stringent upper limits to the GWB yet. These are obtained assuming a spectral model for the signal which weights the data and optimises the search. A novel, model-independent estimation scheme is also presented. The method applied to LISA provides the first full-likelihood map estimator from LISA data. The aim of this work is to deliver the framework of the LISA GWB map-making pipeline which will be progressively refined and updated until data become available.

Declaration

I hereby declare that the contents of this thesis is my own work, except where specific reference is made to the work of others. This includes all Figures and Tables.

Chapter 2 provides a literature review on the broad theme of stochastic gravitational-wave backgrounds. Specific references are provided throughout.

Chapter 3 presents the derivation of optimal map solutions in the context of gravitational-wave observations. These methods are developed in [1, 2] by this author.

Chapter 4 describes ground-based interferometers and presents the application of methods laid out in Chapter 3 to data from this class of detectors. The work presented here is mainly drawn from [1, 3, 4] by this author.

Chapter 5 describes the Laser Interferometer Space Antenna and presents the application of methods laid out in Chapter 3 to data from this detector. The work presented here is mainly drawn from [2] by this author.

Chapter 6 outlines future prospects for the methods developed in this thesis, providing conclusions to this body of work. Specific contributions by this author on the topics discussed in this chapter may be found in [5, 6].

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Arianna Renzini

September 2020

For Edward Tasker. Wish you were here.

Contents

List of figures **11**

1 Introduction **13**

2 Stochastic Gravitational-Wave Backgrounds **19**

2.1 Astrophysical Backgrounds 24

2.1.1 Spectral Properties of Astrophysical Backgrounds 27

2.2 Primordial Backgrounds 28

2.3 Anisotropies in Gravitational-Wave Backgrounds 31

2.4 Strain Signal and Stokes' Parameters 35

2.5 Observational Properties of the Stochastic Background 42

3 Maximum-likelihood map solutions **47**

3.1 Maps from Uncorrelated Data Streams 49

3.2 Maps from Correlated Data Streams 54

4 Mapping the GWB with Ground-based Detectors **59**

4.1 Detector response 62

4.2 Optimal map solutions: generation and testing 71

4.2.1 Data generation 74

4.2.2 Mapping tests 77

4.2.3 Polarisation maps 87

4.3 Maps from the Advanced LIGO O1 and O2 data runs 96

4.3.1 Noise Analysis 98

4.3.2 Model-dependent, broadband limits 103

4.3.3 Model-independent spectral limits 111

5 Mapping the GWB with LISA **117**

5.1 Detector response 119

5.2 Optimal map solutions for LISA 128

5.2.1 Data generation 129

5.2.2 Mapping tests 132

6 Conclusions and Future Prospects **141**

Bibliography **149**

List of figures

1.1 Gravitational-wave detector sensitivities and characteristic signals. 15

2.1 Cosmological GW signals. 30
2.2 Comparison between different astrophysical GWB models. 34
2.3 Time-domain properties of gravitational-wave backgrounds. 43

4.1 Schematic representation of a LIGO-like laser interferometer 60
4.2 Spherical harmonic expansion of the overlap function. 67
4.3 Instantaneous overlap function for ten different baselines. 69
4.4 Angular phase coverage of the 10 baseline case plotted in galactic coordinates and angular scale sensitivity of each baseline. 70
4.5 LIGO Livingston smoothed PSD for 60 second time segment. 75
4.6 Input and output maps for the high ANL cases. 80
4.7 Input and output maps for the low ANL cases. 81
4.8 Input and output maps for the Galaxy Poisson simulation. 82
4.9 Residual maps for high and low ANL. 83
4.10 Input vs. reconstructed power spectra for high and low ANL. 84
4.11 Evolution of the condition number of the beam-pattern matrix with integration time for different baseline combinations. 85
4.12 Map standard deviation for the LIGO-VIRGO, three baseline simulations with different ANL. 85
4.13 Noise map for the ten baseline case after 24 hours of integration. 86
4.14 Input and output polarised I, V maps. 91
4.15 Input and output polarised I, Q, U maps. 93

4.16 Input and output polarised I, V, Q, U maps. 94
4.17 Distribution of noise spectra (PSD) for the Advanced LIGO O1 and O2 data runs.101
4.18 One and two dimensional distributions in the noise fitting parameters for the O1 and O2 data runs. 102
4.19 SNR and noise maps of the SGWB from LIGO data. 105
4.20 Evolution of the standard deviation of the maps with integration time. 106
4.21 Comparison of upper limit value of the GWB from aLIGO data to expected value from literature. 107
4.22 Upper limits on the GWB angular power spectrum from aLIGO data. 108
4.23 Upper limits per frequency bin on the monopole signal of the GWB from aLIGO data. 112
4.24 SNR and noise maps for the frequency bins used in the model-independent limits. 113
4.25 SNR and noise maps for the frequency bins used in the model-independent limits. 114

5.1 Illustration of the TDI configurations for LISA. 120
5.2 Normalised autocorrelated response of LISA TDI channel X. 123
5.3 LISA sky-integrated auto-correlated and cross-correlated responses. 124
5.4 LISA relative sky sensitivity as a function of frequency. 127
5.5 Data power sample from the autocorrelated X channel. 131
5.6 Example input map and reconstruction at high SNR with LISA. 133
5.7 Average reconstructed power spectra and transfer functions with LISA. 134
5.8 Average reconstructed power spectra and transfer functions with LISA. 134
5.9 Input and reconstructed output maps in a low SNR case with LISA. 135
5.10 Sky distribution of the noise in the SSB frame. 136
5.11 Sky inputs and reconstructed outputs for a Milky Way - like GWB distribution with LISA. 137

6.1 Future gravitational-wave detector sensitivities and characteristic signals. . . . 142

Chapter 1

Introduction

The first direct measurement of a gravitational wave (GW) announced in early 2016 by the Laser Interferometer Gravitational-wave Observatory (LIGO) collaboration inaugurated the detection era for GW astronomy, opening an uncharted path to the unknowns of our Universe. This novel observational outlook allows us to tackle astrophysical and cosmological puzzles independently from radiation, which up to now has been the only source of information in the sky. Cosmology, as the study of the Universe as a whole and of its origins, calls for statistical methods and measurements; it is clear then that once existing and planned GW observatories reach sufficient sensitivity to detect a population of signals, we will enter the era of GW cosmology. Extracting cosmological information from GW detections requires the introduction of versatile and untested methods at all levels of analysis, building on the tools we use today in electromagnetic-wave cosmology. It is in this context that this thesis finds its ultimate motivation.

There are several proposed ways to detect gravitational waves, however the only confirmed detection to date has been using multiple ground-based, kilometre-scale interferometers. Eleven direct detections of GW signals have been confirmed thus far [7], all by the Laser Interferometer Gravitational-wave observatory (LIGO) and Virgo collaboration (LVC) which counts three detectors: the two Advanced LIGO (aLIGO) observatories [8] in Hanford (WA) and Livingston (LA), United States, and the Virgo detector in Pisa, Italy [9]. Many more detections are being announced throughout the current Advanced LIGO observing run [10, 11]. These

detectors monitor the frequency range between $10 - 10^3$ Hz, where the primary expected GW signals are compact binary inspiral and merger signatures. All detected signals have been interpreted as transient signals emitted during the final moments of binary stellar-mass black holes and neutron stars coalescing. These signals depend on the distance, masses, spins, separation and coalescence stage of the two binary objects. Such events are identified in the timestream via signal-to-noise threshold triggers and then classified through matched-filtering against predicted signal templates by varying the aforementioned signal parameters to obtain the best fit. As such they represent a tiny fraction of the data collected by the network of detectors. Whilst it is very difficult to identify signals of individual events below the detection threshold, given enough sensitivity and integration, it is expected that these events will eventually lead to the statistical detection of a *gravitational-wave background* (GWB) made up of the superposition of many undetected events. Indeed, all the sub-threshold compact binary coalescence events in the Universe will contribute to an astrophysical background, along with other secondary astrophysical GW sources [see e.g. 12]. Cosmological and primordial backgrounds may also be present, albeit at levels far below astrophysical ones at accessible frequencies in the foreseeable future [see e.g. 13].

The detection of an astrophysical background will be challenging with LIGO class detectors although the number of such detectors is expected to double over the next decade, and the LIGO instruments are expected to reach *design* sensitivity within a few years [14]. This, together with the guarantee of longer integration times may bring the overall, integrated sensitivity to a point where a background detection may be feasible. Irrespective of this, pulsar timing arrays (PTAs) are expected to yield direct detection of a stochastic background of continuous GWs from massive black-hole binaries in the nano-hertz frequency band over the next 10 years [15, 16]. PTAs are catalogues of pulsars monitored specifically to detect gravitational waves and other elusive perturbations in the Universe. Pulsars are highly magnetized, rotating neutron stars which emit a collimated beam of electromagnetic radiation that may be observed at periodic intervals as the star rotates in a "lighthouse" fashion. Pulsars are typically categorised based on the duration of their pulse period, which ranges from seconds to milliseconds, making very stable clocks such that GW signatures may

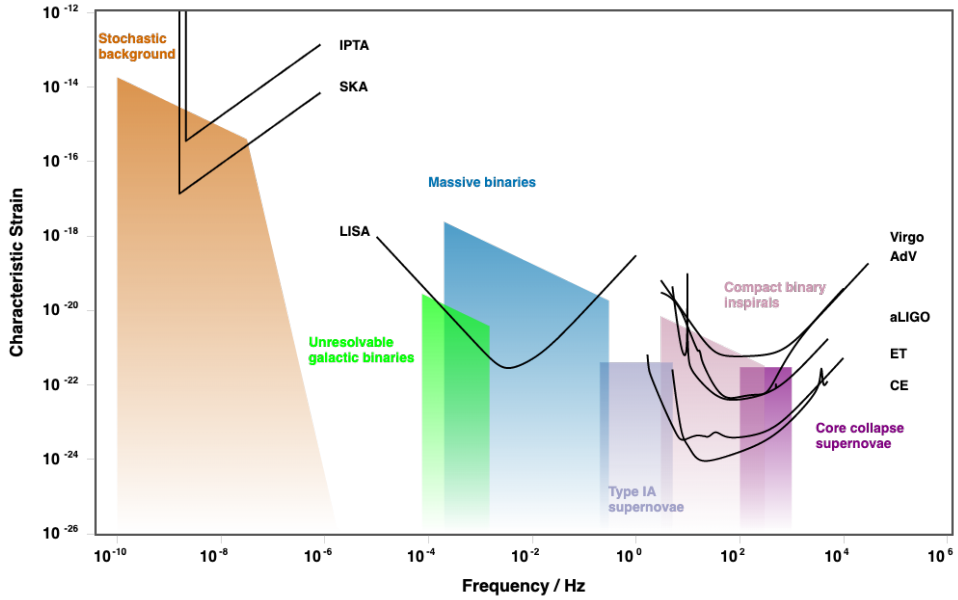


Figure 1.1: Gravitational-wave detector sensitivity curves and characteristic signal estimates in their relative frequency bands from [19]. The characteristic GW *strain*, i.e. the fractional length deformation imprinted on the spacetime metric due to a passing wave, is plotted as a function of frequency for typical GW sources. The signals are plotted here as simple power-laws with characteristic amplitudes corresponding to typical signal estimates in the literature. From left to right: the International Pulsar Timing Array (IPTA) sensitivity and the expected sensitivity of the Square Kilometre Array (SKA) to millisecond pulsar timing residuals; the LISA sensitivity curve; the Virgo, Advanced Virgo (AdV), Advanced Design LIGO (aLIGO), ET and CE curves all probing the higher end of the frequency range.

be found in the timing residuals of their pulse periods [see e.g. 17]. Specifically, PTAs are well suited to characterising a GWB of stochastic origin [18].

On longer time scales, the ESA Laser Interferometer Space Antenna (LISA) mission [20, 21] with an expected launch date in the late 2030s, is virtually guaranteed to reach the sensitivity required to detect a background of galactic and extra galactic sources. In fact, the presence of these backgrounds constitutes a challenge for the LISA analysis pipeline as the presence of a statistical background complicates the problem of signal and noise estimation. These problems are similar to ones encountered in Cosmic Microwave Background (CMB) data analysis [see e.g. 22] and the analogy is particularly close when considering the case of low frequency interferometric measurements of the CMB using coherent detectors in radio interferometry [see e.g. 23–25]. Beyond LISA, there are also planned *third generation* (3G) ground-based interferometers such as the Einstein Telescope (ET) [26] and Cosmic Explorer (CE) [27] which will probe the same frequency range as second generation detectors with a

highly increased sensitivity. These detectors are expected to resolve almost every compact binary coalescence in the observable Universe [28]. This provides a positive perspective on the possibility of measuring the background with this class of detectors, and motivates further specialistic studies. A summary plot of current and future GW detector sensitivity curves as a function of the frequency ranges probed may be seen in Figure 1.1, from [19]. This includes the typical predicted signals in each range.

Other than direct detection, for many years cosmologists have searched for signatures of gravitational waves in the Cosmic Microwave Background. These may be observed as a contribution to the density fluctuations in the primordial plasma which emitted the last scattering surface. At present these contributions are difficult to separate out, so the total observed fluctuations place an upper limit on the size of gravitational-wave fluctuations at recombination.

A detection and characterisation of both stochastic and primordial GWBs is an exciting prospect. The average value of stochastic backgrounds, together with any spectral dependence, will yield important constraints on the redshift distribution and nature of the source population. Beyond this, any detection of anisotropies in the backgrounds, particularly for extra-galactic sources, will constitute a biased tracer of the underlying matter distribution and will open a new *multi-messenger* window on studies of large scale structure. The detection of a cosmological GWB in particular is considered to be the ultimate challenge in cosmology as GWs are the only signal expected to reach us from before recombination and the consequent generation of the cosmological microwave background (CMB), providing unique insight on the origins of the Universe. The detection of an average gravitational wave density would constitute a groundbreaking result and any information on anisotropies could yield potential evidence of an inflationary epoch. Efforts to calculate the expected anisotropies in a variety of GWBs are ongoing [29–31], and lately there has been particular emphasis on understanding the time-domain characteristics of GWB signals as these will be crucial when analysing the detector timestream data [32, 33].

In light of this and of the expected growth in detector numbers and sensitivity we have developed a wide-scoped program aimed at building generalised, maximum-likelihood, grav-

itational wave map-making algorithms for the purpose of mapping *incoherent* and *stochastic* backgrounds. We have used a bottom-up approach to deal with data and have worked directly in the sky–pixel domain building on the experience of similar efforts in CMB analysis, focusing on the present ground-based interferometer network and on the future LISA space mission. Algorithms to map the gravitational wave density and its directional dependence in the sky frame have been considered in a number of works in the literature including [34–43], which have strongly influenced this work. Actual constraints on the background energy density of gravitational waves $\Omega_{\rm GW}$ and its anisotropies have been published by the LIGO collaboration in [44–48]. These have employed a number of methods such as sidereal folding of the data [see also 49], and radiometer or spherical harmonic domain methods [40] using the frame imposed by the LIGO baseline to reduce the sky rotation to a single phase modulation. More generalised methods have been explored by [35, 43]. These have included a well-developed "radiometer" method aimed at mapping the sky signal of individually occurring sources along with algorithms that constrain directly the spherical harmonic coefficients of the GWB [38, 40]. The work presented here focuses on solving for maximum-likelihood maps in the *pixel* domain directly instead, as this is found to be more natural for GWs and minimises the information loss.

Much work has also been done to assess the ability of LISA to detect and characterise GWBs [35, 36, 50–53], and map-making with space-based detectors is discussed in [37, 54]. These approaches reduce the mapping problem to a least-squares solution and solve for the spherical harmonic coefficients of the signal by looking at the individual cross-correlations between data-streams. Conversely, the method developed here solves the full-likelihood problem, in the pixel domain, giving the first optimal framework for map-making with LISA.

Additionally, it has been proposed to solve for phase-coherent backgrounds [42], which is effectively not possible due to both the characteristics of the signal and the nature of the measurement. In fact, any GWB that may be observed is phase-incoherent: in the case of astrophysical backgrounds and most cosmological ones, the total signal will necessarily be the accumulation of unsynchronised bursts and continuous waves with random phases; only a truly primordial background generated by modes re-entering the horizon simultaneously after inflation will be originally coherent, but it will suffer decoherence in its propagation

through the Universe [55].

This ***thesis*** is organised as follows: Chapter 2 offers an overview of stochastic GWBs and their characteristics, with particular attention to observational properties which will be crucial in determining optimal detection techniques. Chapter 3 describes the maximum likelihood solutions developed and applied in this thesis. Chapter 4 reviews the response of ground-based interferometers and shows the application of the maximum likelihood mapping method optimised for LIGO-like detectors both on simulated and real data from the LIGO detectors. Chapter 5 reviews the response of LISA and shows the application of the maximum likelihood mapping method optimised for the latter on mock data. Chapter 6 closes this thesis with the future prospects of this work, presenting the opportunities and challenges that lay ahead in gravitational-wave map-making.

Chapter 2

Stochastic Gravitational-Wave Backgrounds

All direct detections of gravitational waves to date have been coherent measurements of resolved waveforms in the detector timestream data which may be traced back to point-like sources in the sky. These make up a tiny fraction of the gravitational-wave sky, and the vast collection of unresolved signals add up incoherently and correspond to multiple point sources or extended sources fall under the category of gravitational-wave backgrounds (GWBs). A variety of different backgrounds is expected given the range of GW sources in the Universe, however regardless of their origin most of these share the feature of being *stochastic*, as they may be described by stochastic amplitudes, and are hence referred to as stochastic gravitational-wave backgrounds (SGWBs). As will be shown in this Chapter, some backgrounds are stochastic by generation process, whereas others are stochastic due to the characteristics or limitations of the specific detector used to observe them. An SGWB of this latter nature is by definition at the threshold of detection, making it effectively a detector-dependent observable. Determining whether a signal is a background and also whether it is a truly stochastic quantity is then an iterative process, where first all signals present in the detector timestream undergo signal-to-noise ratio (SNR) estimation which determines their resolvability, and then the background is estimated as the cumulative sum of only the unresolved and/or sub-threshold signals. In practice each GW detector will be able to access

qualitatively different backgrounds, depending on the noise levels and frequency ranges probed.

This Chapter reviews the expected sources of GWB typically considered in the literature and outlines the different properties of the backgrounds in relation to the detectors that measure GWs today and those that will probe them in the foreseeable future. Particular attention is given to the possible anisotropic features of these backgrounds which may be investigated with the mapping algorithms presented in this thesis. To start, it is necessary to provide the theoretical framework assumed when describing the generation and propagation of gravitational waves in the Universe.

Structure and Energy Content of the Universe

The geometry of our homogenous, isotropic, and expanding Universe is well-described by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric,

$$ds^2 = -dt^2 + a^2(t)\left(\frac{dr^2}{(1-Kr^2)} + r^2 dS^2\right). \tag{2.1}$$

Here t is coordinate time and all the time dependence of the spatial part of the metric is contained in the scale factor $a(t)$. The curvature constant K can assume the discrete values $\{-1, 0, 1\}$ which correspond to an open, flat, and closed universe respectively; for our Universe, K has been measured to be consistent with 0, hence to very good approximation it may be considered flat. In this case, the comoving distance between two objects embedded in the metric is equal to the comoving coordinate distance, r, and the *proper* distance evolves with time and is $d(t) = a(t)\,r$ at time t. These two measures play a key role in understanding the *location* of events and objects and thus the layout of structure in the Universe.

The FLRW metric is an exact solution to the Einstein field equations[1] [56],

$$R_{\mu\nu} + \frac{1}{2} R\, g_{\mu\nu} = 8\pi G T_{\mu\nu}\,, \tag{2.2}$$

where $T_{\mu\nu}$ is the matter energy–momentum tensor and $G = 6.67 \times 10^{-11}$ m^3 kg^{-1} s^{-1} is the universal gravitational constant. These are a set of uniquely powerful equations which

[1] Note: these ignore the contribution of a possible cosmological constant.

connect the energy content of the Universe, $T_{\mu\nu}$, to the geometry of the Universe itself and thus set the laws of information propagation in Space. Note that the units are chosen such that the speed of light $c = 1$. Simplifying the description of the energy content, i.e. the *matter* content, to that of a perfect fluid, the energy–momentum tensor may be chosen as $T_{\mu\nu} = \mathrm{diag}(-\rho, p, p, p)$, where p is the pressure and ρ is the energy density of matter. Writing Equation 2.1 in metric form, calculating the relevant Riemann and Ricci tensor components and plugging them into the Einstein equations returns the Friedmann equations,

$$\left(\frac{\dot{a}(t)}{a(t)}\right)^2 \equiv H^2(t) = \frac{8\pi G}{3}\rho + \frac{K}{a(t)^2}\,, \tag{2.3}$$

$$\frac{\ddot{a}(t)}{a(t)} = -\frac{4\pi G}{3}(\rho + 3p)\,, \tag{2.4}$$

which are the cornerstone of modern cosmology. $H(t)$ is the *Hubble rate* which quantifies the expansion of the Universe at time t. For our flat Universe the matter energy density today is then

$$\rho_c = \frac{3H_0^2 c^2}{8\pi G}\,, \tag{2.5}$$

which is referred to as the *critical energy density*, i.e. the energy density required to close the Universe today. Here $H_0 = 70\ \mathrm{km\ s^{-1}\ Mpc^{-1}}$ is the Hubble rate today[2]. The density parameter

$$\Omega_\mathrm{M} = \frac{\rho}{\rho_c} \tag{2.6}$$

may be defined, where M stands for matter. Assuming the vast majority of matter behaves as pressureless dust, the conservation of the energy–momentum tensor implies $\rho \sim \Omega_\mathrm{M} \sim a^{-3}$. Furthermore, it is useful to introduce the *redshift parameter*

$$z(t) = \frac{a_0}{a(t)} - 1\,. \tag{2.7}$$

Equation (2.3) may then be recast in the useful form

$$H(z) = H_0\left(\Omega_\mathrm{M}(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda\right)^{1/2}, \tag{2.8}$$

[2]Note that the subscript "$_0$" denotes quantities measured today.

where the density parameter for the cosmological constant Λ has been introduced for completion. Essentially, Equation (2.8) links together the Universe's expansion with its content, with respect to the observer at $z = 0$, and will invariably be part of any cosmological path or length calculation in our Universe, from our perspective.

At this stage, one may ask: where do gravitational waves fit in? If gravitational waves had a strong impact in the content and evolution of the Universe as a whole, these would certainly appear in the equations above, however experimental evidence shows otherwise. It is then reasonable to consider GWs as perturbations in the Universe, which do not influence the bigger picture but obey the propagation rules the structure and energy content of the Universe have set out. Nevertheless, they carry energy which may be characterised by the dimensionless energy density $\Omega_{\rm GW}$, similarly to the other components of the Universe above. $\Omega_{\rm GW}$ is then the ensemble energy of GWs, which is essentially that of the GWB, minus the infinitesimal fraction carried by resolved sources. The spectral dependence of this measure is given by the physical energy density of the GWB per logarithmic frequency interval as [57]

$$\Omega_{\rm GW}(f) = \frac{1}{\rho_c}\frac{d\rho_{\rm GW}}{d\ln f}. \tag{2.9}$$

Note that the frequency f here is the *measured* frequency, assuming an observer and the intervention of an instrument which mediates the true signal and the measurement. Beyond the isotropic value of the GWB over the entire sky it is possible to include any anisotropy by adding a directional dependence $\Omega_{\rm GW}(f) \to \Omega_{\rm GW}(f, \hat{\boldsymbol{n}})$ where $\hat{\boldsymbol{n}}$ is the unit vector of a line-of-sight.

It is now possible to recover the cosmological information contained in the $\Omega_{\rm GW}$ parameter, which is effectively the cumulative GW energy history weighted per redshifted frequency bin, as presented in [58]. This may be seen by noting that homogeneity and isotropy throughout the universe imply

$$\rho_{\rm GW} = \int_0^\infty \frac{df}{f}\int_0^\infty dz \frac{N(z)}{1+z} f_{\rm r}\frac{dE_{\rm GW}}{df_{\rm r}}, \tag{2.10}$$

where z is redshift, $N(z)$ is the number of GW emitters as a function of redshift, and $f_{\rm r}$ is the emission frequency of gravitational waves in the rest frame of the emitter, $f_{\rm r} = f(1+z)$. $dE_{\rm GW}/df_{\rm r}$ is then the spectral energy density of the source population. Equation (2.9) may

simply be re-written as

$$\rho_{\rm GW} = \int_0^\infty \frac{df}{f} \rho_{\rm c}\, \Omega_{\rm GW}(f)\,, \tag{2.11}$$

hence equating these frequency by frequency yields

$$\Omega_{\rm GW}(f) = \frac{1}{c^2 \rho_c} \int_0^\infty dz \frac{N(z)}{1+z} \left[f_{\rm r} \frac{dE_{\rm GW}}{df_{\rm r}} \right]_{f_{\rm r}=f(1+z)}. \tag{2.12}$$

The interpretation of Equation (2.12) is clear: the fractional energy density of GWs per measured logarithmic frequency intervals, at redshift zero, is the integral over the cosmic history of the comoving number density of the sources multiplied by the emitted energy fraction of each source, in the appropriately redshifted frequency bin. This draws a simple connection between the observable $\Omega_{\rm GW}$ and the individual source event rate as a function of redshift $N(z)$. In particular, in the case of several different types of sources the latter may be substituted by a sum over all contributors,

$$N(z) \left[f_{\rm r} \frac{dE_{\rm GW}}{df_{\rm r}} \right]_{f_{\rm r}=f(1+z)} \longrightarrow \sum_i N_i(z) \left[f_{\rm r} \frac{dE_{\rm GW,i}}{df_{\rm r}} \right]_{f_{\rm r}}. \tag{2.13}$$

GWBs fall broadly into two distinct types based on the underlying generation mechanism. The first is a superposition of the signal of astrophysical sources, both galactic and extra-galactic, that are not individually detected or resolved. This type of background will thus be unresolved and incoherent, in that the temporal phase of the signal is expected to be random, and is expected to be *stochastic* in the limit where the number of sources N is very high. The second category is a cosmological or primordial GWB. This category covers GWBs generated by the evolution of vacuum fluctuations during an inflationary epoch that end up as super-horizon tensor modes at the end of inflation [59, 60]. It also includes GWBs generated by non-linear phenomena at early times such as phase transitions or via emission by topological defects [61–63]. Assuming homogeneity and isotropy, primordial backgrounds are also expected to be stochastic due to the generation process, the propagation effects, or both [64, 65]. In particular, stochasticity here implies there is no information in the temporal phase of the background [55]. The randomness in the emission parameters for an astrophysical background implies also that these are in general expected to be unpolarised, as an equal distribution in the polarisation modes will cancel out any local asymmetry. This

applies also to most primordial backgrounds, with some exceptions as explained below. The main distinction between different background components is then in the spectral dependence of the source energy density in Equation (2.12).

In this diverse picture of the GW Universe the observer plays an essential role, in that all measurements are necessarily performed at redshift $z = 0$. This will necessarily lead to non-negligible selection effects, and implies that *qualitatively* different backgrounds will be contributing from different redshift shells. From an outside, "universal" perspective however, at least for an astrophysical background it is possible to imagine three distinct scales where the GW signal will be qualitatively similar: the local, astrophysical scale, where individual source parameters will play a defining role; the galactic scale, where each galaxy may be described by its effective GW luminosity, linked to galactic parameters such as mass and average metallicity; finally, the cosmological scale, where each galaxy is a point-like source comoving with the cosmic flow. Information from *all* these scales will define a *detected* signal of the GWB, as will emerge in the discussion throughout this Chapter.

The rest of this Chapter provides an overview of the above properties of stochastic backgrounds typically considered in the literature and sets the mathematical framework which will be used throughout this thesis. First, a broad characterisation of astrophysical and primordial backgrounds is given, focusing on the similarities and differences the various types of backgrounds present so as to provide some simple tools which will be used in the Chapters that follow. Then, a detailed mathematical description of the strain signal of GWBs is provided, which lies at the heart of GWB modelling and data analysis.

2.1 Astrophysical Backgrounds

Astrophysical GWBs are the collection of all gravitational waves generated by astrophysical processes which are either individual sub-threshold signals or are so numerous that they add up incoherently and form a continuous signal in the timestream. These processes are typically the inspiral and merger of compact binary systems including black hole binaries, neutron star binaries, white dwarf binaries and systems counting a mixed pair of these objects. Black hole binaries in particular are a vast category of sources, as the mass of each black

hole in the binary spans several orders of magnitude, from scales considered *massive* ($M_{\rm BH} \sim 10^5 - 10^{10}\mathscr{M}_\odot$), *intermediate* ($M_{\rm BH} \sim 10^2 - 10^5\mathscr{M}_\odot$), and *stellar mass* ($M_{\rm BH} \sim 5 - 10^2\mathscr{M}_\odot$), above the Tolman–Oppenheimer–Volkoff limit [66]. There are binaries that have a mass ratio close to one, where the two black holes in the binary fall within the same mass category, and there are also so-called extreme mass ratio (EMR) binaries where the two objects belong to different categories. These include extreme mixed binaries such as massive black hole–neutron star pairs. Each compact binary's life is divided in two main stages: the inspiral phase, during which the emitted GW signal has a simple cosine description and the energy emission increases very gradually as the binary's orbit slowly shrinks; and the final coalescence, which is marked by an intense GW burst or *chirp* signalling an abrupt release of gravitational energy and is typically modelled numerically. Given these descriptive differences, these two signals are typically treated as different types of events.

Other astrophysical sources of gravitational waves include asymmetric supernovae explosions and the rotation of neutron stars which are not spherically symmetric and may produce tri-axial emission [12]. These are however minor sources of GWs compared to compact binaries and are also considerably harder to model, hence the discussion that follows is mainly focused on the latter; nevertheless, many equations presented are general and may be applied to these secondary sources as well.

The $\Omega_{\rm GW}$ from astrophysical events is a compact measure of the history of the galaxy- and star-forming universe. To see this in more detail, following [58], let N be the number of events of a particular signal type in a comoving volume. This may be decomposed in the event rate $\dot{N}$ multiplied by the cosmic time slice dt_r/dz,

$$N(z) = \dot{N}\frac{dt_r}{dz} \equiv \frac{\dot{N}}{(1+z)H(z)}\,. \tag{2.14}$$

Here $\Omega_{\rm M}$, Ω_k, and Ω_Λ are the matter, curvature and cosmological constant density parameters respectively. Then as in [67] the fractional energy density may be re-written as

$$\Omega_{\rm GW} = \frac{f}{c^2\rho_{\rm c}}\int_0^{z_{\rm max}} dz\,\frac{\dot{N}(z)}{(1+z)H(z)}\left\langle \frac{dE_{\rm GW}}{df_{\rm r}}\bigg|_{f_{\rm r}=f(1+z)}\right\rangle, \tag{2.15}$$

where the angle brackets indicate an averaging over the particular source population parameters θ, such that

$$\left\langle \frac{dE_{\rm GW}}{df_{\rm r}} \right\rangle = \int d\theta\, p(\theta) \frac{dE_{\rm GW}(\theta; f_{\rm r})}{df_{\rm r}}, \tag{2.16}$$

and $p(\theta)$ is the probability distribution of the source parameters[3]. In general, one would expect the event rate $\dot{N}$ to locally increase with redshift. The event rate for stellar mass compact binary coalescences, for example, has been found to increase with redshift when analysing the direct detections reported by LIGO [68]. However, the rate can not increase monotonically out to infinity and must have a turn-around point at some peak redshift $\tilde{z}$, and then decay to zero as star formation ceases in the early Universe. For a general event rate, this may be described with a phenomenological model [69] as

$$\dot{N} = \mathscr{C}(\alpha, \beta, \hat{z}) \frac{\dot{N}_0 (1+z)^{\alpha}}{1 + \left(\frac{1+z}{1+\hat{z}}\right)^{\alpha+\beta}}, \tag{2.17}$$

such that $\dot{N} \propto (1+z)^{\alpha}$ at low redshift, reaches a maximum at $z \sim \hat{z}$, then $\dot{N} \propto (1+z)^{-\beta}$ at high redshift and $\mathscr{C}$ is the normalisation constant $\mathscr{C}(\alpha, \beta, \hat{z}) = 1 + (1+\hat{z})^{-(\alpha+\beta)}$ which ensures $\dot{N}(z=0) = \dot{N}_0$. Determining $\hat{z}, \alpha$ and β would give new insight on binary formation history and consequently star formation history as well. In practice, however, GW detectors inevitably probe the redshift-integrated observable $\Omega_{GW}(f)$ and the redshift history may only be reconstructed with careful modelling and the cross-correlation with independent measurements, such as in the case presented in [67]. Furthermore, it is important to note that the amplitude of the GW signal is directly related to the mass and density of the compact objects merging, which implies for example that binary black hole signals will have a more intense characteristic signal than binary neutron stars, and so forth. This may then be linked to the detection threshold via the SNR of each signal, which will set the maximum expected redshift depth at which each type of signal may be observed by the detectors probing that frequency range. A good example of this is the background generated by double white dwarfs, which is expected to be observable by the LISA detector only at very low redshift [70] such that LISA will only be sensitive to DWDs from the Milky Way and, perhaps, to those in the local

[3] In the case of a stochastic background one would expect the parameter space to be fully sampled by the population, however selection effects imposed by the detector may break this limit.

group. These selection effects need to be taken into account when attempting to correctly reconstruct the redshift distribution of events.

2.1.1 Spectral Properties of Astrophysical Backgrounds

It is possible to investigate the expected spectral dependence of $\Omega_{\rm GW}$ induced by the integrands in Equation (2.12) by plugging in different population models. This is particularly useful in the case of an astrophysical background as the behaviour of N and $E_{\rm GW}$ will be directly linked to the star formation history of the Universe, and astrophysical GWs from different source populations will contribute to $\Omega_{\rm GW}$ in different amounts at different times. This may lead to a way to differentiate between astrophysical backgrounds and also test different stellar and galaxy evolution models [71].

An example of astrophysical background with a simple expected spectral dependence is that of massive black hole binaries in the slow, adiabatic inspiral phase [72]. In this case it is sufficient to model the GW radiation at leading quadrupole order, hence [73]

$$\frac{dE_{\rm GW}}{d\ln f_{\rm r}} = \frac{\pi^{2/3}}{3}\mathscr{M}^{5/3} f_{\rm r}^{2/3} \tag{2.18}$$

where $\mathscr{M}$ is the *chirp mass* as defined in [73],

$$\mathscr{M} = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}} \tag{2.19}$$

where m_1 and m_2 are the two black hole masses. The comoving number density will depend on the chirp masses in the binary distribution, and may be re-written as an integral over $\mathscr{M}$,

$$N = \int_{\mathscr{M}_{\rm min}}^{\mathscr{M}_{\rm max}} \frac{dN}{d\mathscr{M}} d\mathscr{M}\,, \tag{2.20}$$

where the integration extrema will depend on the source population characteristics. Rearranging the terms in Equation (2.15) one obtains

$$\Omega_{\rm GW}^{\rm CB} \propto f^{2/3} \int_0^{z_{\rm max}} dz \frac{1}{(1+z)^{1/3}} \int_{\mathscr{M}_{\rm min}}^{\mathscr{M}_{\rm max}} d\mathscr{M} \frac{dN}{d\mathscr{M}} \mathscr{M}^{5/3}\,, \tag{2.21}$$

where the average over the population parameters is simply reduced to the integral over all possible chirp masses. Here the superscript CB labels compact binaries generically. This provides a simple spectral dependence for $\Omega_{\rm GW}^{\rm CB}$ which may be condensed as

$$\Omega_{\rm GW}^{\rm CB} = \Omega_{\rm GW}^{\rm CB}(f_0)\left(\frac{f}{f_0}\right)^{2/3}, \tag{2.22}$$

where $\Omega_{\rm GW}^{\rm CB}(f_0)$ is the appropriately normalised energy density calculated at a reference frequency f_0.

This simple model may be extended to all general inspiral-dominated GWBs as the same principles apply, and has been reprised all over the literature and used in almost every GWB search in the data available. The same principle, for example, has been used in the estimation of the expected background from stellar mass black hole and neutron star binaries in the LIGO detector frequency band [74], with some tweaks to account for the chirp signal from merging black holes which gives a substantially different energy output per frequency.

2.2 Primordial Backgrounds

There is a range of potential physical processes occurring in the primordial Universe which would lead to a rich variety of GW signals. The most noteworthy are reported here and are drawn from [13], which is a broad review of the subject.

First and foremost, all inflationary models include an irreducible GWB as a by-product of inflation. Any inflationary process leads to the stretching of quantum fluctuations which are parametrically amplified into classical density perturbations, producing in particular tensor perturbations which translate into gravitational waves through the Einstein equation. The amplification process causes the perturbation wavelengths to expand until they reach the size of the horizon, when they exit causal contact and their evolution remains frozen until the Universe's horizon returns to their scale in the post-inflationary era and they may re-enter. At time of re-entry, the GWB is formed by standing waves and their phases are correlated between modes with opposite momenta [75]. In most models the perturbations are adiabatic, Gaussian and approximately scale-invariant, leading to a similarly defined GWB. In this case,

the fractional energy density will be approximately frequency-independent,

$$\Omega_{\rm GW}^{\rm Infl.}(f) \approx \text{const.}\,, \tag{2.23}$$

and will be primarily characterised by its amplitude. The perturbations are also expected to imprint a pattern of B-modes in the polarization of the CMB [76], which has not been detected yet and is the aim of numerous CMB experiments. This non-detection imposes an upper bound on the inflationary GWB amplitude, setting it well below the reach of present or planned GW detectors. The prevalent opinion in the community is that precision measurements of the CMB are still the best chance to attain any evidence of inflation and discriminate between models.

Beyond the irreducible background, however, the inflationary era may also source a secondary GWB which in certain cases will rise above the primary component, often presenting strongly scale-dependent features or characteristic peaks in the power spectrum. Note that these other effects are referred to as secondary here because they are highly model dependent and often require fine-tuned conditions, not because they are the result of higher order contributions. The details of GW production during inflation can change significantly if additional fields other than the inflaton are present during inflation, as these may have interactions leading to strong particle production or may act as spectator fields and exhibit a sub-luminal propagation speed. The presence of additional symmetries in the inflationary sector would also modify GW emission, as these would lead to the breaking of space reparametrizations, allowing the graviton to acquire a mass. Similarly, alternative theories of gravity other than general relativity governing the inflationary era would have an impact on all aspects of primordial GWs. Finally, also the scalar density perturbations enhanced during inflation may source GWs, either at second order in perturbation theory or by acting as seeds for primordial black holes. All these possible inflationary scenarios are effectively extremely hard to probe, and GWs present a unique window into this epoch. See for example [77] for a discussion of inflationary backgrounds in the LISA band, and how a positive detection would revolutionise the current science of inflation.

Other than inflation, the other main sources of primordial backgrounds investigated in the literature are related to phase transitions which may have occurred as the primordial Universe

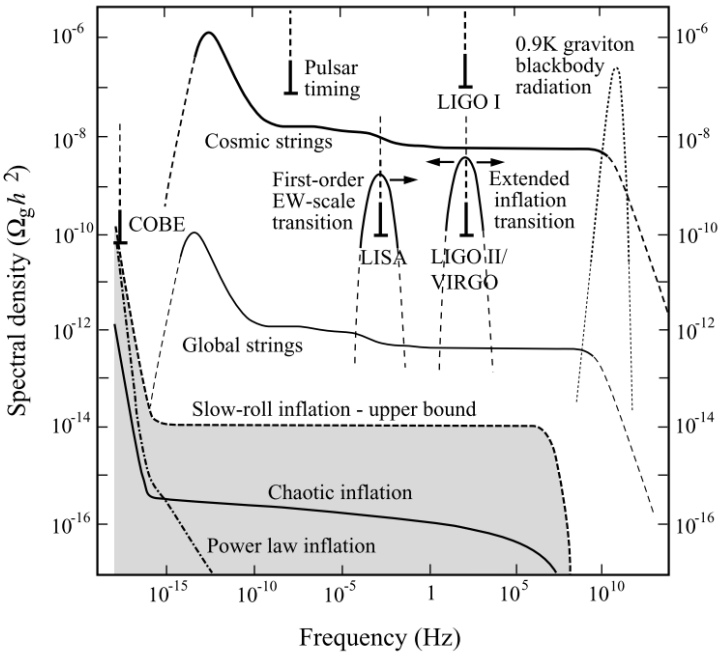


Figure 2.1: A summary plot with the expected amplitudes of cosmological signals discussed here and the expected detector sensitivities, from [62]. These include various inflationary scenarios, primordial first-order phase transitions and cosmic strings. The sensitivity of the following experiments/detection methods are shown: the Cosmic Background Explorer (COBE) CMB measurements; pulsar timing; LISA; LIGO and Virgo and their advanced (II) configurations.

cooled down in its post-inflationary adiabatic expansion. First order electro-weak phase transitions in particular would prompt the nucleation of true-vacuum bubbles, which when colliding would produce GWs [59, 61]. These are characterised by a power spectrum which presents a sharp peak at the frequency related to the energy scales of the transition. Phase transitions may also give rise to stable topological defects, the most prominent example of which are cosmic strings [62, 63], which radiate GWs when oscillating and forming loops. Topological defects also produce an irreducible background as the number density of defects must be preserved within a causal volume, providing a continuous source of GWs in the Einstein tensor. A summary of primordial GW signals and experimental sensitivities is shown in Figure 2.1, drawn from [62].

All GWBs emitted in the primordial Universe are considered stochastic under the *ergodic* assumption, which equates the ensemble average of an observable with its time and/or spatial average. This essentially implies that by observing large enough regions of the Universe today, or a given region for long enough time, one has access to many independent realisations of

the system. The assumption rests on two fundamental premises of cosmology, namely that the universe is statistically homogeneous and isotropic, such that the initial conditions of the GW-generating process may be considered the same at every point in space, and that the GW source fulfils causality, operating at a time when the typical size of a region of causal contact in the Universe was smaller than today. It can be shown [13] that the correlation scale today of a GW signal from the early Universe is tiny comparable to the present Hubble scale, verifying this requirement.

Generally speaking, stochasticity would also imply the lack of a preferred polarisation mode in the signal, as it assumes the configuration space is fully sampled in the observation set. While this is intuitive for astrophysical backgrounds, the argument for primordial GWs is more delicate and, while a background may be stochastic for the reasons described above, there may be specific formation channels which will violate parity and induce an asymmetry in the polarisation modes, effectively generating *circularly polarised* waves. This effect may be found in alternative theories of gravity, such as Chern-Simons gravity [78] or certain quantum gravity models [see e.g. 79]. The study of parity-violating theories of gravity is inspired by grand unification, as the electro-weak sector of the standard model is *chiral* and thus maximally violates parity [80, 81]. The specifics of these theories go beyond the scope of this work, however the possibility of a polarised GWB motivates the study of polarisation detection carried out in Chapter 4.2.3.

2.3 Anisotropies in Gravitational-Wave Backgrounds

GWB anisotropies will depend on both the inhomogeneities in the GW source mechanisms, whether of cosmological or astrophysical origin, and the GW propagation through the Universe, which will accumulate line-of-sight effects [29]. For astrophysical backgrounds, it is clear that the spatial distribution of GW events will be a biased tracer of the underlying light matter distribution, while primordial backgrounds will present anisotropies linked to the particular physical processes which occurred in the early Universe. To investigate these effects it is useful to define the fractional GW energy density per solid angle $d\hat{\boldsymbol{n}}$,

$$\Omega_{\rm GW}(f,\hat{\boldsymbol{n}}) = \frac{1}{c^2\rho_c}\frac{d\rho_{\rm GW}}{d\ln f\, d\hat{\boldsymbol{n}}}, \tag{2.24}$$

such that

$$\overline{\Omega}_{\rm GW}(f) = \frac{1}{4\pi}\int_{S^2} d\hat{\boldsymbol{n}}\,\Omega_{\rm GW}(f,\hat{\boldsymbol{n}})\,, \tag{2.25}$$

as in [31, 82]. $\overline{\Omega}_{\rm GW}(f)$ is the GWB monopole, where the bar is introduced to highlight it is an isotropic all-sky average over the directional fluctuations. These may be described by the GW density contrast,

$$\delta_{\rm GW}(f,\hat{\boldsymbol{n}}) = \frac{\Omega_{\rm GW}(f,\hat{\boldsymbol{n}}) - \overline{\Omega}_{\rm GW}(f)}{\overline{\Omega}_{\rm GW}(f)}\,, \tag{2.26}$$

in analogy with the formalism used to quantify the level of anisotropy in the CMB temperature [83] and in galaxy clustering [84]. This dimensionless parameter allows to describe GWB anisotropies as a fractional variation with the respect to the monopole value, providing an unambiguous metric to compare different results.

The density contrast may be decomposed into three distinct contributions related to the origin of anisotropies,

$$\delta_{\rm GW} = \delta^{\rm s}_{\rm GW} + \delta^{\rm los}_{\rm GW} + \mathscr{D}\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{v}}_{\rm obs}\,. \tag{2.27}$$

Here, $\delta^{\rm s}_{\rm GW}$ is the result of anisotropies at the GW source, for example due to the specific distribution of events for astrophysical backgrounds, while $\delta^{\rm los}_{\rm GW}$ is the anisotropy accumulated along the line-of-sight [29]. The latter will include terms similar to those encountered in CMB anisotropy calculations, such as the Sachs-Wolfe effect, due to the local curvature at emission, and the integrated Sachs-Wolfe effect, due to the curvature perturbations encountered along the line-of-sight [85]; a Doppler term due to the source's peculiar velocity; and higher order effects such as gravitational lensing and redshift-space distortions [86]. The final term is the dipole of magnitude $\mathscr{D}$ induced by the observer's peculiar velocity $\boldsymbol{v}_{\rm obs}$. All terms in Equation (2.27) have been the subject of intense study in recent years and both the astrophysical and primordial contributions have been modelled extensively. See for example [87, 88] for independent calculations of astrophysical background anisotropies, and [65, 82, 89] for examples of primordial ones. The propagation term $\delta^{\rm los}_{\rm GW}$ is often assumed to be negligible in the case of astrophysical backgrounds, or it is directly included as part of $\delta^{\rm s}_{\rm GW}$; in [90] it is expressly estimated and shown to account for at most $\sim 10\%$ of the total anisotropy. In the case of a truly primordial background dating back as far as or before the CMB however this

term will be considerably larger, and will be comparable to the intrinsic anisotropies at large scales as shown in [29].

A detailed review of these specific cases goes beyond the scope of this work, however it is useful to analyse the general scenario where the density contrast at the source behaves as a Gaussian random field, as this will be mostly assumed throughout. In this case, $\delta^{\rm s}_{\rm GW}$ may be uniquely characterised by its two point correlation function,

$$C_{\rm GW}(f,\cos\theta) = \left\langle \delta^{\rm s}_{\rm GW}(f,\hat{\boldsymbol{n}})\, \delta^{\rm s}_{\rm GW}(f,\hat{\boldsymbol{n}}')\right\rangle , \tag{2.28}$$

in analogy with CMB anisotropies, as the first moment $\left\langle \delta^{\rm s}_{\rm GW}\right\rangle$ is zero by definition, and all higher order moments will either vanish or may be expressed in terms of $C_{\rm GW}$ [82]. Here $\cos\theta = \hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{n}}'$ and the angular brackets represent an averaging over all $\hat{\boldsymbol{n}}$, $\hat{\boldsymbol{n}}'$ direction pairs of aperture θ. $C_{\rm GW}(f,\cos\theta)$ will then be different from zero where there are consistent spatial correlations on the sky associated to θ, at frequency f. To see this formally, it is common practice to expand the two point function into angular multipoles as

$$C_{\rm GW}(f,\cos\theta) = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} C^{\delta}_{\ell}(f) P_{\ell}(\cos\theta) , \tag{2.29}$$

where P_ℓ is the ℓth Legendre polynomial. It is then possible to invert this equation to obtain the angular power spectrum $C^{\delta}_{\ell}(f)$ for the GW energy density contrast,

$$C^{\delta}_{\ell}(f) = 2\pi \int_{-1}^{+1} d\cos\theta\, C_{\rm GW}(f,\cos\theta)\, P_{\ell}(\cos\theta) , \tag{2.30}$$

which may be estimated from observations. Note that the quantity $\ell(\ell+1)C_\ell/2\pi$ is roughly the contribution per logarithmic ℓ bin to the variance of $\delta^{\rm s}$, hence this will often be the quantity shown in plots quantifying the C^{δ}_{ℓ} power spectrum from different signal models and in GW data sets. Alternatively, it is possible to investigate the angular power spectrum of the $\Omega_{\rm GW}$ parameter itself, first by decomposing it into its spherical harmonic compoments,

$$\Omega_{\rm GW}(f,\hat{\boldsymbol{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{+\ell} \Omega^{\rm GW}_{\ell m}(f)\, Y_{\ell m}(\hat{\boldsymbol{n}}) , \tag{2.31}$$

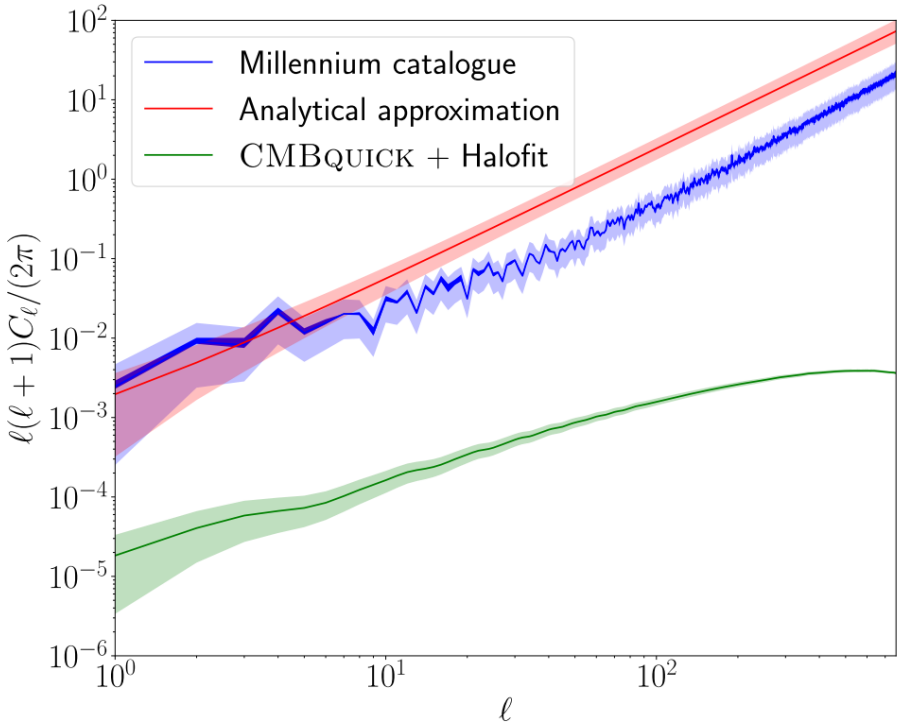


Figure 2.2: Different predictions of the astrophysical background angular power spectrum from [82]. Analytic and numerical predictions calculated in [82] in red and blue respectively, whereas in green the predictions presented in [87]. These predictions don't seem to agree and the main difference in the approaches is the treatment of highly non-linear low redshift effects.

and then by defining

$$C_\ell^\Omega = \frac{1}{2\ell+1}\sum_{m=-\ell}^{+\ell}\langle\Omega_{\ell m}^{\rm GW}(f)\,\Omega_{\ell m}^{\rm GW}(f)\rangle \tag{2.32}$$

as will be done in Chapters 4 and 5 from mock data and real data sets from the LIGO observatories. The information encapsulated in C_ℓ^Ω and C_ℓ^δ is the same and they are both dimensionless quantities, however note that the numerical scale of C_ℓ^Ω is roughly that of $\Omega_{\rm GW}^2$, whereas C_ℓ^δ shows the fractional contribution of each mode to the power spectrum with respect to the monopole $\overline{\Omega}_{\rm GW}$.

To quote the expected levels of anisotropy from calculations in the references mentioned above, see Figure 2.2 for astrophysical backgrounds. The difference in the estimates highlighted here stems probably from different treatments of the source effects at very low redshift, where the clustering at small scales induces non-linear effects in the field which the overall anisotropy level is very sensitive to. Cosmological backgrounds are expected to present CMB-like ansotropies of $\delta_{\rm GW}^{\rm s}\sim 10^{-5}$ [29]. Note that the *measurement* of any GWB angular power spectrum assumes averaging over the magnetic number m, which is necessary to

compare the correlations at different angular scales, however it is not trivial at an operational level. In practice, as will be shown in the next Chapters, in the case of GW backgrounds it is not possible to estimate an angular power spectrum directly from the data as the detector response function is highly inhomogenous and the m dependence is necessary to appropriately project the data onto the sky, estimating a map of the background before integrating over m.

Finally, note that to properly assess the sourced GW anisotropies from direct measurements it is necessary to subtract out the kinematic dipole term induced by the observer's peculiar velocity as seen in Equation (2.27). This may be done just as in CMB analysis, potentially using the independent estimate of $\boldsymbol{v}_{\rm obs}$ from CMB data themselves.

2.4 Strain Signal and Stokes' Parameters

Irregardless of their origin, stochastic background signals may be described in a common mathematical framework. It is worth laying this structure out, starting from the basic principles of general relativity, to set a solid foundation for the calculations that follow.

In the framework of linearised gravity [91], a *gravitational wave* propagating through empty space is defined as a linear perturbation of the background metric $g_{\mu\nu}$,

$$\tilde{g}_{\mu\nu} = g_{\mu\nu} + h_{\mu\nu}\,, \tag{2.33}$$

where $h_{\mu\nu}$ is *small* such that $h_{\mu\nu} = \epsilon H_{\mu\nu}$, $\epsilon \ll 1$ and is referred to as the GW strain tensor. To see why, it is necessary to linearise gravity, such that only terms $\mathcal{O}(\epsilon)$ are considered, and solve the Einstein equations in this context. The standard in linearised gravity is to fix the background to be flat Minkowski spacetime, $g_{\mu\nu} \equiv \eta_{\mu\nu} := \mathrm{diag}(-1,1,1,1)$. By inspection, one sees that $h_{\mu\nu}$ behaves as a rank two tensor, hence it transforms as

$$h'_{\mu\nu} = \Lambda^{\alpha}{}_{\mu}\Lambda^{\beta}{}_{\nu} h_{\alpha\beta} \tag{2.34}$$

under a Lorentz transformation Λ. To find the solution to the Einstein equations 2.2 it is necessary to calculate the Ricci tensor and scalar. First, the linearised Riemann tensor is

calculated as [92]

$$R^{\alpha}{}_{\mu\nu\beta} = \frac{1}{2}\left(\partial_\mu\partial_\nu h^{\alpha}{}_{\beta} + \partial^\alpha\partial_\beta h_{\mu\nu} - \partial_\mu\partial_\beta h^{\alpha}{}_{\nu} - \partial^\alpha\partial_\nu h_{\mu\beta}\right), \tag{2.35}$$

where the first index has been raised via the background metric. Contraction to obtain the Ricci tensor is then simple to carry out and gives

$$R_{\mu\nu} = -\frac{1}{2}\left(\partial_\mu\partial_\nu h + \Box h_{\mu\nu} - \partial_\mu\partial_\alpha h^{\alpha}{}_{\nu} - \partial_\nu\partial_\alpha h^{\alpha}{}_{\mu}\right), \tag{2.36}$$

where the d'Alambertian is defined as $\Box = \eta^{\alpha\beta}\partial_\alpha\partial_\beta$ and $h = \eta^{\alpha\beta}h_{\alpha\beta}$ is the trace of the strain tensor. Contracting again gives the first-order Ricci scalar,

$$R = \partial_\mu\partial_\nu h^{\nu\mu} - \Box h\,. \tag{2.37}$$

At this stage, it is possible to perform a coordinate transformation which will considerably simplify the calculation:

$$x'^{\mu} = x^{\mu} + \epsilon^{\mu}(x)\,, \tag{2.38}$$

$$\frac{\partial x^{\mu}}{\partial x'^{\nu}} = \delta^{\mu}{}_{\nu} - \partial_\nu\epsilon^{\mu} + \mathcal{O}(\epsilon^2)\,, \tag{2.39}$$

chosen such that $\partial_\nu\epsilon^\mu$ is of the order of h. Then the coordinate transformation applied on the strain tensor reduces to

$$h'_{\mu\nu} = h_{\mu\nu} - \partial_\nu\epsilon^{\mu} - \partial_\mu\epsilon^{\nu}\,, \tag{2.40}$$

similarly to the gauge transformation of the electromagnetic vector field A_μ: $A'_\mu = A_\mu + \partial_\mu\Lambda$, where Λ is a scalar field. In electromagnetism the gauge may be fixed by choosing Λ, the typical choice being the Lorenz gauge $\partial_\mu A^\mu = 0$, which significantly simplifies Maxwell's equations. Similarly, it is possible to fix the gauge of the GW strain; the common choice is the De Donder gauge condition

$$\partial_\mu h^{\mu}{}_{\nu} = \frac{1}{2}\partial_\nu h\,, \tag{2.41}$$

where the primes have been dropped for simplicity. Imposing the De Donder gauge and plugging the Ricci tensor and scalar into Equation 2.2, the Einstein equations in vacuum

become [56]

$$\Box\left(h_{\mu\nu}-\frac{1}{2}\eta_{\mu\nu}h\right)=-16\pi GT_{\mu\nu}=0\,, \tag{2.42}$$

where in the final equality the energy–momentum tensor is set to 0 to satisfy vacuum conditions. Equation (2.42) is simply the (relativistic) wave equation, hence the solution may be constructed as a linear superposition of *plane waves*,

$$h_{\mu\nu}=Ae_{\mu\nu}\,e^{i\left(k^{\alpha}x_{\alpha}\right)}\,, \tag{2.43}$$

where $e_{\mu\nu}$ is the normalised *polarisation* tensor and k_{α} is the wave four-vector. Inserting this solution in the equation above one gets the dispersion relation $k_{\alpha}k^{\alpha}=0$ which implies that the waves travel at the speed of light. Note that the De Donder gauge does not fix all the degrees of freedom (DOFs): the initial symmetric rank two tensor has 10 DOFs and the gauge conditions fix 4 of these leaving 6 residual DOFs. It can be shown that 4 of the 6 residual DOFs are not physical and it is possible to perform a residual transformation, which is equivalent to fixing 4 more conditions on the polarisation,

$$\begin{aligned} e_{0i}&=0\,,\\ e^{\alpha}{}_{\alpha}&=0\,, \end{aligned} \tag{2.44}$$

where i are the three spatial indices. The conditions imposed by Equations (2.41) and (2.44) are known as the *transverse-tracelesess* (TT) gauge, where the gravitational wave $h_{\mu\nu}$ has manifestly two physical DOFs.

By choosing the TT gauge, $h_{\mu\nu}$ may be reduced to a purely spatial perturbation, h_{ij}, which will carry two DOFs or equivalently two independent, spin-2, polarisation states. Generalising Equation (2.43) [see e.g. 35], the GW strain at time t and position vector $\boldsymbol{x}$ given by an infinite superposition of plane waves incoming from all directions

$$\hat{\boldsymbol{n}}=(\sin\theta\cos\phi,\sin\theta\sin\phi,\cos\theta) \tag{2.45}$$

in standard angular coordinates (θ, ϕ) on the 2-sphere may be written as

$$h_{ij}(t, \boldsymbol{x}) = \int_{-\infty}^{+\infty} df \int_{S^2} d\hat{\boldsymbol{n}} \sum_{P=+,\times} h_P(f, \hat{\boldsymbol{n}})\, e^P_{ij}(\hat{\boldsymbol{n}})\, e^{i2\pi f(\hat{\boldsymbol{n}}\cdot\boldsymbol{x}-t)}\,, \tag{2.46}$$

where the spatial wave vector is written explicitly as $\boldsymbol{k} = 2\pi f \hat{\boldsymbol{n}}$, P labels the polarisation states and h_P are the respective amplitudes assigned to each state. Choosing the *linear* polarisation basis $\{+, \times\}$, the orthogonal polarisation base tensors e^P may be chosen as

$$\boldsymbol{e}^+ = \boldsymbol{e}_\theta \otimes \boldsymbol{e}_\theta - \boldsymbol{e}_\phi \otimes \boldsymbol{e}_\phi\,, \tag{2.47}$$

$$\boldsymbol{e}^\times = \boldsymbol{e}_\theta \otimes \boldsymbol{e}_\phi + \boldsymbol{e}_\phi \otimes \boldsymbol{e}_\theta\,, \tag{2.48}$$

with

$$\boldsymbol{e}_\theta = (\cos\theta\cos\phi, \cos\theta\sin\phi, -\sin\theta)\,, \tag{2.49}$$

$$\boldsymbol{e}_\phi = (-\sin\phi, \cos\phi, 0)\,, \tag{2.50}$$

As defined in [93].

For a stochastic signal h_P represents a random amplitude, and here the working assumption is that the amplitudes are drawn from a Gaussian probability distribution. This is most likely a good approximation for a cosmological GWB, as explained above. For a stochastic GWB of astrophysical sources this assumption may break down but the central limit theorem will guarantee that the statistics approach that of a Gaussian random field if the signal is sourced by a sufficiently large number of independent events and that any high signal-to-noise outliers have been subtracted from the detector timestreams. Under the Gaussian assumption, the statistical properties of the amplitudes are then characterised solely by the second order moments $\langle h_P(f, \hat{\boldsymbol{n}}) h^\star_{P'}(f', \hat{\boldsymbol{n}}')\rangle$, which, assuming statistical homogeneity,

correspond to ensemble averages as [43]

$$\begin{pmatrix} \langle h_+ h'^{\star}_+ \rangle & \langle h_+ h'^{\star}_\times \rangle \\ \langle h_\times h'^{\star}_+ \rangle & \langle h_\times h'^{\star}_\times \rangle \end{pmatrix} = \frac{1}{2}\delta(\boldsymbol{n}-\boldsymbol{n'})\,\delta(f-f')\times \begin{pmatrix} I(f,\hat{\boldsymbol{n}})+Q(f,\hat{\boldsymbol{n}}) & U(f,\hat{\boldsymbol{n}})-iV(f,\hat{\boldsymbol{n}}) \\ U(f,\hat{\boldsymbol{n}})+iV(f,\hat{\boldsymbol{n}}) & I(f,\hat{\boldsymbol{n}})-Q(f,\hat{\boldsymbol{n}}) \end{pmatrix}, \tag{2.51}$$

where the Stokes parameters I, the intensity, Q and U, giving the linear polarisation, and V, the circular polarisation have been introduced. The four Stokes parameters completely describe the polarisation of the observed signal in analogy with electromagnetic Stokes parameters for the photon. The difference here is that whilst the electromagnetic Q and U Stokes parameters transform as spin-2 quantities with respect to rotations, their strain counterparts transform as spin-4, as shown below. In both cases the intensity I behaves as a scalar under rotations, while V transforms as a pseudo-scalar. The four Stokes parameters may then be written as combinations of the second order moments of h_P, specifically

$$I(f,\hat{\boldsymbol{n}}) = \langle h_+(f,\hat{\boldsymbol{n}})\,h^{\star}_+(f,\hat{\boldsymbol{n}})\rangle + \langle h_\times(f,\hat{\boldsymbol{n}})\,h^{\star}_\times(f,\hat{\boldsymbol{n}})\rangle\,, \tag{2.52}$$

$$Q(f,\hat{\boldsymbol{n}}) = \langle h_+(f,\hat{\boldsymbol{n}})\,h^{\star}_+(f,\hat{\boldsymbol{n}})\rangle - \langle h_\times(f,\hat{\boldsymbol{n}})\,h^{\star}_\times(f,\hat{\boldsymbol{n}})\rangle\,, \tag{2.53}$$

$$U(f,\hat{\boldsymbol{n}}) = -2\langle \mathbb{R}\big(h_+(f,\hat{\boldsymbol{n}})\,h^{\star}_\times(f,\hat{\boldsymbol{n}})\big)\rangle\,, \tag{2.54}$$

$$V(f,\hat{\boldsymbol{n}}) = -2\langle \mathbb{I}\big(h_+(f,\hat{\boldsymbol{n}})\,h^{\star}_\times(f,\hat{\boldsymbol{n}})\big)\rangle\,. \tag{2.55}$$

The GW strain may be equivalently decomposed in the circular polarisation basis $\{\boldsymbol{e}_R, \boldsymbol{e}_L\}$ where

$$\boldsymbol{e}_{R/L} = \frac{\boldsymbol{e}_+ \pm i\boldsymbol{e}_\times}{\sqrt{2}}\,, \tag{2.56}$$

and

$$h_{R/L} = \frac{h_+ \mp i h_\times}{\sqrt{2}}\,, \tag{2.57}$$

such that [94, 95]

$$I(f,\hat{\boldsymbol{n}}) = \langle h_R(f,\hat{\boldsymbol{n}})\, h_R^\star(f,\hat{\boldsymbol{n}})\rangle + \langle h_L(f,\hat{\boldsymbol{n}})\, h_L^\star(f,\hat{\boldsymbol{n}})\rangle\,, \tag{2.58}$$

$$Q(f,\hat{\boldsymbol{n}}) = 2\langle \mathbb{R}\big(h_R(f,\hat{\boldsymbol{n}})\, h_L^\star(f,\hat{\boldsymbol{n}})\big)\rangle\,, \tag{2.59}$$

$$U(f,\hat{\boldsymbol{n}}) = 2\langle \mathbb{I}\big(h_R(f,\hat{\boldsymbol{n}})\, h_L^\star(f,\hat{\boldsymbol{n}})\big)\rangle\,, \tag{2.60}$$

$$V(f,\hat{\boldsymbol{n}}) = \langle h_R(f,\hat{\boldsymbol{n}})\, h_R^\star(f,\hat{\boldsymbol{n}})\rangle - \langle h_L(f,\hat{\boldsymbol{n}})\, h_L^\star(f,\hat{\boldsymbol{n}})\rangle\,. \tag{2.61}$$

This notation is particularly useful when dealing with circularly polarised GWs, as the meaning of the V parameter is very clear.

These expressions for the Stokes' parameters provide an operative definition starting from the measurement of the amplitudes of the h fields. However, the spin properties of the fields are not immediately manifest and these will be an important factor to take into account in case of detection. To investigate these further, it is useful to use the R/L basis as it can be shown that h_L and h_R have spin ± 2 respectively. These may then be decomposed into spherical harmonics as [96]

$$h_{R/L}(f,\hat{\boldsymbol{n}}) = \sum_{\ell} \sum_{m=-\ell}^{\ell} a_{\ell m}^{R/L}(f)\; {}_{\pm 2}Y_{\ell m}^\star(\hat{\boldsymbol{n}})\,, \tag{2.62}$$

where $a_{\ell m}^{R/L}$ are the spherical harmonic coefficients of the R and L fields respectively and ${}_{\pm 2}Y_{\ell m}$ are the spin-2 weighted spherical harmonic functions. The spin-2 nature of the fields implies that $a_{\ell m}^{R/L} = 0$ for $\ell < 2$. In analogy with the E and B modes in electromagnetism, it is possible to construct *gradient* G and *curl* C modes from the spin-weighted coefficients as

$$a_{\ell m}^{G} = a_{\ell m}^{R} + a_{\ell m}^{L}\,, \tag{2.63}$$

$$a_{\ell m}^{C} = i\big(a_{\ell m}^{R} - a_{\ell m}^{L}\big)\,, \tag{2.64}$$

hence it is theoretically possible to estimate the power- and cross-spectra of the G and C modes of GW signals from measurements of $h_{R/L}$. However, given the coherent nature of GW detectors, it is not useful to measure stochastic quantities directly in the strain as any averaging of this quantity will yield zero. Then, assembling the Stokes' parameters as in Equations (2.58-2.61) with Equation (2.62), it appears clear that I and V yield spin-0 quantities

as the auto-correlations of the $h_{R/L}$ fields cancel the spin-2 dependence out, whereas these sum in the Q and U parameters yielding two pure spin-4 fields as [97]

$$(Q \pm iU)(f, \hat{\boldsymbol{n}}) = \sum_{\ell} \sum_{m=-\ell}^{\ell} b^{\pm 4}_{\ell m}(f)\, {}_{\pm 4}Y^{\star}_{\ell m}(\hat{\boldsymbol{n}}), \tag{2.65}$$

where ${}_{\pm 4}Y_{\ell m}$ are the spin-4 weighted spherical harmonic functions, and $b^{\pm 4}_{\ell m} = 0$ for $\ell < 4$ as above. These expressions will reveal useful when carrying out a full polarisation analysis in Chapter 4. For a more detailed discussion on GWB Stokes' parameters and their spin properties, see [96].

Finally, the intensity parameter I may be related to the fractional energy density parameter $\Omega_{\rm GW}(f)$ as [93]

$$\Omega_{\rm GW}(f) = \frac{4\pi^2 f^3}{\rho_c G} I(f). \tag{2.66}$$

This is the fundamental relation which allows to connect GWB observations to the cosmological implications of the background, and may easily be extended to a directional statement via Equation (2.25) as

$$\Omega_{\rm GW}(f, \hat{\boldsymbol{n}}) = \frac{4\pi^2 f^3}{\rho_c G} I(f, \hat{\boldsymbol{n}}). \tag{2.67}$$

As for the other Stokes' parameters, these do not contribute to $\Omega_{\rm GW}$ and will be equal to zero in the case of a truly stochastic, and thus unpolarised, background. This may be easily gleaned by their definitions, in Equations (2.53-2.55), as the second moments of the h_+ and $h_\times$ fields should be equal and the expectation value of any cross correlations between the two should vanish.

$\Omega_{\rm GW}$ will then inherit the spectral properties from the GWB intensity as per Equation (2.66). As explained in the previous sections, often the spectral dependence may be reduced to a power law, in which case it is common practice to write

$$\Omega_{\rm GW}(f) = \Omega_{\rm GW}(f_0) \left(\frac{f}{f_0}\right)^{\alpha}, \tag{2.68}$$

where α encapsulates the spectral shape of the signal. Then one may also assume, adding in also the directional dependence,

$$I(f,\hat{\boldsymbol{n}}) = I(f_0,\hat{\boldsymbol{n}})\left(\frac{f}{f_0}\right)^{\alpha-3}. \tag{2.69}$$

This will be particularly useful when map-making as it allows to integrate out the spectral dependence of the reconstructed signal and solve for an intensity sky map. Note that in the ensemble averaging required to recover the Stokes' parameters all the temporal phase information in the GWs is discarded; as argued above, no physical information may be recovered from the coherent phase component.

2.5 Observational Properties of the Stochastic Background

The observational properties of the SGWB such as its time-domain characteristics and selection effects imposed by the observation method merit special attention as these will be crucial towards developing an effective map-making method. For the sake of simplicity, this discussion is focused on astrophysical backgrounds where events, i.e. bursts and continuous waves, are clearly defined and are not extremely model-dependent. However, the following will in general apply to primordial backgrounds as well and has intuitive extension to backgrounds from topological defects and primordial black holes for example.

Three main regimes are identified in the literature based on the time-domain properties of GW signals, and are referred to as *shot noise*, where the GWs appear as isolated bursts in the timestream, *Gaussian*, where the timestream is packed with numerous overlapping signals and the central limit theorem applies, and *popcorn* which lies somewhere in between. A visualisation of these is provided in Figure 2.3 drawn from [12], where the distinguishing factor between regimes is parametrised by the *duty cycle*, which is the mean ratio between event duration and the time interval between consecutive events. This will also be equal to the average number of events present at the detector at a given observation time. To understand whether signals as in Figure 2.3 are truly unresolved, stochastic backgrounds it is necessary to clarify the meaning of *stochasticity* and *resolvability* of a signal, in an operational sense. A clear approach is assumed in [33], where it is said that a signal is *stochastic* if it is more reliable,

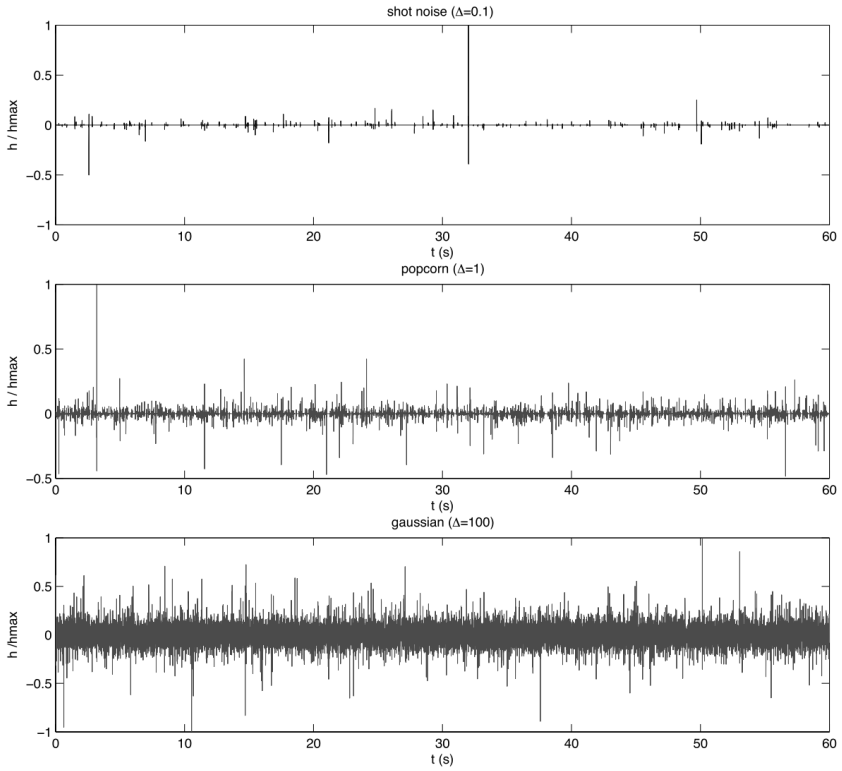


Figure 2.3: Examples of interferometer timestreams populated by GWBs with different time-domain properties, from [12]: shot noise (top panel), popcorn (middle panel), gaussian (bottom panel). The strain on the y-axis is normalised to show the fractional contribution of each burst/wavelet to the signal, while the Δ parameter is the *duty cycle*, i.e. the mean ratio between event duration and the time interval between consecutive events.

in the Bayesian sense of the word, to search for it in the data using a stochastic model for the waveform as opposed to a deterministic one. Furthermore, a signal is defined as *resolvable* if it can be decomposed into non-overlapping and individually detectable signals, in either the time or frequency domain. This highlights the importance of the temporal behaviour of the signal as a multitude of time-coincident waves will, probably, not be resolvable and will also probably be stochastic. However, a series of unresolvable, time-separated bursts in the data may or may not be stochastic, depending on what the best method is to confirm a detection based on these measurements. In the case of astrophysical GW signals, there is a continuous distribution of events where the highest SNR are expected to be resolved, and beneath a certain detector-imposed cut-off the events accumulate to form the background signal. The boundary between the resolved and unresolved population is not clear-cut, and a few bright signals stand out above the confusion background. When these are resolvable, deterministic signals, they should be subtracted from the data, leaving a residual background

which may be analysed with stochastic tools. In [33] the authors present simulations designed in order to assess whether different signals are better suited to stochastic or deterministic modelling. They find, similarly to [98] and somewhat unsurprisingly, that deterministic signal models are favoured whenever the number of sources contributing to the background is small and stochastic models are preferred for larger source numbers. The boundary between regimes however is not well defined, and hybrid models comprised of a deterministic signal model for individual bright sources and a Gaussian-stochastic signal model for the confusion background are generally most effective to extract the population parameters.

These analyses highlight the distinction between *event rate*, i.e. the actual rate of binaries merging in the visible Universe, and *detection rate*, which will be a function of the binary parameters and of the detector characteristics. Quantifying this distinction is an ever-growing effort; current event rate estimates based on astrophysical models informed by the LIGO detections show that there is approximately one binary neutron star merger every 13 seconds and one binary black hole merger every 223 seconds [74]. Clearly, the vast majority of these lie well beneath the threshold for detection. This is quantified for example in [99] where the detection rate per year is encoded as a function of the merging binary masses. It is shown that there are several astrophysical models for compact binaries which fit the detection rate of ~ 10 events per year from the aLIGO O1 and O2 data runs. It is clear that the rate of "intermediate" sub-threshold, SNR > 1 events from these models is still low, remaining far from the stochastic limit as defined above.

There is a further degree of complication caused by these unresolved, bright signals in the timestream which behave as shot noise and may hinder the correct interpretation of a detection. These signals represent the random, Poissonian fluctuations in the finite population of sources and will induce a white noise component in the power spectrum. The observational effect that arises is linked to cosmic variance and the selection bias induced by the fact that nearby, bright sources approximately randomly distributed around the observer will mask the underlying structure of the Universe. This is analysed and quantified in [100, 101] for the monopole and anisotropies of an astrophysical background respectively. Here the Poisson white noise is modelled with a free parameter representing the cut-off distance from Earth beyond which the GW signal may be considered a stochastic background, by the definition

given above. The cut-off naturally depends on the sensitivity of the detector array under consideration. This is analysed in [101] for the anisotropic background, and it is found that the white noise component dominates the current LIGO and Virgo detectors, however the planned upgrades of these detectors and the future 3G detectors may be sufficient to curb this effect. Furthermore, as GW sources are finite in time, then Poisson statistics dictates that this noise decays as the inverse of the observation time, hence long observing runs will progressively mitigate this effect.

It's worth pointing out that the time-domain features of astrophysical backgrounds discussed here may themselves provide insight in compact binary population parameters such as the binary coalescence rates for different binary species, as detailed in [32]. Furthermore, the time dependence of astrophysical backgrounds in general may be a powerful discriminating factor between these and backgrounds of primordial origin.

Chapter 3

Maximum-likelihood map solutions

This chapter explores the generalised problem of finding the maximum-likelihood map solution from the time domain data of a scanning experiment involving an array of detectors [see e.g. 102]. This has been considered by a number of authors [40, 42, 45] for gravitational wave telescopes and these have made different choices which include solving for maps performing either a broad- or narrowband frequency integration, carrying out either all-sky or directional searches. These different options give rise to four different approaches to map-making: broadband sky searches, which typically describe the signal with a smooth frequency spectrum; narrowband sky searches, which relax this assumption to each frequency bin; broadband directional searches, which involve optimal filters for small patches of sky and expect a broad spectrum of gravitational waves; finally, narrow-band directional searches, which pick out specific frequencies in the data associated to a particular category of sources. The directional approach, both narrow- and broadband, is known in the literature as *radiometer* method [38]. All these approaches necessarily assume Fourier-transforming the timestream data into the frequency domain where the data are complex, as opposed to the matched-filtering analyses done for the resolved event detections, and the reality of the timestream implies that all the information is contained in half of the frequency range (i.e. positive frequencies). While these approaches will give rise to different optimal maps, the nature of the solution remains the same, as the likelihood describing the data is independent from directional filters and integration extrema.

An important distinction can be made with respect to solving for a map in sky coordinates (pixels) or in the Fourier domain (spherical harmonic coefficients). Usually, in interferometry, the choice has been to obtain a Fourier domain solution since the measurements directly constrain this domain [see e.g 25, for an application to CMB interferometry]. The reconstruction of the actual sky "image" is then left as a separate inversion problem, and is considered straightforward. However, as will be discussed in Chapter 4, in the case of GW detectors there is a non-trivial coupling of spherical modes in the sky response due to the non-compactness of the beam which indicates that adopting this approach may have significant drawbacks. In particular, there is no isotropic kernel for a primary beam describing the coupling of individual pairs of directions. Therefore, limiting the number of spherical modes in the solution in a piecewise fashion could make the problem more ill-conditioned than it actually is. This is the main motivation for the alternative approach adopted in [1–4] and in this thesis, which is to solve for the map directly in the pixel domain. Again, while working in these two domains will give rise to possibly radically different output maps, this will not influence the structure of the method.

The map-making problem can be stated as an inversion of a general observing equation,

$$\boldsymbol{D} = \boldsymbol{A}\,\boldsymbol{s} + \boldsymbol{N}\,, \tag{3.1}$$

where $\boldsymbol{D}$ is an object containing one or multiple observations, the linear operator $\boldsymbol{A}$ projects the signal "map" or array of "maps" $\boldsymbol{s}$ into the domain of the observations, and $\boldsymbol{N}$ is a noise component. The observations span the time domain and constitute a data timestream whilst the signal component spans a separate domain, which may be either pixel or spherical harmonic. To be completely general this domain may be referred to as the *map domain*, however subsequently the pixel domain will be explicitly assumed throughout. $\boldsymbol{A}$ will typically depend on both time and frequency and will describe the directional dependence of the detector response, which will be expressed in the same domain as the signal. The noise component spans the same domain as the observations.

In gravitational-wave science, any measurement of a stochastic signal involves cross correlation of multiple data streams. This stems from the necessity of making quadratic measurements in the strain, where the confused phase information cancels out and just the overall

intensity is preserved. There is no stochastic GW signal which carries coherent information in the phase on the sky, as reviewed in Chapter 2.5. Hence, the true observables $\boldsymbol{s}$ in Equation (3.1) will be the Stokes' parameters, and each component of $\boldsymbol{D}$ will in general be either an auto-correlation or a cross-correlation of two detectors. The Stokes' parameters as seen in Chapter 2 are *real* quantities, and so is the timestream data. These reality conditions will be useful when carrying out the map-making procedures explicitly.

This Chapter tackles the mapping problem in two different scenarios: one where the data streams from the array of detectors are *not correlated*, i.e. they may be considered independent measurements and the noise in each detector is uncorrelated; the other where the data streams are *correlated*, i.e. the detectors used should be considered to share noise. The signal and noise components are considered to be Gaussian; in the limit that the signal is non-Gaussian, the maximum–likelihood map solutions derived below remain unchanged, but the interpretation of their covariance would be affected by the presence of higher order cumulants in the underlying probability densities. The quantities effectively modelled as Gaussian random fields and described by a Gaussian likelihood function actually differ in the two cases: in the case of uncorrelated detectors, it is possible to construct a residual with the cross-correlated data hence the Gaussian variables are the cross-correlated timestream segments from detector pairs, whereas in the case of correlated data the Gaussian variables are the individual data segments from each correlated detector. The different map-making methods laid out here are then applied to ground-based interferometers and the LIGO detectors in Chapter 4, and to LISA in Chapter 5.

3.1 Maps from Uncorrelated Data Streams

Given an array of detectors, in general the data stream of a detector A may be written as $d_A = R_A h + n_A$, where R_A describes the specific response of detector A to the signal "outside the detector" h. In the case of GWBs, h will always average to zero in any sky direction and may thus not be considered a useful observable, and it is necessary to construct a quadratic estimator which involves the cross-correlation of the data streams in the array, for example $D_{AB} = d_A d_B^\star$ for detectors A, B. If the data streams from the detector array may be considered

uncorrelated, then at least to first order the expectation value of the correlated data of two *distinct* detectors in the array will be

$$\langle D_{AB}\rangle = \langle d_A d_B^\star\rangle = R_A R_B^\star \langle h\, h^\star\rangle + \langle n_A\, n_B^\star\rangle \approx A_{AB}\, \boldsymbol{s}\,, \tag{3.2}$$

where A_{AB} is defined as the cross-correlated response of the two, and quantifies the ability of the detector pair to constrain the GW intensity directly. This assumption is especially useful in the case of very low signal-to-noise (SNR) measurements, as the direct cross-correlation of the data sets allows to average out a portion of the individual detector noise and put the signal below within detection reach. Regardless of SNR, it is then possible to build an estimator which is based on the statement that the *residual* $D_{AB} - \langle D_{AB}\rangle$ averages statistically to zero:

$$\langle D_{AB} - \langle D_{AB}\rangle\rangle = 0\,. \tag{3.3}$$

The detector labels can be substituted with the *baseline* index $\boldsymbol{b}$, which labels the vector connecting the pair of detectors in 3D space. To lighten the notation, let $\boldsymbol{D}$ be the observation made for the baseline $\boldsymbol{b}$. Similarly, let $\boldsymbol{A}$ represent the response of the baseline, which will in general be Hermitian by construction.

Under the assumption that each observation $\boldsymbol{D}$ is Gaussian and obeys Equation (3.3) and the noise component of each detector is a zero–mean, Gaussian variate with covariance $\boldsymbol{N} = \langle \boldsymbol{n} \otimes \boldsymbol{n}^\star\rangle$ the maximum-likelihood solution for the signal map is obtained by minimizing the Gaussian likelihood $\mathscr{L}$ of a map $\boldsymbol{s}$ given the residual,

$$\mathscr{L}(\boldsymbol{D}|\boldsymbol{s}) = \frac{1}{(2\pi)^{\mathscr{N}/2}|\boldsymbol{C}|^{1/2}} e^{-\frac{1}{2}(\boldsymbol{D}-\langle\boldsymbol{D}\rangle)^\dagger \boldsymbol{C}^{-1}(\boldsymbol{D}-\langle\boldsymbol{D}\rangle)}\,, \tag{3.4}$$

where $\boldsymbol{C}$ is the covariance of the residual, explicitly [50]

$$\boldsymbol{C_b} = \left(\langle \boldsymbol{D_b}\boldsymbol{D_b}^\dagger\rangle - \langle \boldsymbol{D_b}\rangle\langle \boldsymbol{D_b}^\dagger\rangle\right) \equiv \langle \boldsymbol{N_b}\boldsymbol{N_b}^\dagger\rangle\,, \tag{3.5}$$

where $\mathtt{N}$ is the noise covariance. In other words, the covariance of the residual does not contain the GW signal, hence it is simply the covariance of the noise. Maximising the likelihood is

equivalent to minimising the χ^2, which up to a normalisation factor is

$$\chi^2 = -\ln \mathscr{L}(\boldsymbol{D}|\boldsymbol{s}) \propto \frac{1}{2}\left[(\boldsymbol{D}-\langle\boldsymbol{D}\rangle)^\dagger \boldsymbol{C}^{-1}(\boldsymbol{D}-\langle\boldsymbol{D}\rangle) + \mathrm{Tr}\ln \boldsymbol{C}\right]. \tag{3.6}$$

Noticing that

$$\frac{\partial \boldsymbol{D}}{\partial \boldsymbol{s}} = \boldsymbol{A}, \qquad \frac{\partial \boldsymbol{D}^\dagger}{\partial \boldsymbol{s}} = \boldsymbol{A}^\dagger, \tag{3.7}$$

then

$$\frac{\partial \chi^2}{\partial \boldsymbol{s}} = \boldsymbol{A}\boldsymbol{C}^{-1}\left(\boldsymbol{D}^\dagger - \boldsymbol{A}^\dagger \boldsymbol{s}\right) + \boldsymbol{A}^\dagger \boldsymbol{C}^{-1}\left(\boldsymbol{D} - \boldsymbol{A}\boldsymbol{s}\right) = 0, \tag{3.8}$$

as $\boldsymbol{C}^{-1}$ is symmetric. All the quantities and operators here are written as matrices in bold as they will in general span multiple dimensions, i.e. time and frequency, and include multiple baselines. This is just the sum of two objects where one the conjugate of the other, and leads to the solution for the estimated map

$$\tilde{\boldsymbol{s}} = \left(\boldsymbol{A}^\dagger \boldsymbol{C}^{-1}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^\dagger \boldsymbol{C}^{-1}\boldsymbol{D}. \tag{3.9}$$

In practice, this solution may need to be implemented iteratively as the noise covariance is not known *a priori*, and needs to be estimated from the data itself. This iterative estimation is carried out for example with CMB data, when estimating maps polluted by foreground noise. However, in the case of a noise dominated timestream, this solution becomes a very simple and effective way of estimating the optimal map solution, as shown below.

Low Signal Case

In the case of low signal, Equation (3.9) becomes a simple closed-form solution of the mapping problem, assuming that the noise component in the data dominates the measurement. In practice, this implies that the residual covariance $\boldsymbol{C}$ is well estimated by the product of the noise power spectra *in each data segment*, $\mathtt{N} = N_A \cdot N_B$, which is quartic in the data dimension. Hence the solution becomes

$$\tilde{\boldsymbol{s}} = \left(\boldsymbol{A}^\dagger\ \mathtt{N}^{-1}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^\dagger \mathtt{N}^{-1}\boldsymbol{D}. \tag{3.10}$$

This solution is used to estimate maps of the GWB from cross-correlated LIGO Hanford and LIGO Livingston data in several papers from within the LIGO collaboration, for example [48, 103–105], and in our own works [1, 3, 4] which are presented in Chapter 4.

It is useful to define the two operators in Equation (3.10),

$$\boldsymbol{z} = \boldsymbol{A}^{\dagger}\mathbb{N}^{-1}\boldsymbol{D}\,, \tag{3.11}$$

and

$$\boldsymbol{M} = \boldsymbol{A}^{\dagger}\ \mathbb{N}^{-1}\boldsymbol{A}\,, \tag{3.12}$$

and examine their properties. $\boldsymbol{z}$ is referred to as the *projection* operator and effectively projects the noise-weighted data stream through the cross-correlated response, and the *beam pattern* matrix $\boldsymbol{M}$ corresponds to the inverse noise covariance matrix [106] which contains the noise-weighted correlation of each pair of directional indices, which may be either pixel ps or spherical harmonic ℓms. To investigate these further it becomes necessary to make operative choices and choose a domain to work in. Stochastic GW signals, as reviewed in Chapter 2, are modelled and measured in Fourier space hence the observed timestream is Fourier transformed so that $\boldsymbol{d}$ spans the frequency domain. In practice it is convenient to divide the timestreams into smaller time segments and Fast-Fourier-transform (FFT) each segment, which is treated as an independent measurement. This reduces the cost of Fourier transforming, excludes frequencies that are far into the $1/f$ tail and ensures that the noise properties are stationary within each segment. In the following the segments are labelled using the superscript τ. With the choice mentioned above and adopted here of $\boldsymbol{s}$ spanning the sky pixel domain Equation (3.1) then becomes

$$D^{\boldsymbol{b}}_{\tau,f} = \sum_{p} A^{\boldsymbol{b}}_{fp} s_p + N^{\boldsymbol{b}}_{\tau,f}\,, \tag{3.13}$$

where p runs over all pixels on the sky. The time dependence of A_{fp} is implicit through the pixel index, as the pixel response will be different at different times. In the frequency domain, assuming stationarity, the noise contribution takes on a particularly simple form with

$$\mathbb{N}^{\boldsymbol{b}}_{ff'} = \delta(f - f')P^{\boldsymbol{b}}(f)\,, \tag{3.14}$$

where $P^{\boldsymbol{b}}$ is the product of the two power spectra associated with the detectors in the array which make up the baseline $\boldsymbol{b}$. In the low signal case, these are simply the power-spectrum densities (PSDs) of the data sets. The power spectrum of the baseline noise will, in general, be different for separate time frames τ due to long–term drifts in detector characteristics. To account for this the noise spectra must be re-estimated for each time segment.

When the observations include multiple baselines $\boldsymbol{b}$ the expressions are modified by summing over the individual contributions in such a way that

$$z_p = \sum_{\boldsymbol{b},\tau,f} A^{\boldsymbol{b}\,\dagger}_{\tau,f} \left(\mathtt{N}^{\boldsymbol{b}}_{\tau f}\right)^{-1} D^{\boldsymbol{b}}_{\tau,f}\,, \tag{3.15}$$

and

$$M_{pp'} = \sum_{\boldsymbol{b},\tau,f} A^{\boldsymbol{b}\,\dagger}_{\tau,f} \left(\mathtt{N}^{\boldsymbol{b}}_{\tau f}\right)^{-1} A^{\boldsymbol{b}}_{\tau,f} \tag{3.16}$$

where the superscripts and subscripts denote the individual contributions to each frequency, time frame and baseline. This combination follows from extending the likelihood formulation (3.4) by noting that

$$\mathscr{L}(\boldsymbol{D}|\boldsymbol{s})_{\rm tot} = \prod_{\boldsymbol{b}} \mathscr{L}(\boldsymbol{D}_{\boldsymbol{b}}|\boldsymbol{s})\,, \tag{3.17}$$

which translates into a sum in the log likelihood. This will hold as long as the individual data sets have uncorrelated noise, and the times and frequencies may be considered independent. Although the times and frequencies are not explicitly labelled in Equation (3.17), the same principle holds. The map then obtained by combining Equations (3.15-3.16),

$$\tilde{s}_p = M^{-1}_{pp'} z_p\,, \tag{3.18}$$

is *optimal* under these assumptions [106], as it returns the maximum likelihood estimate for the signal minimizing the signal variance.

The frequency dependence of the solution $\tilde{\boldsymbol{s}}$ is not manifest here, and a broadband approach may be non-trivial to define. Typically, a directionally-independent functional form is assumed and is integrated out, effectively becoming part of the response $\boldsymbol{A}$. This type of approach is explained in detail in Chapter 4, where specific cases are presented. A way around this problem is to consider narrowband integrations or single frequencies to probe

the frequency dependence of the maps, but the conditioning of the beam-pattern matrix may become prohibitive. This is also explored in Chapter 4.

3.2 Maps from Correlated Data Streams

In the case of correlated data, the simplification implied in Equation (3.2) no longer holds, and the likelihood (3.4) is no longer appropriate as it is not possible to define a residual. The maximum likelihood map estimator with correlated data streams will still necessarily be quadratic in the strain as any linear component will vanish over time and frequency averaging. It is possible then to take a step back, and write a valid description of the data by simply assuming that the signal and noise components *in each individual detector* behave as Gaussian variables, and thus use the Gaussian likelihood

$$\mathscr{L}(\boldsymbol{d}|\boldsymbol{s}) = \frac{1}{(2\pi)^{\mathscr{N}/2}|\boldsymbol{C}|^{1/2}} e^{-\frac{1}{2}\boldsymbol{d}^{\dagger}\boldsymbol{C}^{-1}\boldsymbol{d}}, \tag{3.19}$$

which is simply the probability of observing the data vector $\boldsymbol{d}$ given the map $\boldsymbol{s}$. $\mathscr{N}$ represents the degrees of freedom in the problem. Here the data vector spans the detector array space, and each component is linear to the strain h as

$$\boldsymbol{d} = \boldsymbol{R}h + \boldsymbol{n}, \tag{3.20}$$

where $\boldsymbol{n}$ is the noise vector and $\boldsymbol{R}$ is the linear response vector. The noise is assumed to have zero mean and covariance $\boldsymbol{N} = \langle \boldsymbol{n} \otimes \boldsymbol{n}^{\star} \rangle$, whereas $\boldsymbol{C}$ is the covariance of the data, $\boldsymbol{C} = \boldsymbol{A}\tilde{\boldsymbol{s}} + \boldsymbol{N}$, which describes the correlations in the array. $\tilde{\boldsymbol{s}}$ is the map estimate, and $\boldsymbol{A}$ is the cross-correlated response matrix of the detector array as above. While at a glance this may seem like a simpler problem to solve and rests on fewer assumptions than that presented in the previous Section, note that in this case the covariance is the full covariance of the data, not of the residual, hence in general it will be necessary to estimate this iteratively to obtain an optimal estimate of the map. Furthermore, in this case the full correlation between data streams is being considered, including the auto-correlations for each detector; these may not

be removed from the likelihood however when considering noise-dominated data streams it is extremely convenient to exclude these, as explained previously.

The solution which maximises $\mathscr{L}$ is straightforward to derive. To start, the solution will be the same to that which minimises

$$\chi^2 = \boldsymbol{d}^\dagger \boldsymbol{C}^{-1} \boldsymbol{d} - \log(\det \boldsymbol{C}) \,, \tag{3.21}$$

which is the associated χ^2 up to a normalisation constant. Due to the linearity and cyclic property of the trace, this may be rewritten as

$$\chi^2 = \mathrm{Tr}\left[\boldsymbol{D} \cdot \boldsymbol{C}^{-1}\right] + \mathrm{Tr}\log \boldsymbol{C} \,, \tag{3.22}$$

where the trace occurs over the array components, and the '$\cdot$' is a matrix product. The data matrix $\boldsymbol{D} = \boldsymbol{d}^\star \otimes \boldsymbol{d}$ is analogous to the one used in the uncorrelated case and is quadratic in the strain. The derivative of the χ^2 with respect to the map is

$$\frac{\partial \chi^2}{\partial \tilde{\boldsymbol{s}}} = -\mathrm{Tr}\left[\boldsymbol{D}\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\boldsymbol{C}^{-1}\right] + \mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\right] , \tag{3.23}$$

where the relation

$$\frac{\partial \boldsymbol{U}^{-1}(x)}{\partial x} = -\boldsymbol{U}^{-1}\boldsymbol{U}'\boldsymbol{U}^{-1} \tag{3.24}$$

was used. Introducing a $\boldsymbol{C}^{-1}\boldsymbol{C}$ term in the second trace on the left-hand side and equating the whole expression to zero, one gets

$$\mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\boldsymbol{C}^{-1}\boldsymbol{D}\right] = \mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}C^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\right]\tilde{\boldsymbol{s}} + \mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\boldsymbol{N}\right] , \tag{3.25}$$

where again the cyclic property of the trace was used and the covariance was explicitly decomposed as $\boldsymbol{C} = \frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}}\tilde{\boldsymbol{s}} + \boldsymbol{N}$. The maximum-likelihood solution for the map $\tilde{\boldsymbol{s}}$ may then be written as

$$\tilde{\boldsymbol{s}}_p = \mathscr{F}^{-1}_{pp'} \cdot \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}_{p'}}\boldsymbol{C}^{-1}(\boldsymbol{D}-\boldsymbol{N})\right] \tag{3.26}$$

$$\mathscr{F}_{pp'} = \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}_p}\boldsymbol{C}^{-1}\frac{\partial \boldsymbol{C}}{\partial \tilde{\boldsymbol{s}}_{p'}}\right] \tag{3.27}$$

where the pixel indices are chosen and explicitly written down for clarity. $\mathscr{F}$ is typically referred to as the *Fisher information* matrix, whereas the trace in Equation (3.26) is simply the gradient term. This is utterly analogous to a procedure commonly used in CMB mapping, for example presented in [107, 108]. This is not a closed-form solution and requires an iterative approach: an initial guess $\tilde{\boldsymbol{s}}_{\rm in}$ is plugged into $\boldsymbol{C}$ and Equation (3.26) returns a first estimate of $\tilde{s}_p$, then the calculation is repeated until convergence is reached.

As in the case explained previously, it is convenient to segment the data sets into smaller time chunks and FFT each chunk to work directly in frequency space. Assuming these chunks are independent, the total likelihood will be a product of Gaussian likelihoods of the form (3.19), and this will translate into a sum over time labels τ and frequencies f in the solution,

$$\tilde{\boldsymbol{s}}_p = \mathscr{F}^{-1}_{pp'} \cdot \frac{1}{2} \mathrm{Tr} \left[\sum_{\tau,f} \boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{\boldsymbol{s}}_{p'}} \boldsymbol{C}^{-1}_{\tau,f} (\boldsymbol{D}_{\tau,f} - \boldsymbol{N}_f) \right], \tag{3.28}$$

$$\mathscr{F}_{pp'} = \frac{1}{2} \mathrm{Tr} \left[\boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{\boldsymbol{s}}_{p'}} \boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{\boldsymbol{s}}_{p'}} \right]. \tag{3.29}$$

The noise model $\boldsymbol{N}$ will generally not depend on time, but it is easily possible to include the dependency by introducing it above. As before, it is also possible to consider sums over short frequency intervals or single frequencies to probe the frequency dependence of the maps, if the conditioning of the Fisher matrix permits.

Note how the solution in Equations (3.28-3.29) substantially differs from the solution presented in the previous Section, Equations (3.15-3.16), such that the former do not simply reduce to the latter when the noise correlations are negligible. This is due to the fact that the data model itself is different in the two cases. In particular, in the case of uncorrelated, noise-dominated timestreams cross-correlating the individual detector measurements is a necessary step to eliminate the local noise in the data and reach detection. This is typically the case considered for ground-based GW detecors such as LIGO as described in the next Chapter. On the other hand, the solution presented in Section 3.1 breaks down as soon as the assumption in Equation (3.2) no longer holds; specifically, in the case of ground-based GW detectors it is possible that faint noise sources, such as Schumann resonances from the Earth's magnetic field, may correlate the noise in the interferometers. It is thus possible that the

second method presented in Section 3.2 may become useful for ground-based interferometer networks as well.

To conclude this section, some remarks on analogies and differences between GWB and CMB analysis, which will be helpful to contextualise the following Chapters.

The analysis of interferometric data from GW detectors for the purpose of constraining GWBs may be split into two main stages: first, a compression step from time/frequency domain to sky coordinates, with the methods for example explained in this Chapter; then, it is possible to isotropise the measurement to obtain the angular power spectrum (see Equation (2.30)). This is utterly similar to the process in CMB analysis, where there is initially a map estimation and subsequently a final C_ℓ estimate.

In CMB observations made with incoherent (bolometric) detectors, the first stage uses a χ^2 estimate to obtain a maximum likelihood map [102]. This is possible because the signal component of the data does not vanish in the ensemble mean, which allows the definition of a residual whose covariance is just the noise covariance. This is similar to the method derived in Section 3.1 for GWB mapping in the case of uncorrelated data streams, and in particular for the low signal case which gives a closed-form solution. This type of solution is well-suited for arrays of ground-based interferometers as these are set far apart on the Earth's surface. In the case of CMB mapping however the noise is often estimated iteratively, as the signal component introduces a bias in the noise estimation.

This is not possible for data obtained from noise correlated, coherent detectors such as LISA. In this case both signal and noise components in the time-domain vanish in the ensemble mean, and no residual can be defined due to the correlated nature of the data. It is then necessary to use the full likelihood of the data as shown in Equation 3.19. There is no closed-form solution for the maximum likelihood in this case and iterative methods such as quadratic estimators must be used to estimate the signal [107], as explained above. This is analogous to the second stage in CMB analysis, i.e. the C_ℓ estimation, where the map signal component vanishes in the ensemble mean limit and therefore the full data likelihood is used to estimate the signal.

There is a final, important point to be made regarding the power spectrum estimation for GWBs and the analogy with other methods, for example used in CMB measurements and radio interferometry. In most astronomical experiments, the beam is well-modelled by a single, relatively narrow Gaussian on the sky, hence it is possible to use a flat-sky approximation applied to each individual sky observation, and "piece the sky" back together in a second step. The flat-sky approximation allows the unequivocal use of the spherical harmonic expansion, as in the limit where the observed sky is a small flat patch and each observation is independent the observation will not introduce extra correlations in the sky Fourier domain. Even when the beam is extended, the Gaussianity of the beam ensures it is isotropic, averaging over the harmonic coefficients. Working in the Fourier domain is then extremely useful as it is possible to compactify information and truncate calculations at the required scale, optimising the computation of a sky map. This is not possible with GW observations given the highly non-Gaussian and non-compact characteristics of the beam-patterns for each baseline. Effectively, the flat-sky approximation does not apply, hence working in pixel space becomes preferable as the Fourier domain contains correlations which are tricky to account for and may pollute the measurement. This is shown in Chapter 4 and 5, and the implications are explored in depth.

Chapter 4

Mapping the Gravitational-Wave Background with Ground-based Detectors

The scientific community is actively instituting a fast-growing, world-wide gravitational-wave detector network, with several presently operating, planned and proposed Earth based instruments. Most past and present designs aim to carry out interferometric measurements of the strain impressed by a wave in transit on an array of laser beams arranged in *arms*, and are often referred to as laser interferometer observatories or telescopes. There are also alternative instruments designed for detecting GWs, such as resonant masses or *Weber bars*: long, solid metal antennae isolated from outside vibrations which resonate at GW frequencies; see e.g. AURIGA in Padova, Italy [109]. However, the work presented in this thesis focuses on the former style of detector as it is the only that has been successfully employed to detect gravitational waves.

There are currently four functioning units which have comparable sensitivity ranges: two detectors which form the Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) array, one in Hanford (WA) and one in Livingston (LA); the Advanced Virgo detector in Cascina, Italy, analogous to LIGO; finally, and most recently, the Kamioka Gravitational Wave Detector (KAGRA) in the Gifu prefecture, Japan. These cycle through observing runs and shut-down periods for maintenance and upgrades, and a systematic global effort is made to synchronise observations and maximise detection potential. These interferometers, often referred to as *second generation* (2G) detectors, share similar technologies and have

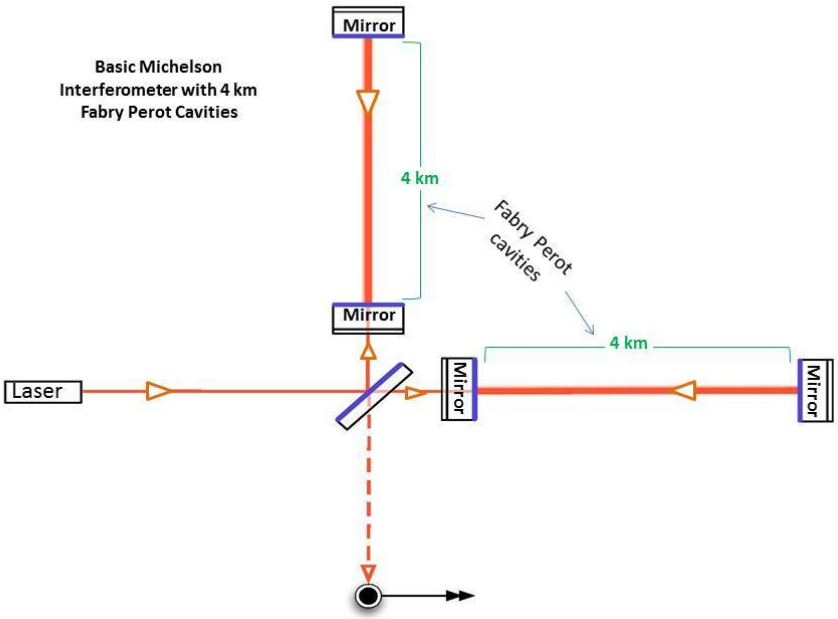


Figure 4.1: Schematic representation of a LIGO-like laser interferometer, from the official Caltech LIGO Lab website[1]. Note the *beam splitter* at the junction of the Fabry-Perot arms, which splits the laser beam, and the mirrors which amplify the signal.

orthogonal Fabry–Pérot arm cavities of the order of a few kilometres long. As such, they are most sensitive to signals with frequencies between 10 Hz to 10^3 Hz (for reference, see Figure 1.1). A schematic sketch of a LIGO-like detector is shown in Figure 4.1. Other operational facilities include the Australian International Gravitational Observatory (AIGO) in Gingin, Australia and the GEO600 detector in Hanover, Germany; these do not match the observational abilities of those listed above given their reduced size and are used as testing grounds for new technologies [110].

The Indian Initiative in Gravitational-wave Observations (IndIGO or LIGO India) is planned to be installed in the Hingoli District, India, and is currently the only confirmed new GW detector that will be added to the network. There are also proposals for future, *third generation* (3G) detectors, such as the Einstein Telescope (ET), based in Europe, and Cosmic Explorer (CE), in the USA. The provisional design for ET consists of a three-, equal-armed detector with ten kilometre arms, whereas CE is planned to have two orthogonal, forty kilometre arms, sharing the LIGO design style. Given their characteristics these detectors should be

[1] https://www.ligo.caltech.edu

approximately sensitive to the same frequency range as the second generation detectors, but with a considerably improved sensitivity – further details are given in Chapter 6.

As mentioned above, 2G ground-based instruments such as the ones installed at the LIGO facilities are L-shaped Fabry-Perot vacuum tubes with "free" masses hanging at the ends of each arm. A passing gravitational wave changes the relative distance between the masses, and is quantified in the interference of the laser beams travelling along each arm [111] (see Figure 4.1). The original polarisation of the passing wave may be reconstructed through multiple separated measurements of the same wave, through amplitude ratio analysis [112]. With these state-of-the-art instruments, the LIGO and Virgo science collaboration (LVC) has successfully identified multiple binary compact object merger events within redshift $\sim$1 by directly detecting the gravitational-wave signal emitted in the final stages of the merger. These signals, in the 2G frequency band, last up to the order of a minute and populate the timestream sparsely, and a refined matched-filtering analysis in the time domain is employed to recognise the waveforms in the data. The potential of both present and future ground-based interferometers goes beyond these searches; merger events are necessarily the outlying events of their distribution, and a stochastic GWB of more frequent, longer and much lower signal-to-noise events is expected, as detailed in Chapter 2.

In this Chapter the detection and mapping method for the GWB presented in Chapter 3 is tailored to specific ground-based interferometers and tested on both mock and real data. To start, the detector response is introduced and analysed both in the pixel and spherical harmonic domain. Then, the optimal map solution is written including the detector response explicitly and a method to inject SGWB data and reconstruct maps is tested. This includes a novel section with reconstruction of polarised GWBs, i.e. the recovery of all four Stokes' parameters described in Chapter 2. Finally, the method is tested on the aLIGO data runs O1 and O2. The results are in good agreement with those from the LVC collaboration and provide only upper limits for a GWB. Almost the entirety of this Chapter is drawn from this author's papers [1, 3, 4], with some original unpublished additions.

4.1 Detector response

At present, ground-based interferometers are comprised of two orthogonal kilometre-scale Fabry-Perot cavities, generically referred to as the interferometer arms of length L. A single laser beam is split between the arms and travels about 300 times back and forth in each arm before interfering. This process effectively elongates the optical path, turning a 4km arm such as LIGO's into a 1200km arm, thus increasing the detector sensitivity as there are more photons interfering at any given time. The "extended" arm length $\tilde{L}$ sets the range of GW frequencies the detector is most sensitive to, i.e. $f = c/\tilde{L} \sim 10^2$ Hz, whereas the effective scale of the detector L is much smaller than the wavelength of the GWs probed, $L \ll \lambda \sim 10^3$ km. This is referred to as the *small-antenna* limit, or sometimes the *long-wavelength* limit although this may be ambiguous when comparing different length scales [43]. The interferometric readout is synthesised into the *timestream* data, which is in dimensionless strain units.

The timestream $d_A(t, \boldsymbol{x}_A)$ generated by a single, two-armed interferometer labelled A at time t and position $\boldsymbol{x}_A$ may be decomposed in a signal component and a noise component, $d_A = R_A h + n_A$, as explained in Chapter 3. The signal component, which includes the detector response component R_A, can be Fourier expanded to yield the signal as a function of frequency [see e.g. 35]:

$$\tilde{d}_A(f, \boldsymbol{x}_A) = \int_{S^2} d\hat{\boldsymbol{n}} \sum_{P=+,\times} h_P(f, \hat{\boldsymbol{n}})\, F_A^P(f, \hat{\boldsymbol{n}})\, e^{i2\pi f\, \hat{\boldsymbol{n}}\cdot\boldsymbol{x}_A} + n_A(f)\,, \tag{4.1}$$

where $h_P(f, \hat{\boldsymbol{n}})$ is the Fourier component at frequency f of the gravitational-wave strain just outside the detector and F_A^P is the *polarisation response function*. The latter is defined through the contraction of the polarisation tensor and the detector tensor as

$$F_A^P(\hat{\boldsymbol{n}}) = \frac{1}{2}\left(\hat{\boldsymbol{u}}_A \otimes \hat{\boldsymbol{u}}_A - \hat{\boldsymbol{v}}_A \otimes \hat{\boldsymbol{v}}_A\right)^{ij} e^P_{ij}(\hat{\boldsymbol{n}})\,, \tag{4.2}$$

where $\hat{\boldsymbol{u}}_A$ and $\hat{\boldsymbol{v}}_A$ are the unit vectors defining the orientation of the two detector arms. The expansion in Equation (4.1) is then simply the strain as written in Equation (2.46), which has spatial indices ij through the polarisation basis, contracted with the detector tensor as defined above. Note how the detector tensor is just the difference between the response of

each detector arm, $\boldsymbol{u}$ and $\boldsymbol{v}$. This basis may be defined with respect to the detector, which could simplify calculations, however this is not convenient for the purposes of map-making as will become clear. Furthermore, the polarisation response does not depend on frequency or arm length. This is an approximation and in general the response would include extra terms, however in the small antenna limit these simplify out. This is not the case for LISA, as will be shown in Chapter 5.

A signal that is composed of an incoherent superposition of many components, as in the case of a stochastic background, will vanish when averaged over time. As explained in Chapter 3, to observe the GWB it is necessary to consider the integration of the square of the signal. This could be done by squaring the signal of a single detector but the signal-to-noise Ratio (SNR) level can be improved greatly by considering the *cross-correlation* of two or more detectors under the assumption that their individual noise contributions are not correlated. This may be considered true to good approximation in the case of ground-based interferometers, as the main sources of noise are seismic perturbations and disturbances in the optical bench [113], which may both be considered local if the detectors are far apart on the Earth's surface. It's important to note however that there are potential sources of correlated noise between LIGO interferometers which, even though sub-dominant with respect to the uncorrelated noise sources in the LIGO band, may become significant when carrying out cross-correlated stochastic searches. The primary candidate for this type of noise are Schumann magnetic resonances caused by the Earth's electromagnetic field, which may mimic a correlated stochastic signal across the network and be comparable to the sensitivity of the detectors to the SGWB [114, 115].

A pair of interferometers at locations $\boldsymbol{x}_A$ and $\boldsymbol{x}_B$ identifies a unique *baseline* $\boldsymbol{b} = \boldsymbol{x}_A - \boldsymbol{x}_B$ which may be used to label the cross-correlated data stream,

$$\begin{aligned}\langle D_{\boldsymbol{b}}(f)\rangle &\equiv \langle \tilde{d}_A(f)\tilde{d}_B^{\star}(f)\rangle \\ &= \int_{S^2} d\hat{\boldsymbol{n}}\, d\hat{\boldsymbol{n}}' \sum_{\substack{P=+,\times \\ P'=+,\times}} \langle h_P h_{P'}^{\star}\rangle F_A^P F_B^{P'}\, e^{i2\pi f(\hat{\boldsymbol{n}}\cdot\boldsymbol{x}_A - \hat{\boldsymbol{n}}'\cdot\boldsymbol{x}_B)}.\end{aligned} \tag{4.3}$$

Expanding in the polarisation bases and inserting the relation between the second order moments of the strain and the Stokes parameters defined in Chapter 2, Equation (2.51), the

observing equation for the frequency domain signal becomes

$$\langle D_{\boldsymbol{b}}(f)\rangle = \int_{S^2} d\hat{\boldsymbol{n}} \sum_W \gamma^W_{\boldsymbol{b}}(f,\hat{\boldsymbol{n}})\, W(f,\hat{\boldsymbol{n}})\, e^{i2\pi f\,\hat{\boldsymbol{n}}\cdot\boldsymbol{b}}\,, \tag{4.4}$$

where W labels all the Stokes parameters in general, I, Q, U, and V, and $\gamma^W_{\boldsymbol{b}}$ are the cross-correlated responses known as *overlap functions* in the literature and are constructed through different combinations of the beam pattern functions of the two detectors in the baseline,

$$\gamma^I_{\boldsymbol{b}}(\hat{\boldsymbol{n}}) = \frac{5}{8\pi}\left(F^+_A F^{+\star}_B + F^\times_A F^{\times\star}_B\right), \tag{4.5}$$

$$\gamma^Q_{\boldsymbol{b}}(\hat{\boldsymbol{n}}) = \frac{5}{8\pi}\left(F^+_A F^{+\star}_B - F^\times_A F^{\times\star}_B\right), \tag{4.6}$$

$$\gamma^U_{\boldsymbol{b}}(\hat{\boldsymbol{n}}) = \frac{5}{8\pi}\left(F^+_A F^{\times\star}_B + F^\times_A F^{+\star}_B\right), \tag{4.7}$$

$$\gamma^V_{\boldsymbol{b}}(\hat{\boldsymbol{n}}) = -i\frac{5}{8\pi}\left(F^+_A F^{\times\star}_B - F^\times_A F^{+\star}_B\right). \tag{4.8}$$

The overlap functions contain all the information regarding the relative position of the detectors and their orientation with respect to a fixed gravitational wave polarisation basis. Note, again, that these do not depend on the arm- nor baseline- lengths in this approximation. The baseline length comes into play in the exponential term in Equation 4.4. The normalisation of the functions is imposed by the condition that the angular integral of $\gamma^W_{\boldsymbol{b}}$ is unity in the case of two co-located and co-aligned detectors [57]. In analogy with radio interferometry, the observation as written in Equation (4.4) may be considered to be performed directly in the baseline Fourier space, or "uv-space" [116].

In the following the coordinate system and polarisation basis are fixed to the galactic coordinate frame, as the signal is considered to be static in this frame. For Earth based detectors being considered here this means that $\gamma^W_{\boldsymbol{b}}$ will not be stationary. To avoid cumbersome notation this time-dependence will remain implicit in the equations but it will be useful to recall that the baseline vector is not stationary in galactic coordinates, so the time-dependence is easily re-introduced via $\boldsymbol{b} \to \boldsymbol{b}(t)$. In this section the analysis will be limited to an unpolarised background ($Q = V = U = 0$), hence only the intensity Stokes parameter ($W = I$) of the strain will be considered, and the label W will be dropped.

It is useful at this stage to consider the observing characteristics that are particular to gravitational wave detector cross-correlations. Although the maximum likelihood map solution employed in the following sections will solve the sky directly in the coordinate domain, as anticipated in Chapter 3, it is convenient to expand the observations in spherical harmonics since the cross-correlation signal measures the Fourier domain directly. In this sense, Equation (4.4) may be interpreted as the transform of the signal intensity, which is then measured in baseline space as mentioned above. This is analogous to radio interferometry where the measurement can be most naturally represented in the "uv–plane", i.e. the Fourier transform of the xy–plane, where the baseline is embedded. In the case of gravitational wave detectors however, the "beam" (i.e. the overlap functions) with which the sky is convolved is not compact, hence no flat-sky approximation can be used and a gravitational wave measurement can be represented on a "uvw–sphere". The beam pattern is also symmetric with respect to the pointing axis, which implies that the detectors are equally sensitive to gravitational waves of opposite momenta.

To introduce the spherical harmonic basis the intensity field I is expanded over basis functions $Y_{\ell m}(\hat{\boldsymbol{n}})$ in terms of coefficients $a_{\ell m}$ [39, 41],

$$I(f,\hat{\boldsymbol{n}}) = E(f)\sum_{LM} a_{\ell m} Y^{\star}_{\ell m}(\hat{\boldsymbol{n}})\,, \tag{4.9}$$

where the spectral dependence of the quadratic signal $E(f)$ is assumed to be direction independent and has been separated out, implying that the sky signal has the same frequency dependence in every direction. This is a strong simplification which relies on the assumption that the observable background is made up of the same category of sources, which all behave similarly in frequency space, such that their incoherent superposition retains the same spectral behaviour. This has been explored in Chapter 2 for astrophysical and primordial backgrounds and will be discussed again in the next section, which illustrates the mapping algorithm.

Inserting this into (4.4) and making use of the spherical harmonic expansion of the plane wave

$$e^{i\boldsymbol{k}\cdot\boldsymbol{x}} = 4\pi\sum_{\ell m} i^{\ell}\, j_{\ell}(kx)\, Y_{\ell m}(\hat{\boldsymbol{k}})\, Y^{\star}_{\ell m}(\hat{\boldsymbol{x}})\,, \tag{4.10}$$

where $j_\ell(x)$ are spherical Bessel functions, the observing equation for the directly observed uvw–sphere harmonic coefficient $D^b_{\ell m}$ defined via

$$\langle D_{\boldsymbol{b}}(f)\rangle = E(f)\sum_{\ell m} D^b_{\ell m} Y^\star_{\ell m}(\hat{\boldsymbol{b}}) \tag{4.11}$$

is derived as

$$D^b_{\ell m} = i^\ell 4\pi\, E(f) \sum_{\ell' m' LM} j_\ell(2\pi f b)\, a_{\ell' m'} \mathcal{K}^{LM}_{\ell m,\ell' m'}\,. \tag{4.12}$$

The coupling kernel $\mathcal{K}^{LM}_{\ell m,\ell' m'}$ incorporates the Gauntt integral

$$\begin{aligned}\mathcal{K}^{LM}_{\ell m,\ell' m'} &= \gamma_{LM}\int_{S^2} d\hat{\boldsymbol{n}}\, Y_{LM}(\hat{\boldsymbol{n}})\, Y^\star_{\ell m}(\hat{\boldsymbol{n}})\, Y^\star_{\ell' m'}(\hat{\boldsymbol{n}})\,,\\ &= \gamma_{LM}\sqrt{\frac{(2L+1)(2\ell+1)(2\ell'+1)}{4\pi}}\begin{pmatrix} L & \ell & \ell' \\ 0 & 0 & 0\end{pmatrix}\times\begin{pmatrix} L & \ell & \ell' \\ M & -m & -m'\end{pmatrix},\end{aligned} \tag{4.13}$$

where the terms in brackets indicate $3j$ coefficients and the overlap function $\gamma(\hat{\boldsymbol{n}})$ has also been expanded in spherical harmonics as

$$\gamma_{\ell m} = \int d\Omega_{\hat{\boldsymbol{n}}}\,\gamma(\hat{\boldsymbol{n}})\, Y^\star_{\ell m}(\hat{\boldsymbol{n}})\,. \tag{4.14}$$

The $3j$ coefficients obey the triangle condition $|L-l| \le l' \le |L+l|$ and $M+m+m'=0$ for any set of quantum numbers and are zero otherwise.

The kernel $\mathcal{K}$ proves that the projection of three-dimensional plane waves onto the two-dimensional sphere, together with a non-uniform, non-compact weighting of the sky, introduces a coupling between ℓ, ℓ', and L in the spherical harmonic expansion. This expansion is useful when considering how spherical modes are reconstructed by the observations. One aspect that is peculiar to gravitational wave detectors is that they are equally sensitive to modes travelling in opposite directions – there is no possibility of beam forming since gravitational waves cannot be focused using detectors on feasible scales. This implies that the overlap functions are symmetric and therefore the detectors may seem insensitive to a mode that is odd with respect to the same axis. To see this explicitly, it is sufficient to isotropise $\gamma_{\ell m}$ and look at the power spectrum

$$\Gamma_\ell = \frac{1}{2\ell+1}\sum_m \left|\gamma_{\ell m}\right|^2. \tag{4.15}$$

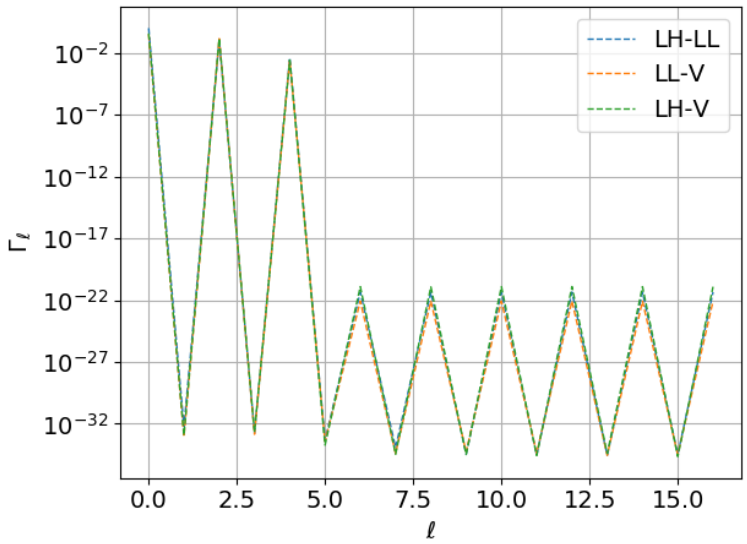


Figure 4.2: Power spectra of the overlap functions γ for different baselines. As the overlap functions are even naturally $\Gamma_\ell = 0$ for odd ℓ. Note there are effectively no differences between different baselines as the baseline length does not enter in this calculation and the different orientations of the detectors only induce differences in the phase m.

Effectively, $\Gamma_\ell = 0$ for odd ℓ as may be seen in Figure 4.2 and one might expect that this in turn means the observations cannot reconstruct $a_{\ell m}$ modes with odd ℓ since a general rotation only couples m and is irreducible in ℓ. This is not the case however since the projection of the plane waves onto the sphere breaks the orthogonality condition – different wave-vectors $\boldsymbol{k}$ contribute to each projected ℓ mode and this breaks the symmetry implied by the overlap functions. This is easily seen in Equation (4.12) considering the observation of an odd mode on the sky $\ell' = 2p+1$ through an even mode of the overlap function $L = 2q$ with $p, q \in \mathbb{Z}^+$. In that case the $3j$ coefficients are non-zero in the range $|2(p-q)-1| \leq \ell \leq 2(p+q)+1$ implying that a measurement at "uvw point" $D^b_{\ell m}$ does not vanish for odd ℓ'.

Given sufficient (angular) phase coverage of the uvw–sphere it should be possible to reconstruct all spherical modes on the sky within a range of angular scales - this should be a well-conditioned inversion problem. For an Earth based detector the phase coverage is provided by the Earth's rotation and the latitudinal positioning of the detectors involved in the baselines. Table 4.1 lists the position and orientation of five detectors considered in this Chapter: the LIGO Observatories (LH and LL), Virgo (V) and Kagra (K) are all operational

Table 4.1: Latitude, longitude and the orientation angle α in degrees of the five currently operating detectors considered as part of various observing combinations in this work. Specifically, α is the angle between the local parallel and the bisector of the aperture of the interferometer [see also 95].

Detector	Label	Latitude	Longitude	α
LIGO Hanford	LH	46.4	-119.4	171.8
LIGO Livingston	LL	30.7	-90.8	243.0
Virgo	V	43.6	10.5	116.5
Kagra	K	36.3	137.2	225.0
GEO600	G	48.0	9.8	68.8

although not all their data sets are open access; the detector GEO600 (G) was included in the analysis to forecast what an extra detector may bring to the network.

Figure 4.3 shows the overlap functions, in galactic coordinates and at arbitrary time, for the ten baseline combinations being considered. Each map is normalised according to the definition of the overlap functions, which depends on position and orientation of the different detectors, so it is not immediately straightforward to compare these. The combination of extensive but inhomogeneous weighting and limited range in latitudes means that, on time scales shorter than 24 hours, the problem of sky reconstruction is ill-conditioned. However, as will be demonstrated in the next section, integrating for sufficiently long time scales improves the conditioning considerably.

Figure 4.4 shows a more useful representation of the coverage of the ten baseline set. The track of each baseline vector direction can be plotted on the sky, which highlights how the restricted range of latitudes of the detectors affects the coverage on the sky. This is equivalent to the definition of "uv coverage" in radio astronomy, hence the title of the plot. To obtain an idea of the sensitivity of each baseline b to different scales one can integrate the spherical Bessel function for each b in frequency to obtain an estimate for the *frequency-integrated sensitivity* $\mathscr{S}$,

$$\mathscr{S} = \int df\, w(f;\sigma,\mu)\, j_\ell(2\pi f b)\,, \tag{4.16}$$

where $w(f;\mu,\sigma)$ is a Gaussian window function that accounts for the effective sensitivity to different frequencies. $\mu = 200$ Hz and $\sigma = 50$ Hz were chosen to approximate the sensitivity of a LIGO-type baseline.

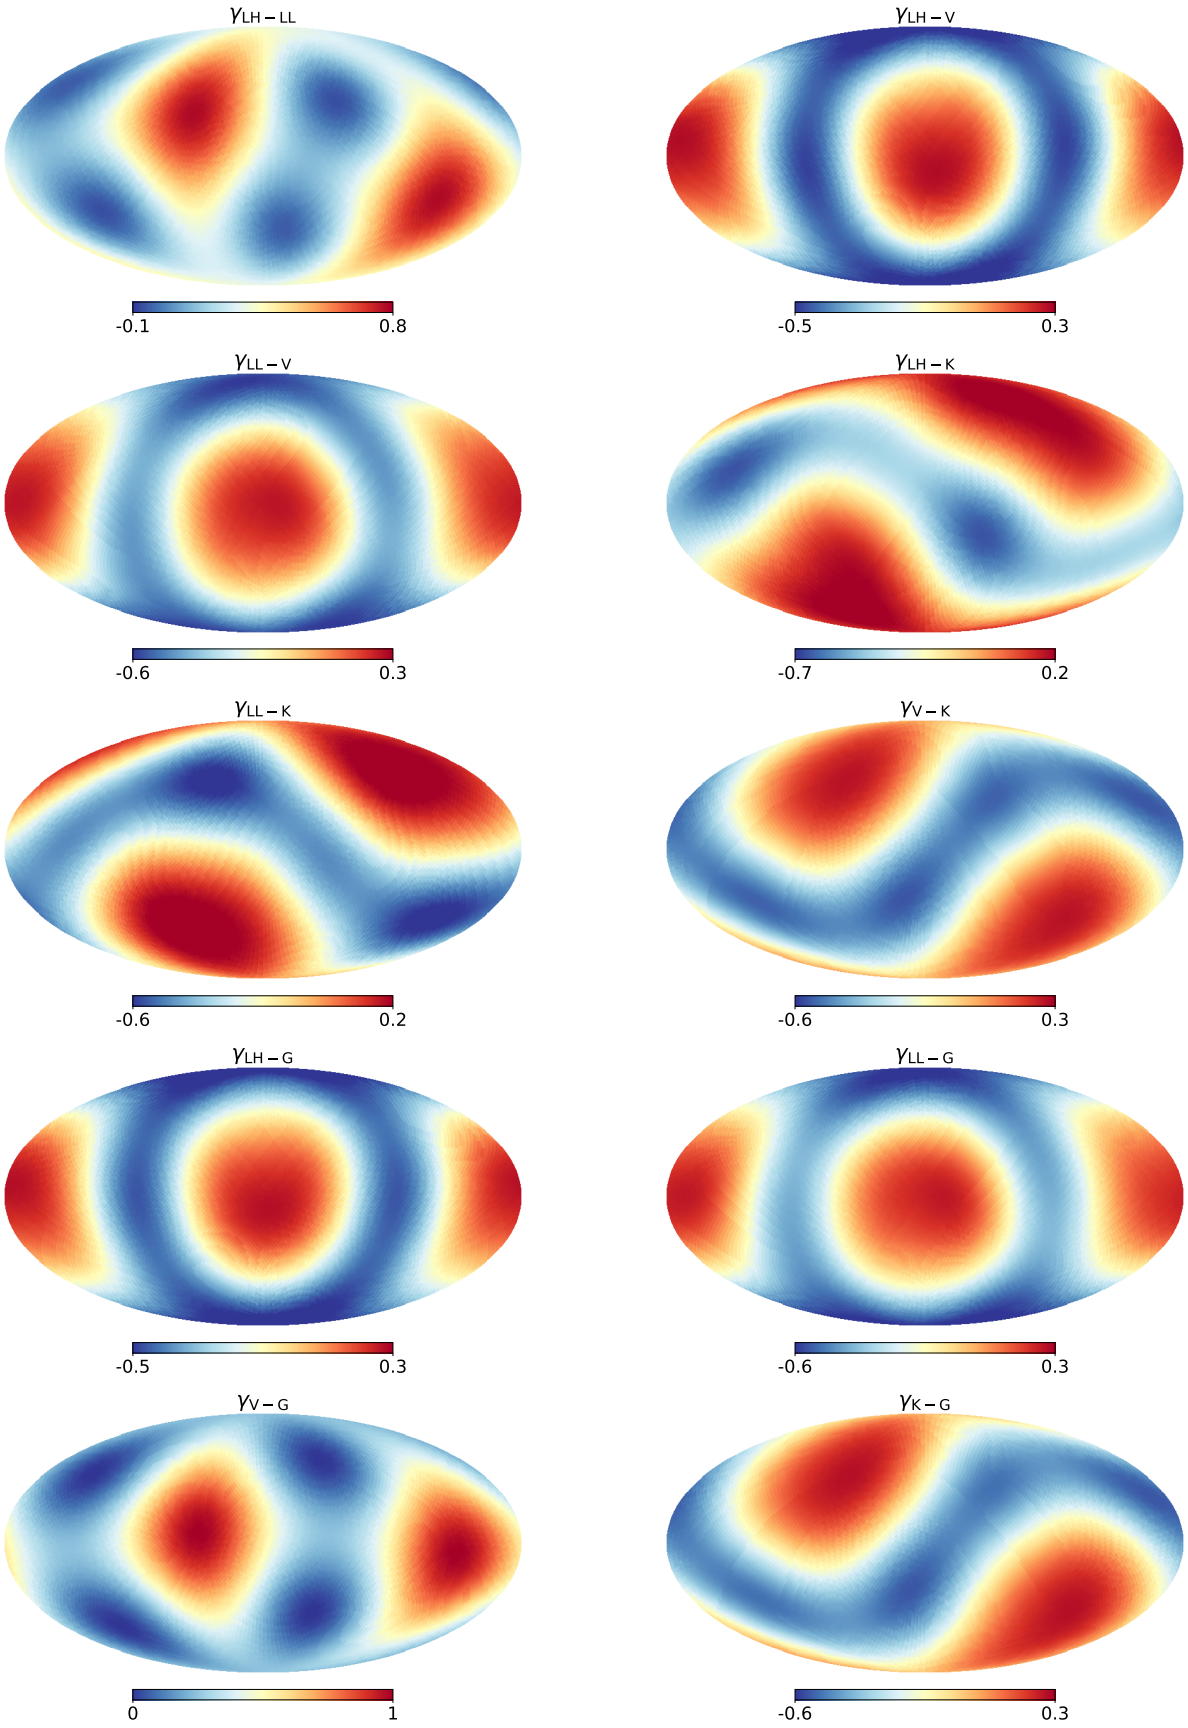


Figure 4.3: Instantaneous overlap function $\gamma(\hat{\boldsymbol{n}})$ for ten different baselines constructed by pairing the detectors listed in Table 4.1, in galactic coordinates. These may be interpreted as instantaneous beams and shows how the strain intensity is integrated at an arbitrary phase of the Earth's rotation.

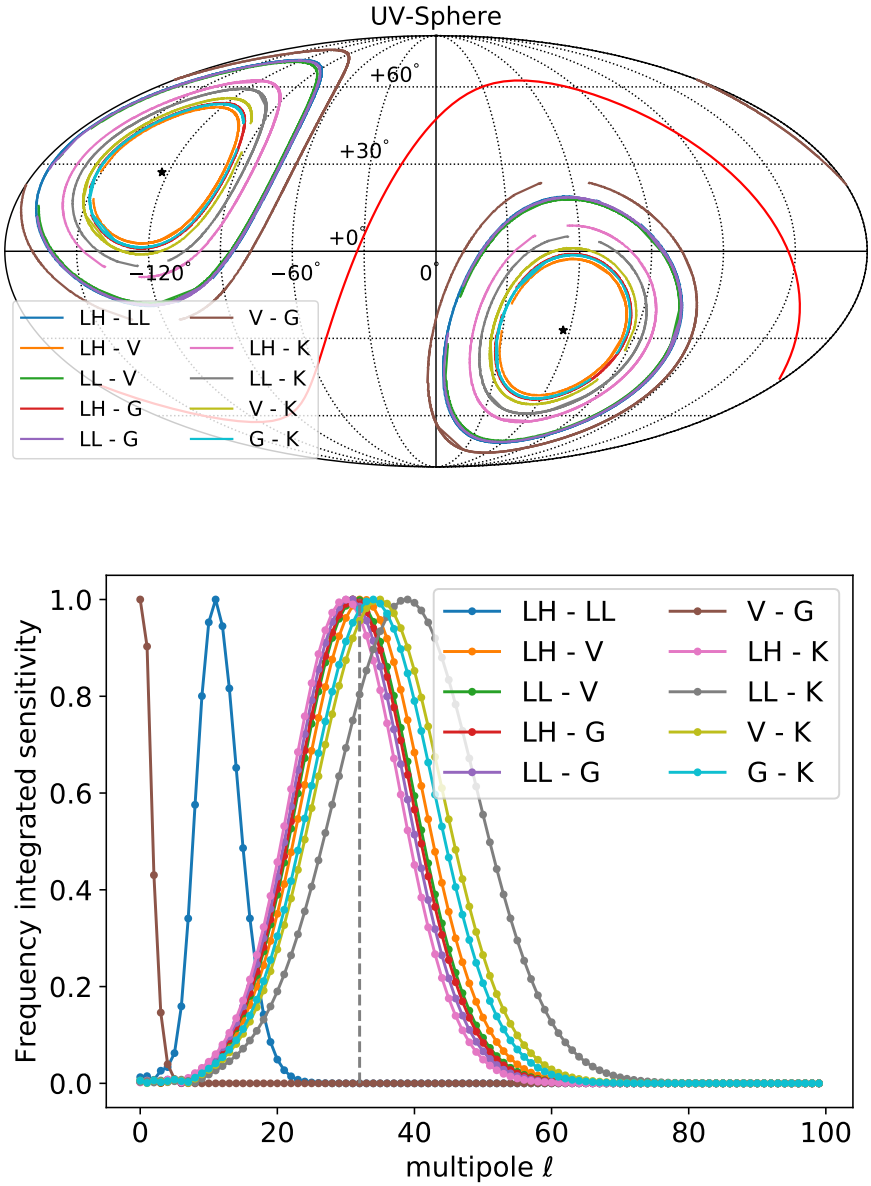


Figure 4.4: Angular phase coverage (top panel) of the 10 baseline case plotted in galactic coordinates. Each baseline track is plotted symmetrically about the celestial equator (red, solid curve). The tracks span 48 hours of simulated data following the segmentation and flagging of the actual "O1" LIGO data - only small gaps in the angular phase coverage of each baseline remain. An approximate view of the angular scale sensitivity of each baseline is also shown (bottom panel). The grey (dashed) vertical line at $\ell = 32$ gives an indication of the limiting resolution in the reconstructed maps at given by the working pixelisation scale.

This estimate does not include the effect of the coupling kernel $\mathcal{K}^{LM}_{\ell m,\ell' m'}$ in Equation (4.12). The coupling significantly complicates the interpretation of the baseline resolution in terms of an overall resolution scale of any final map. For example, the $3j$ coefficients in Equation (4.13) correlate a multipole ℓ with multipoles ℓ' in the range $|\ell - \ell'| \leq 2\ell$. This implies that the only correct way to estimate sky sensitivity for GW detectors is through *mapping*, as introduced at the end of Chapter 3, as any m–compactification of the data which skips this step will ignore these correlations.

In the tests which follow, the O1 release aLIGO time stream data [117] was used as a template for the segmentation and flagging of all the simulated runs. The data contains gaps for periods where either detectors forming each baseline were off-line or for periods when the data is flagged in such a way that it cannot be used for analysis. Some of these gaps can be seen in the baseline tracks in Figure 4.4 that cover a single 48 hour period. Since the gaps are not expected to be particularly correlated with diurnal phase, their overall effect should diminish as integration time increases beyond a few days.

4.2 Optimal map solutions: generation and testing

In searches for stochastic signals with ground-based detectors, it is widely assumed that the local noise is uncorrelated between detectors. In fact, even though the data streams in interferometers such as LIGO are heavily dominated by noise and instantaneous random correlations are highly probable, when performing stochastic analyses these will average down over time, as the assumption made is that the ensemble averages of the correlated noise components are zero,

$$\langle D_{\boldsymbol{b}} \rangle = \boldsymbol{A_b} \langle h\, h^{\star} \rangle + \langle n_A\, n_B^{\star} \rangle \approx \boldsymbol{A_b} I\,, \tag{4.17}$$

as explained in Chapter 3.1. Here A and B label the detectors which identify the baseline $\boldsymbol{b}$, as in the sections above, and $\boldsymbol{A_b}$ is the generalised quadratic response function obtained by integrating $\gamma_{\boldsymbol{b}}$ over the sky, with appropriate normalisation and exponential factor (see Equation 4.4). The appropriate mapping method in the case of ground-based detectors is then the one presented in Chapter 3.1 for uncorrelated data streams. Furthermore, in the

case of aLIGO data [105] the timestream is highly dominated by the noise component, hence the noise-dominated, closed-form solution presented in Equation 3.10 may be used.

In this section the general mapping equation is tailored to the case of ground-based detectors, and results from mapping simulations are presented. The simulated data is generated directly in the frequency domain as cross-correlated data from pairs of ground-based interferometers. Equivalently, each data stream may be thought of as the read-out from a single baseline. Mimicking the treatment of real aLIGO data in [118], the data is generated for consecutive 60-second time segments throughout which the rotation of the baseline with respect to the galactic frame is negligible. In the case of real data, this reduces the cost of Fourier transforming, excludes frequencies that are far into the $1/f$ tail and ensures that the noise properties are stationary within each segment. The `HEALPix`[2] [119] hierarchical pixelisation scheme is used which discretises the sky into equal area elements. The FFT time-scale is crucial as it sets the upper limit of the pixelisation scale: throughout this time, each direction vector $\hat{\boldsymbol{n}}$ will point to the same pixel. To avoid edge effects the scale is chosen such that each pixel is over-sampled in consecutive time segments. The equal area scheme allows to discretise the observation Equation (4.4) for a single baseline vector $\boldsymbol{b}^\tau$, at time segment τ, as

$$D_f^\tau = E_f \Delta_f \frac{4\pi}{N_{\rm pix}} \sum_p I_p \gamma_p^\tau e^{i2\pi f \boldsymbol{b}^\tau \cdot \hat{\boldsymbol{p}}} + n_f^\tau, \tag{4.18}$$

where $N_{\rm pix}$ is the total number of sky pixels, Δ_f is the discretisation interval in frequency, I_p is the intensity on the sky at pixel p with unit direction $\hat{\boldsymbol{p}}$, and γ_p^τ is the discretised overlap function. The notation highlights that in the galactic frame adopted here the baseline vector and overlap function are both rotated with respect to the sky at each different time segment τ. This expansion assumes that all of the sky signal has the same frequency dependence E_f, which factors out from the sky dependence all together. This will be assumed in all the simulations in this section, however in reality this will not be the case and the model will have to account for a number of sky signals, each having a distinct spectral dependence. For a broad band analysis such as that carried out here one would have to solve for multiple maps. Alternatively, an analysis where the data is separated into narrow frequency bands could deal with multiple signals through a later component separation stage. An example of this will be

[2] http://Healpix.sourceforge.net

shown in Section 4.3, where narrowband maps from the aLIGO O1 and O2 data releases are shown.

The projection $\boldsymbol{z} = \boldsymbol{A}^\dagger \boldsymbol{N}^{-1}\boldsymbol{D}$ required in Equation (3.11) may now be identified with the discretised form as

$$\boldsymbol{z} \rightarrow z_p^\tau = \Delta_f \frac{4\pi}{N_{\rm pix}} \sum_\tau \sum_f \gamma_p^\tau \frac{E_f d_f^\tau}{P_f^\tau} e^{-i2\pi f\, \boldsymbol{b}^\tau\cdot\hat{\boldsymbol{p}}}, \tag{4.19}$$

and it is possible to identify the quadratic response operator $\boldsymbol{A}$ as

$$A_{f,p}^\tau = \Delta_f \frac{4\pi}{N_{\rm pix}} \gamma_p^\tau E_f\, e^{-i2\pi f\, \boldsymbol{b}^\tau\cdot\hat{\boldsymbol{p}}}. \tag{4.20}$$

It is important to note that the reality of the observed timestreams implies the condition $D_f^\star = D_{-f}$ such that the operation can be carried out as a sum over the positive frequency domain only by adding the complex conjugate terms. In this case it is useful to carry out the computation directly and obtain a simple expression for z_p as

$$\begin{aligned} z_p^\tau =& \Delta_f \frac{4\pi}{N_{\rm pix}} \sum_\tau \sum_{f=0}^{\infty} \gamma_p^\tau \frac{E_f d_f^\tau}{P_f^\tau} e^{-i2\pi f\, \boldsymbol{b}^\tau\cdot\hat{\boldsymbol{p}}}\, d_f^\tau] + \text{c. c.} \\ =& \Delta_f \frac{8\pi}{N_{\rm pix}} \sum_\tau \sum_{f=0}^{\infty} \gamma_p^\tau \frac{E_f d_f^\tau}{P_f^\tau} \left[\cos(2\pi f\, \boldsymbol{b}^\tau\cdot\hat{\boldsymbol{p}})\mathbb{R}\left(d_f^\tau\right) - \sin(2\pi f\, \boldsymbol{b}^\tau\cdot\hat{\boldsymbol{p}})\mathbb{I}\left(d_f^\tau\right)\right]. \end{aligned} \tag{4.21}$$

Similarly we can identify the beam pattern $\boldsymbol{M} = \boldsymbol{A}^\dagger \boldsymbol{N}^{-1}\boldsymbol{A}$ in Equation (3.12) as

$$\begin{aligned} \boldsymbol{M} \rightarrow M_{pp'}^\tau &= \Delta_f^2 \frac{16\pi^2}{N_{\rm pix}^2} \sum_\tau \sum_{f>0} \gamma_p^\tau \gamma_{p'}^\tau \frac{E_f^2}{P_f^\tau} e^{2\pi f\, \boldsymbol{b}^\tau\cdot(\hat{\boldsymbol{p}}-\hat{\boldsymbol{p}}')} + \text{c. c.} \\ &= \Delta_f^2 \frac{16\pi^2}{N_{\rm pix}^2} \sum_\tau \sum_{f>0} \gamma_p^\tau \gamma_{p'}^\tau \frac{E_f^2}{P_f^\tau} \cos\left[2\pi f\, \boldsymbol{b}^\tau\cdot\left(\hat{\boldsymbol{p}}-\hat{\boldsymbol{p}}'\right)\right]. \end{aligned} \tag{4.22}$$

When the observations cover multiple times τ and baselines $\boldsymbol{b}$ Equations (3.15) an (3.16) apply and the expressions are modified by summing over the individual contributions in such a away that

$$z_p = \sum_{\boldsymbol{b},\tau} z_p^{\boldsymbol{b},\tau}, \tag{4.23}$$

and

$$M_{pp'} = \sum_{\boldsymbol{b},\tau} M_{pp'}^{\boldsymbol{b},\tau}, \tag{4.24}$$

where the superscripts denote the individual contributions to each baseline. The maximum-likelihood map estimate is then given by

$$\tilde{I}_p = \sum_{p'} M^{-1}_{pp'} z_{p'} \,. \tag{4.25}$$

The calculation of $z^{\boldsymbol{b},\tau}$ and $M^{\boldsymbol{b},\tau}$ will involve a rotation of the baseline dependent overlap function $\gamma^{\tau,b}_p$ maps and baseline vectors $\boldsymbol{b}^\tau$ to the galactic frame. The `QPoint`[3] library of [120] is employed to apply the location dependent transformation from detector to sky frames using quaternions. As these keep track of also the individual orientation of the polarisation basis for each direction, this library will be essential for the polarised GWB analysis in Section 4.2.3. Furthermore, this generalised treatment of the directionality which does not require fixing the basis to a detector-dependent frame is extremely useful as the algorithm is easily extended to include the multiple baseline scenario.

For real data, the estimate of the power spectrum of the noise in each baseline P^τ_f will, in general, be different for separate time frames τ due to long-term drifts in detector characteristics. To account for this the noise spectra are re-estimated for each time segment, as explained below.

4.2.1 Data generation

In order to test this procedure in the most realistic setting from the outset the algorithm has been based around the LIGO open data format. The aLIGO O1 data frames are sequentiallly loaded and stored with a sampling rate of 4096 Hz, where both LH and LL detectors were simultaneously in operation. When simulating data from more detectors these are assumed to be operational for the same data frames. The main use of this procedure is to extract a realistic noise curve from the real data to use for data generation, however the set-up of this pipeline is also extremely useful for the actual data analysis which follows in Section 4.3.

The data is parsed using the LIGO flag system in order to reject frames where injection events were present, periods when the detectors were being tested, or sections of the data were the LIGO pipeline has verified that one or either of the instruments was operating

[3] http://github.com/arahlin/qpoint

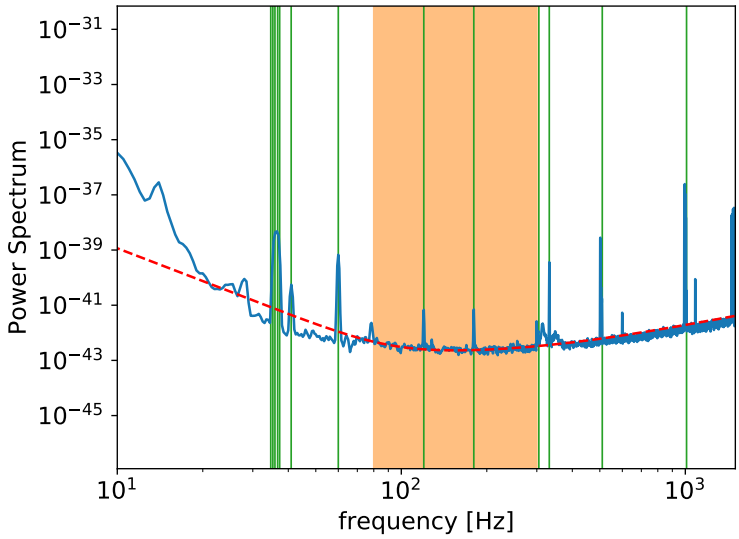


Figure 4.5: LIGO LL, smoothed power spectrum for a single 60 second time segment. The vertical lines show the positions of the notches applied to the frequency data in order to remove known harmonics. The dashed curve shows the model fit to the spectrum within the frequency range shown in the shaded area. The fit is carried out for every 60 second time segment to all detector timestreams. When fitting the spectrum the data is gapped around the notched frequencies to avoid a biasing the weighting function used in the map–making and in the estimate of the pixel noise covariance matrix. The frequency range adopted here is a conservative one that limits $1/f$ and f tail contributions that can fluctuate significantly between segments of the data.

outside of the established parameter range. The remaining data are then subdivided into segments of 60 second duration. These constitute the time segments labelled as τ where, for the angular scales being targeted here, the motion of the Earth may be considered to be negligible. Each segment τ is edge–tapered with a cosine window over the first and last three seconds before Fourier transforming to the frequency domain. The frequency data is then notch filtered to remove biases due to known harmonics (see Table 4.2 in Section 4.3 for details) and band–passed to the range $f \in [80.0, 300.0]$ Hz. This procedure is applied to all detector timestreams involved in the number of baselines being considered. The data segments are analysed at the native aLIGO open data down sampled rate of 4096 Hz.

A power spectrum is computed from the filtered data and a three parameter model[4]

$$P(f) = \mathscr{A}\left[\left(\frac{\mathscr{B}}{0.1+f}\right)^4 + \left(\frac{0.04f}{\mathscr{C}}\right)^2 + (0.07)^2\right], \tag{4.26}$$

[4]from the GWOSC website: https://www.gw-openscience.org

is fitted to the data spectra, within the band–pass, to obtain a functional form of the noise spectra of each detector timestream. The function in Equation (4.26) which may be observed in Figure 4.5 is simply a concave curve with an offset $\mathscr{A}$ and high and low knee frequencies which enclose the minimum, identifying the frequencies where the detector is most sensitive. The parameters $\mathscr{B}$ and $\mathscr{C}$ are directly related to the knee frequencies through an appropriate scaling. To avoid any residual bias due to harmonic lines the data is gapped around known harmonics when fitting with a width of 5σ either side of the notch. The notched frequencies and corresponding σ values are presented in the next Section, where the noise properties of the data are analysed in detail. Since LIGO data are noise dominated the signal is not expected to bias this noise estimation procedure. It is worth noting however that at much higher signal-to-noise ratio (SNR) such as that expected for LISA this step will require an iterative procedure to obtain a noise spectrum that is not biased by any signal contribution, as discussed in Chapter 5. If the fit fails, or the parameters are outside a given range, the time segment is discarded. The failure typically happens for time segments where the $1/f$ tail is larger than expected or if there is an unexpected line present in the spectra. The product of the fitted models is then used to obtain the weighting function P_f used in the maximum-likelihood solution,

$$P_f^{\tau,\boldsymbol{b}} = P_f^{\tau,A} P_f^{\tau,B}, \tag{4.27}$$

where A and B denote the two detectors defining the baseline labelled $\boldsymbol{b}$. This is effectively the product of the two noise power spectra in the baseline, and is thus a power of four in the GW strain. Figure 4.5 shows the spectral density for a random segment τ of aLIGO Livingston O1 data along with the notching centres and resulting fit. The data segments are band-passed in order to exclude the large $1/f$ and f-noise tails.

To generate mock-data, after all filtering steps, the frequency domain data is substituted with a scan of an input intensity map $\boldsymbol{I}^{\rm in}$. This is obtained using the transposition of the projection operator

$$\begin{aligned} D_f^{\tau,\boldsymbol{b}} &:= \sum_p A_{f,p}^{\tau,\boldsymbol{b}} I_p^{\rm in} + N_f^{\tau,\boldsymbol{b}} \\ &= \Delta_f E_f \frac{4\pi}{N_{\rm pix}} \sum_p \gamma_p^{\tau,\boldsymbol{b}} I_p^{\rm in} e^{i2\pi f \boldsymbol{b}^\tau \cdot \hat{\boldsymbol{p}}} + N_f^{\tau,\boldsymbol{b}}. \end{aligned} \tag{4.28}$$

This procedure could be extended to obtain fiducial simulations of the time domain data by transforming back to a timestream in order to add transient systematics. As the stochastic analysis is carried out in the frequency domain where the noise is also assumed to be uncorrelated this is not motivated at this stage. The noise component $N_f^{\tau,\boldsymbol{b}}$ is included by adding an uncorrelated, Gaussian realisation with variance given by the fitted cross-correlation power spectrum P_f^{τ}, at each frequency and time segment. The final product, at each 60 second time segment τ, is a set of simulated signal plus noise frequency data for each baseline. All baselines will have the same noise variance given by the actual aLIGO data noise spectra for each time segment and gaps in the simulations will all mimic those resulting from missing or flagged aLIGO data.

To avoid any aliasing in the simulation procedure the input map $\boldsymbol{I}^{\rm in}$ is generated at a resolution level $N_{\rm side} = 32$ such that the structure in the input is over–resolved with respect to the working resolution of the output maps which is limited by the smoothing scale of most baselines (see Figure 4.4 for reference).

4.2.2 Mapping tests

The map $\boldsymbol{z}$ and operator $\boldsymbol{M}$ are accumulated over each time segment τ and baseline $\boldsymbol{b}$ in the combination being considered. A flat spectral dependence $E_f = 1$ is considered here for all the cases, as these are simple tests of the algorithm and the injected signals are not realistic. The algorithm is parallelised using MPI to distribute individual time segments τ between processors. Since the time segments are of equal length, the parallelisation is well balanced and scales close to linearly with number of MPI processes. Once all the time segments have been analysed Equation (4.25) is used to obtain the final maximum–likelihood map estimate $\tilde{I}$.

To interpret the signal-to-noise of features in the map it is necessary to quantify the pixel noise covariance. This is given by the inverse of the operator $\boldsymbol{M}$, as the working assumption is that the data is Gaussian [121]. In practice however, the operator $\boldsymbol{A}$ here introduces a relative calibration between the signal and noise in Equation (3.1) due to the rescaling induced by the time–domain windowing of each segment τ. This rescaling cancels out in the maximum–likelihood solution (3.9) but needs to be accounted for in calculating the noise covariance

of the map [106]. The scaling factor is evaluated by integrating the *unweighted* operator $\sum_{f,\tau,b} A^{\tau,b}_{fp}$ and is used to renormalise the noise matrix. The map of noise standard deviation

$$\sigma_p = \left[M^{-1}_{pp}\right]^{1/2} \tag{4.29}$$

is also estimated for visual comparison, although note that the correlations in the final maps are non-trivial and σ_p is only a guide for the interpretation.

A number of different choices for input maps are used to test the reconstruction with three separate amplitude vs. noise levels (ANL), i.e. the ratios between the input amplitudes of the signal and the noise:

- **High ANL**, corresponding to a noiseless case with input strain intensity of $h^2 \sim \mathcal{O}(1)$.
- **Medium ANL**, corresponding to a signal level comparable to the typical noise level of the aLIGO baseline. This is chosen such that the maximum–likelihood maps converge after analysis of $\mathcal{O}(1)$ day of data.
- **Low ANL**, where the signal level is chosen such that the estimated maps are still converging after similar time scales.

Note that the high ANL does not bias the noise estimation step described above as the noise is estimated using the actual LIGO data for each segment before it is substituted with the simulated scan.

For each of the ANL cases different types of input maps are considered here. The first are pure **monopole** maps which contain no angular structure. The second case are statistically isotropic **Gaussian** random realisations of an $\ell(\ell+1)C_\ell = const.$ angular power spectrum, up to $\ell_{\rm max} = 8$. This value is lower than the expected limiting resolution for the LIGO baseline as seen in Figure 4.4, however it is chosen as it is sufficiently low that these tests may be carried out quickly with the computing power available to us. The third case are anisotropic **Galaxy** maps obtained by smoothing the dust component Planck satellite map[5] [122]. These are dominated by galactic dust emission and provide an input map where the signal is concentrated along the galactic equator. These are not supposed to simulate any realistic

[5] https://pla.esac.esa.int

GWB but are instead designed to test the algorithm with various SNR regimes and levels of anisotropy. Note that this last set will break the Gaussian assumption for the spatial distribution of the signal, which in principle implies that the estimator will be sub-optimal. A non-stationary case is also examined, where a **Poisson** process is simulated at each time segment τ and in each pixel. This generates input signals of a fixed amplitude according to a Poisson random draw with the mean of the Poisson distribution set to the **Galaxy** input map. The aim of this particular test is to show how the signal of a stochastic, non-stationary background, that may be arriving from distinct directions on the sky simultaneously, is accumulated in the map–making procedure.

Figure 4.6 shows the results obtained for a number of high ANL input maps after the integration of approximately 24 hours of mock LH – LL – V data made up of baseline formed by LIGO Hanford and Livingston detectors, and Virgo, as set out in Table 4.1. A full Earth rotation is sufficient to reconstruct all $\ell \leq 8$ modes included in the statistically isotropic cases accurately, although the effect of gaps in the LIGO data and the limited sky coverage of the simulated LH – LL – V baseline combination is seen in the induced pixel correlation that is most evident in the **monopole** case. The limited resolution of the observations, dictated by the angular scales of the overlap functions, is seen most clearly in the **Galaxy** case where the higher resolution input features are significantly smoothed by the convolution. As for the **Gaussian** input, all $\ell \leq 8$ modes are recovered accurately.

The results for the low ANL cases are shown in Figure 4.7. The signal level for these tests was set so that the residual noise in the maximum–likelihood map is comparable to the signal. As is evident in all three output maps, the residual noise modes add a ripple effect to the reconstructions, at particular angular scales. This is especially evident in the **monopole** case, as the true map is simply constant on the sky. The value of the monopole in this case is biased by the noise by roughly 5%. In the **Gaussian** case it is interesting to observe how the noise disturbs the patches of map where the intensity is closest to zero, as may be expected, adding higher angular structure, whereas the peaks and valleys of the map remain in the correct locations. Finally, in the **Galaxy** output map it is still possible to see the plane of the Galaxy as in the high ANL test, however the fluctuations off the plane have grown larger.

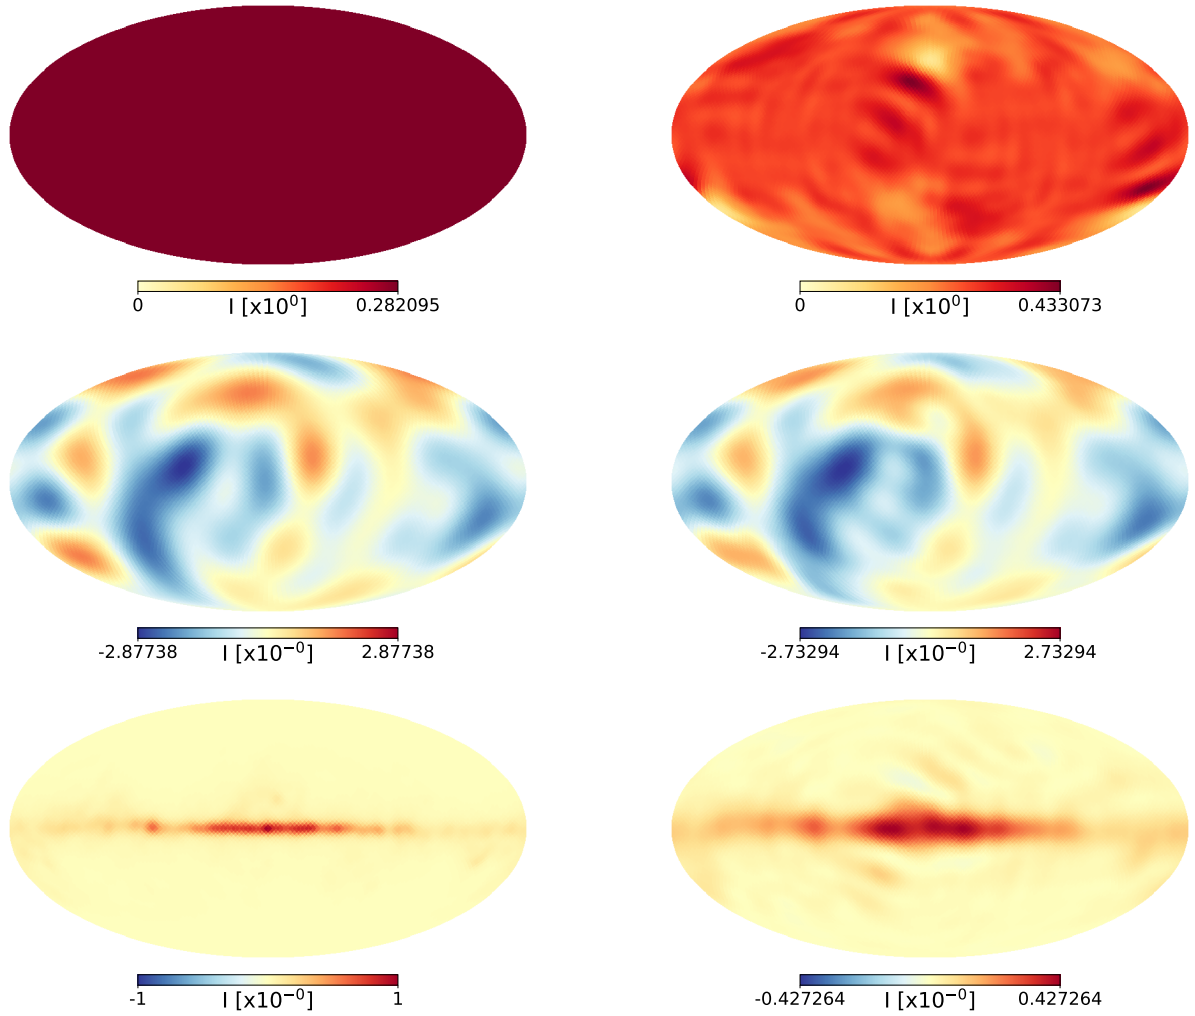


Figure 4.6: Input (left panels) and output (right panels) maps for the high ANL cases. The results are for the analysis of 24 hours, simulated data, mimicking the aLIGO O1 release data structure for the three baseline (LH – LL – V) case. Top: The input **monopole** map with amplitude $1/2\sqrt{\pi} \sim 0.282$ corresponding to a monopole $a_{00} = 1$. The output map recovers the monopole accurately within 0.4%, however the non-trivial "uvw" coverage results in significant correlations in the final map which appear as fluctuations around the mean. Middle: The **Gaussian** statistically isotropic map case. All $\ell \leq 8$ modes are recovered accurately. Bottom: The anisotropic **Galaxy** case. The resolution limit of the observations is most obvious in this case. The maps are reconstructed at a lower pixelisation of $N_{\rm side} = 8$ to lighten the computational load, and are then smoothed with a 10 degree Gaussian beam and over–resolved to $N_{\rm side} = 32$ for presentation purposes.

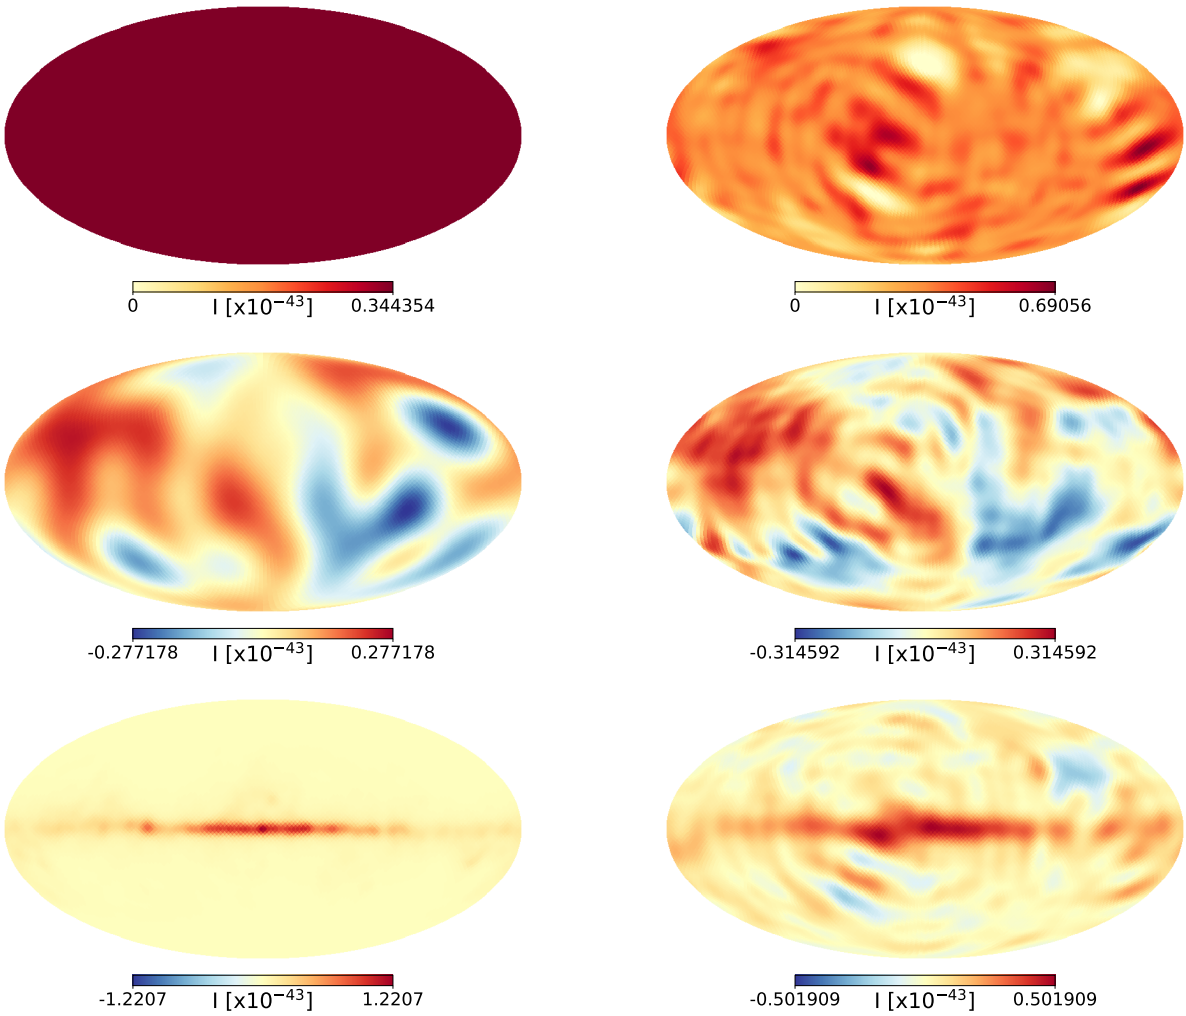


Figure 4.7: Input (left panels) and output (right panels) maps for the low ANL cases for the same baseline combination as in Figure 4.6. The signal level in these simulations was chosen so that the ANL is close to unity for the integration time.

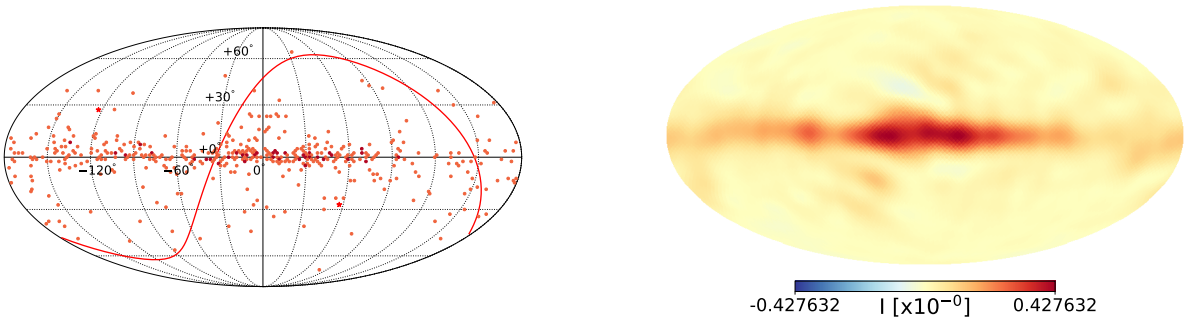


Figure 4.8: The high ANL non–stationary **Poisson** Galaxy input case for approximately 36 hours of analysed mock LH – LL – V data. Accumulated input Poisson map (left panel) in a single minute of the simulation and recovered map (right panel). The increased integration time compared to the other tests is required to accumulate the signal of enough "events". After a sufficient number of "events" the map converges to the stationary case output map seen also in Figure 4.6.

An example of the non-stationary **Poisson** tests may be seen in Figure 4.8. The baseline set-up is the same three-baseline, LH – LL – V as in Figures 4.6 and 4.7. In the right panel, an example of the input for a single minute of the simulation is given. Each dot here represents an "event" which is normalised to unit amplitude and the colour scale indicates the number of "events" accumulated in each pixel (maximum four in this case). As may be seen by eye, this is an incredibly unrealistic case for the detectors considered, as the event rate required for such a sample is orders of magnitude above the expected rate in this frequency band and at this sensitivity. This is discussed from a theoretical standpoint in Chapter 2.

A visual comparison between low and high ANL regimes can be made by comparing the residual maps and the power spectrum reconstruction in the two cases. Figure 4.9 shows the residual maps obtained by subtracting the input and output maps for the **Gaussian** tests shown above for the two ANLs. Note that, in the high ANL case, the pixel residual fluctuation is of the order of 25% of the amplitude of the input map, whereas in the low SNL case it is of the order of the map itself which clearly indicates the output map has not converged yet. Figure 4.10 shows the input and reconstructed C_ℓ power spectra obtained in the two **Gaussian** cases. It appears clear that in the absence of noise the GWB power on the sky is very well reconstructed, at $\ell < 8$, and there is no residual leakage to the higher modes. In the presence of noise instead there is strong aliasing between modes and an irreducible high-ℓ residual due to the pixel-pixel noise is present.

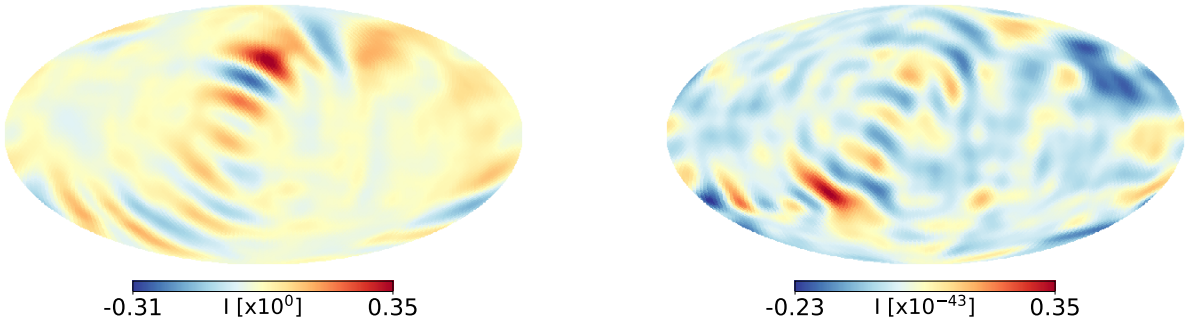


Figure 4.9: Residual maps in high (left panel) and low (right panel) ANL regimes for the **Gaussian** tests shown in Figures 4.6 and 4.7. Note that, in the high ANL case, the pixel residual fluctuation is of the order of 25% of the amplitude of the input map, whereas in the low SNL case it is of the order of the map itself. Also note how the pattern that emerges traces the pattern of the **monopole** reconstructions, especially visible in Figure 4.7, which further indicates that these regions of the sky are more poorly constrained by the detector network.

For current generation computing cores at a sampling rate of 4096 Hz, this algorithm can analyse 0.033 days of baseline data per core day with the cost scaling linearly in τ and b. To understand the convergence of the map estimation it is useful to look at the condition number of the operator $\boldsymbol{M}$ in Equation (4.25) and the standard deviation of the output maps as a function of amount of data analysed (in minutes). A large condition number, typically greater than $\mathcal{O}(10^8)$, is an indication that too many modes are trying to be reconstructed for the given observations, hence the problem is ill–conditioned. The standard deviation of the output maps should decrease approximately as the square root of the amount of time analysed as the noise is integrated down. A convergence of the standard deviation of the map to a constant value indicates that the map has converged to a given signal level which does *not* average down with additional integration in time.

Figure 4.11 shows the condition value of $\boldsymbol{M}$ for the single, three, and ten baseline case run on the high ANL **Gaussian** input case. The detectors and baselines considered are the same as those discussed in Section 4.1. The condition number is initially large for all cases but decreases rapidly over the course of 24 hours of integration time except for the single baseline case. In practice, the inversion in Equation (4.25) is conditioned using a pseudo inverse calculation with a limiting condition number of 10^5. After 24 hours only the single LL-LH baseline case requires the removal of singular modes. This is an indication that for the single baseline case the working $N_{\rm side} = 8$ resolution is sufficient and saturates the number

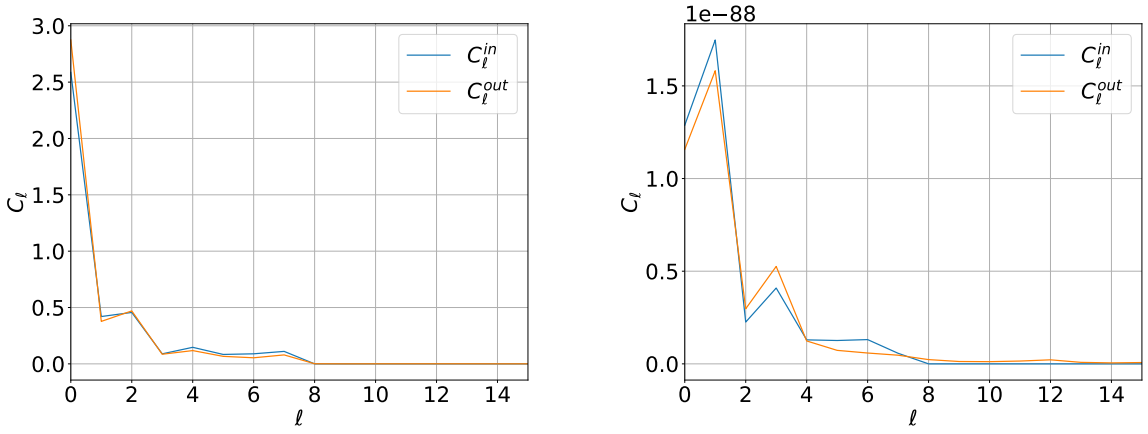


Figure 4.10: Comparison between input and reconstructed C_ℓ power spectra in high (left panel) and low (right panel) ANL regimes for the **Gaussian** tests. These clearly show how precise the sky power reconstruction is in the absence of detector noise, with a clean C_ℓ reconstruction and very little ringing and leakage. Conversely the low ANL case displays mode-to-mode power leakage and in particular there is a residual tail of power at higher ℓs which is a symptom of the irreducible pixel-pixel noise.

of modes that can be reconstructed. The addition of more baselines increases the number of modes that can be reconstructed and the fact that the inversion is carried out without removing singular modes is an indication that this resolution does not saturate the number of modes and could be increased further.

Figure 4.12 shows the evolution of the average standard deviation of the maps $\overline{\sigma} = \sigma_p$ as a function of integration time for various ANL regimes and cases. It is clear that, for the first few hours, the convergence is affected by the ill–conditioning of the problem. However, after a approximately 6 hours of integration time, the standard deviation enters a scaling regime that is approximately close to the $t^{-1/2}$ limit expected until, for the higher ANL cases, it reaches a constant value indicating that the map has converged. The standard deviation also scales roughly as $b^{-1/2}$, given identical signal inputs. This is as expected for the case where each baseline contributes with the same sensitivity, with uncorrelated noise.

Finally, it is interesting to represent the map of noise standard deviation as defined in Equation (4.29) for different baseline combinations. As more accurate, time-integrated maps will be shown in Section 4.3 for the LIGO LL – LH baseline as a result of the full data analysis of the O1 and O2 data sets, in this section it is sufficient to include Figure 4.13 which shows σ_p in the six baseline case (LH-LL-V-K), after 24 hours of integration. Despite the large number

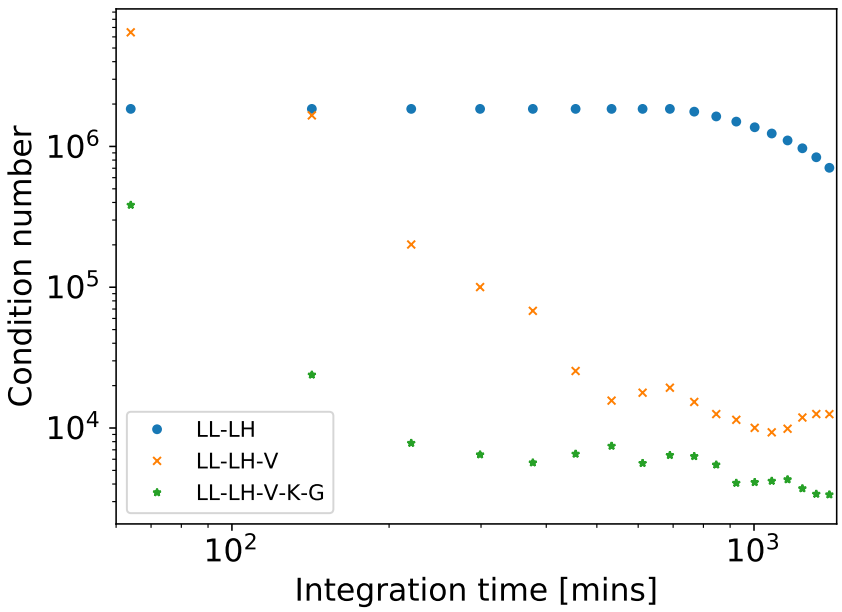


Figure 4.11: Evolution of the condition number of the operator $\boldsymbol{M}$ with integration time for a number of different simulated baseline combinations (single, three, and ten baselines). The conditioning of the maximum-likelihood problem improves rapidly over the first 12 hours except for the one baseline (LL – LH only) case.

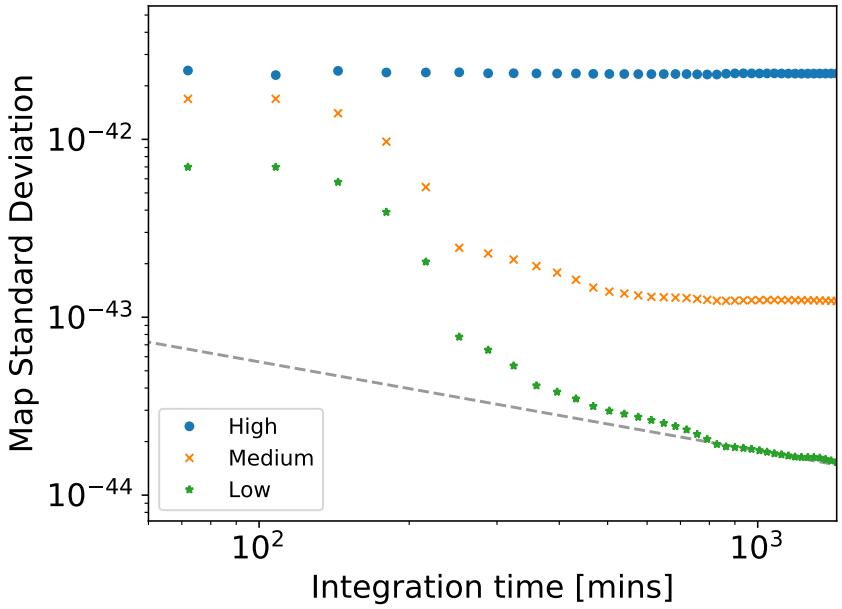


Figure 4.12: Map standard deviation for the LIGO-VIRGO, three baseline simulations with different ANL. The scaling of the standard deviation is initially affected by the ill–conditioning of the problem but enters the expected $t^{-1/2}$ scaling once the rotation of the Earth has improved the conditioning, as illustrated by the grey dashed line. When the standard deviation becomes constant the maps have converged to a detection of the input signal. The only case that has not converged after a full Earth rotation is the low ANL one.

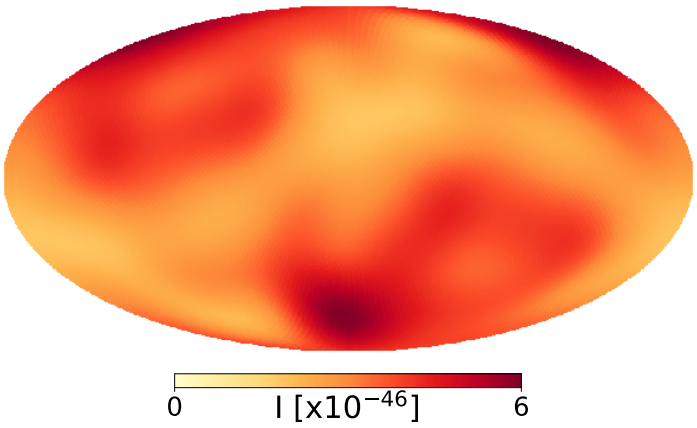


Figure 4.13: Noise map for the six baseline case after 24 hours of integration. The noise is inhomogeneous and there are significant pixel–pixel correlations that are not visualised in the map.

of baselines the noise remains significantly inhomogeneous.

These tests serve as a proof of concept that the mapping method works and is efficient. However, the simple input maps used to simulate the signal component are not of much use in evaluating the potential for observations to detect either the monopole or anisotropic component of a stochastic background. A realistic simulation would require an injection of a properly simulated event rate distribution in the timestream of the various detectors. The first step would be to build a robust simulator of stochastic events from particular populations with the correct correlation on the sky. This work will also require a robust simulation of the time domain noise which entails additional complications with respect to the frequency domain approach taken in this work. This type of work is probably premature, given the current sensitivity of the aLIGO, Virgo and Kagra instruments and the expected event rates from astrophysical populations, as discussed in [123].

The method, with suitable redefinitions of the coordinate transformations and observing operators, can also be applied to PTA data. These provide much longer baselines than those achievable by laser interferometers and could be the first to detect a GWB [15]. However, the array of pulsars is not expected to “scan” the signal, as the “detector” is not moving with respect to the fixed background, hence the time-dependent framework built here is not necessary.

4.2.3 Polarisation maps

This Section explores the possibility of extending the mapping analysis to GWB polarisation reconstruction, and solving for polarised backgrounds. As explained in the previous Chapters, the measurement of a stochastic GW signal may only be carried out in the square of the strain, averaging out the phase information; this is true also for polarised stochastic backgrounds hence this extension aims to include directly the Stokes' parameters $W = Q, U, V$ in the mapping algorithm.

A full polarisation analysis of the GWB would prove useful in the event of a positive background detection, as at the very least it would provide a check of the whole apparatus. Effectively, a stochastic background from astrophysical sources should be unpolarised, hence any spurious Q, U or V component would indicate an error in the analysis or a bias in detector calibration. As for cosmological backgrounds, there are models that expect a circularly polarised signal; see the brief discussion in Chapter 2.

To start, it is instructive to check how the overlap functions γ^W for each of the Stokes' parameters transform under coordinate changes. Specifically, as I and V are spin-0 fields and $Q \pm iU$ are spin-4, the same symmetry is expected to apply to the respective transfer functions. Changing basis will result in 2 rotations: one is equivalent to a coordinate change, e.g. re-writing the detector arm vectors $\boldsymbol{u}_A$, $\boldsymbol{v}_A$ with respect to a new basis, and the remaining rotation is around the local zenith which accounts for the choice of the polarisation basis. This will be the rotation $R(\hat{\boldsymbol{n}}, \Psi)$ which in general will depend on $\hat{\boldsymbol{n}}$ as it results in a rotation around $\hat{\boldsymbol{n}}$ on the sky. For example, if one chooses aligned coordinate systems, such that $\hat{\boldsymbol{n}} \equiv \hat{\boldsymbol{z}}$ in the new frame, then simply

$$R(\hat{\boldsymbol{z}}, \Psi) = \begin{pmatrix} \cos 2\Psi & -\sin 2\Psi & 0 \\ \sin 2\Psi & \cos 2\Psi & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{4.30}$$

The spin-2 property of the polarization basis $e^+, e^\times$ implies that $R(\hat{\boldsymbol{n}}, \pi/4)\, e^+(\hat{\boldsymbol{n}}) = e^\times(\hat{\boldsymbol{n}})$, or, more generally, directly in the $+, \times$ representation,

$$R(\hat{\boldsymbol{n}}, \Psi)\, e^+(\hat{\boldsymbol{n}}) = \cos 2\Psi e^+(\hat{\boldsymbol{n}}) - \sin 2\Psi e^\times(\hat{\boldsymbol{n}}) \tag{4.31}$$

$$R(\hat{\boldsymbol{n}}, \Psi)\, e^\times(\hat{\boldsymbol{n}}) = \sin 2\Psi e^+(\hat{\boldsymbol{n}}) + \cos 2\Psi e^\times(\hat{\boldsymbol{n}})\,. \tag{4.32}$$

To apply a full coordinate change to the measurement equations, this rotation must be applied to the γs by including R^Ψ in the trace in Equation (4.2) as

$$(F_A^+)' = \mathrm{Tr}\left[R^\Psi d_A e^+\right], \qquad (F_A^\times)' = \mathrm{Tr}\left[R^\Psi d_A e^\times\right], \tag{4.33}$$

where d_A is the A detector tensor; note that R^Ψ is orthonormal, d_A is symmetric by definition, and e^P is both. Then

$$\left(F_A^+\right)' = \cos 2\Psi\, \mathrm{Tr}\left[e^+ d_A\right] - \sin 2\Psi\, \mathrm{Tr}\left[e^\times d_A\right] \tag{4.34}$$

$$\left(F_A^\times\right)' = \sin 2\Psi\, \mathrm{Tr}\left[e^+ d_A\right] + \cos 2\Psi\, \mathrm{Tr}\left[e^\times d_A\right]. \tag{4.35}$$

Plugging these into Equation (4.8) the transformed overlap functions are obtained as

$$\left(\gamma^I\right)' = \gamma^I, \tag{4.36}$$

$$\left(\gamma^Q\right)' = \cos 4\Psi\, \gamma^Q - \sin 4\Psi\, \gamma^U, \tag{4.37}$$

$$\left(\gamma^U\right)' = \sin 4\Psi\, \gamma^Q + \cos 4\Psi\, \gamma^U, \tag{4.38}$$

$$\left(\gamma^V\right)' = \gamma^V. \tag{4.39}$$

These properties may also be inferred from the rotational properties of the Stokes' parameters, as the respective γs must obey the same for Equation (4.4) to hold [43]:

$$(Q + iU)' = e^{i4\Psi}\,(Q + iU) \tag{4.40}$$

hence

$$(\gamma^Q + i\gamma^U)' = e^{i4\Psi}\,(\gamma^Q + i\gamma^U)\,. \tag{4.41}$$

With these properties at hand, it is now possible to extend the measurement Equation (4.18) to include the overlap functions for the polarised parameters,

$$D_f^\tau = E_f \Delta_f \frac{4\pi}{N_{\rm pix}} \sum_{\rm P} W_{\rm P} \Gamma_{\rm P}^\tau \, e^{i2\pi f \, \boldsymbol{b}^{\boldsymbol{\tau}} \cdot \hat{\boldsymbol{p}}} + n_f^\tau \,, \tag{4.42}$$

where

$$\boldsymbol{W} = \begin{pmatrix} \boldsymbol{I} \\ \boldsymbol{Q} \\ \boldsymbol{U} \\ \boldsymbol{V} \end{pmatrix}, \qquad \boldsymbol{\Gamma} = \begin{pmatrix} \boldsymbol{\gamma}^I \\ \boldsymbol{\gamma}^Q \\ \boldsymbol{\gamma}^U \\ \boldsymbol{\gamma}^V \end{pmatrix}, \tag{4.43}$$

and the P index spans the $4\times p$ space. This notation is useful as it compactifies the sum over polarisations and the sum over pixels into one. It is now possible to redefine the quadratic response operator as

$$A_{f,{\rm P}}^\tau = E_f \Delta_f \frac{4\pi}{N_{\rm pix}} \Gamma_{\rm P}^\tau \, e^{i2\pi f \, \boldsymbol{b}^{\boldsymbol{\tau}} \cdot \hat{\boldsymbol{p}}} \,, \tag{4.44}$$

and the full solution becomes

$$\tilde{W}_{\rm P} = \sum_{\rm P'} M_{\rm PP'}^{-1} \, z_{\rm P'} \,, \tag{4.45}$$

with

$$z_{\rm P} = \sum_{\boldsymbol{b},\tau} \sum_{f>0} A_{f,{\rm P}}^{\tau\,\dagger} \frac{D_f^\tau}{P_f^\tau} + \text{c. c.} \,, \tag{4.46}$$

and

$$M_{\rm PP'} = \sum_{\boldsymbol{b},\tau} \sum_{f>0} A_{f,{\rm P}}^{\tau\,\dagger} \frac{1}{P_f^\tau} A_{f,{\rm P'}}^\tau + \text{c. c.} \,, \tag{4.47}$$

where the reality condition is imposed by taking the sum over positive frequencies and adding the complex conjugate term. Note that as $V(-f) = -V(f)$ the V part of $M_{\rm PP'}$ is antisymmetric, whereas it is symmetric for all other parameters.

The added complication due to the spin-2 properties of the γ_U and γ_Q functions is that now when estimating the projection in the galactic frame, one needs to take these into account as

$$\Gamma = \begin{pmatrix} \boldsymbol{\gamma}^I \\ \boldsymbol{\gamma}^Q \\ \boldsymbol{\gamma}^U \\ \boldsymbol{\gamma}^V \end{pmatrix}_{\text{Earth}} \longrightarrow \begin{pmatrix} \boldsymbol{\gamma}^I \\ \cos 4\Psi\, \boldsymbol{\gamma}^Q - \sin 4\Psi\, \boldsymbol{\gamma}^U \\ \sin 4\Psi\, \boldsymbol{\gamma}^Q + \cos 4\Psi\, \boldsymbol{\gamma}^U \\ \boldsymbol{\gamma}^V \end{pmatrix}_{\text{Galaxy}} , \tag{4.48}$$

where Ψ is the rotation angle between the frames.

This extension is presented here for the first time and has never been applied in published work. This is partly because a fully stochastic signal is expected to be unpolarised, yielding $Q = U = V = 0$, and partly because a complete decomposition into Stokes' parameters of a stochastic background would only be warranted in the case of positive detection. Nevertheless, it is worth testing the full algorithm here in the hope that it may be of use in the future. Similarly to the tests presented in the previous Section, the injected maps will either be **monopole** only maps to test the ability of the network to constrain a number on the sky, or **Gaussian** maps. In the case of the Q and U parameters the injected maps have been generated with spin-4 weighted spherical harmonics using `HEALPix` tools [119]. These make use of the definitions of the Q and U fields as given in Equation (2.65), injecting a realisation of the spherical harmonic coefficients drawn from a given power spectrum distribution. To avoid conditioning issues, each test only solves for the maps that are expected to be non-zero, i.e. if the injection only includes I and V maps, the reconstruction will only include γ_I and γ_V in the projection and beam such that the size of $\boldsymbol{A}$ and $\boldsymbol{M}$ is scaled to the particular case under consideration.

Figure 4.14 shows the input and reconstructed **monopole** and **Gaussian** maps for the I and V fields using the three-baseline, LH-LL-V network. These results are obtained integrating over 24 hours simulated data, mimicking the aLIGO O1 release data structure. The monopole values of the two fields are reconstructed to percent precision. The injected monopoles are such that $I = 10 \times V$, as realistically the V field will always be lower than the I field. The **Gaussian** injected maps here and in the cases below are generated as random Gaussian fields with $\ell_{\text{max}} = 15$, which is at the boundary of the output pixelisation scale $N_{\text{side}} = 8$. The

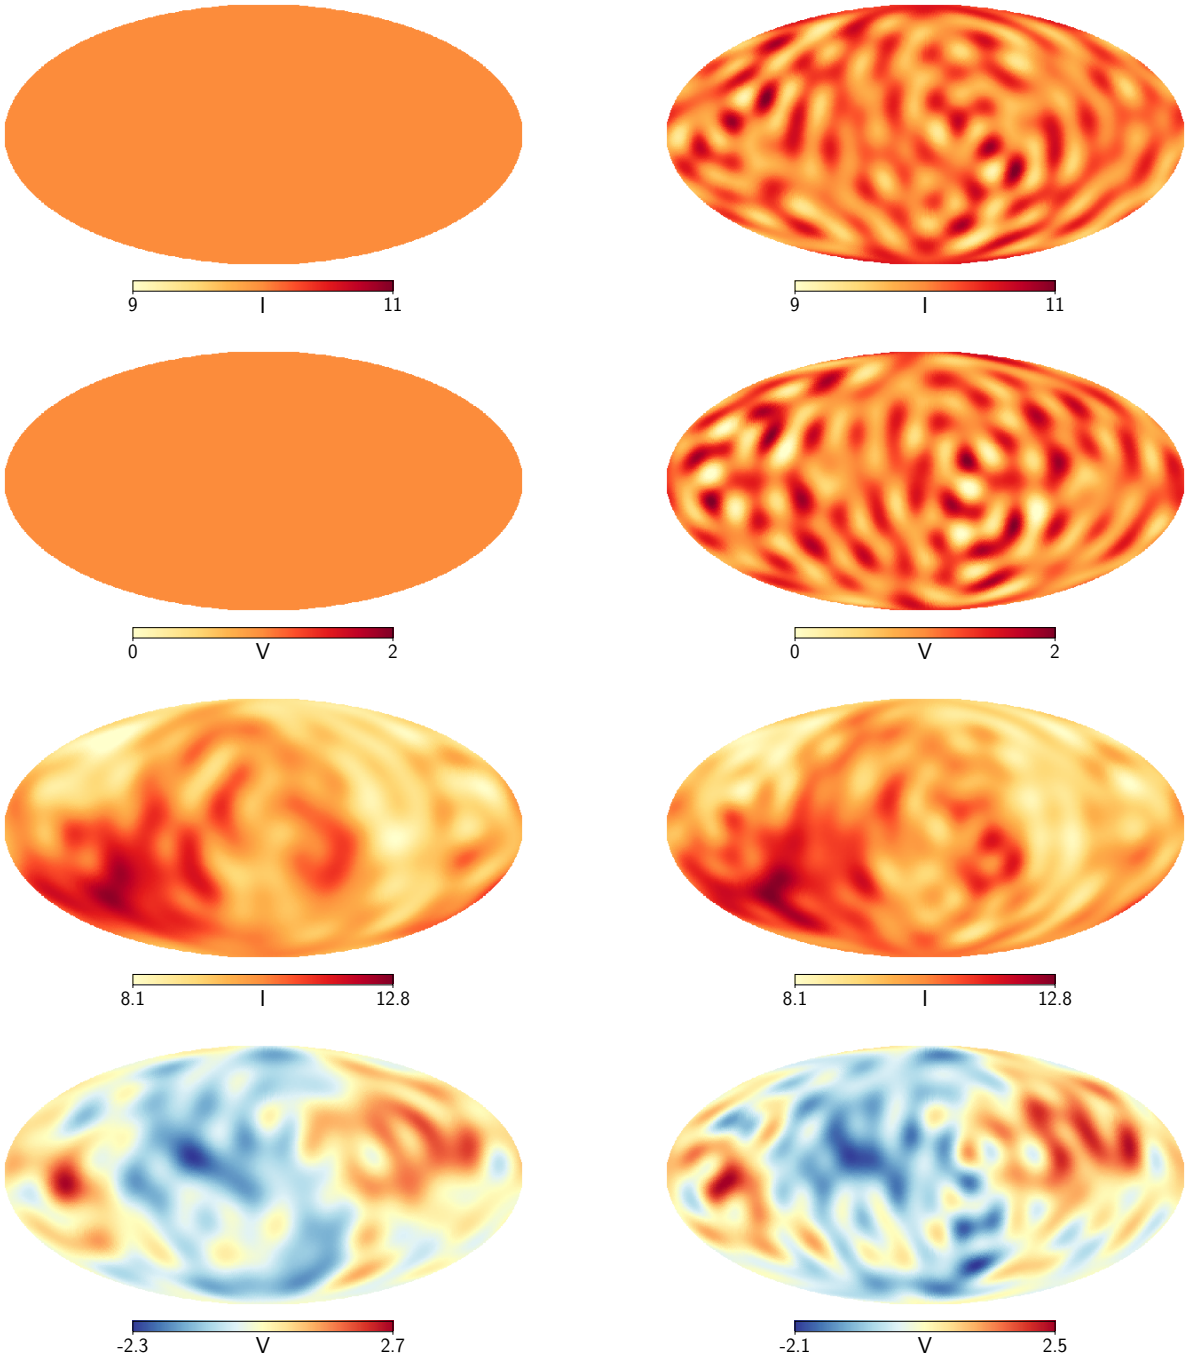


Figure 4.14: The results are for the analysis of 24 hours, simulated data, mimicking the aLIGO O1 release data structure for the three baseline (LH – LL – V) case. Input (left panels) and output (right panels) maps for the I and V Stokes' parameters. The top two cases show the **monopole** injection for both fields, where $I = 10$ is an order of magnitude larger than $V = 1$. The averages of the output maps, $\tilde{I} = 10.03$ and $\tilde{V} = 1.02$ respectively, are in good agreement with the injections. The bottom two cases are the **Gaussian** injections. These are also very well reconstructed. The maps are reconstructed at a lower pixelisation of $N_{\rm side} = 8$ to lighten the computational load, and are then smoothed with a 10 degree Gaussian beam and over–resolved to $N_{\rm side} = 64$ for presentation purposes.

injected signal $\ell_{\max}$ chosen here is higher than that chosen in the previous Section as the computational resources available to us have increased, thus we have taken the opportunity to probe finer angular structure in the background with our mapping technique. Most $\ell \leq 15$ modes are recovered accurately, although some ringing effects remain, due to the limits of reconstructing the higher modes with such a low pixelisation scale.

Figure 4.15 shows the input and reconstructed maps for the I, Q and U fields using the three-baseline network, integrated over 24 hours as above. In this case, the three input maps are generated respectively as a spin-0 temperature map and two spin-4 realisations with the `HealPix` package, in analogy with CMB Stokes' parameters. These are extracted from the same distribution and hence are properly correlated. The I map is amplified to test whether the algorithm is stable with respect to input signals with different orders of magnitude, or whether the solution is subject to mode leakage. In fact, it is found that the reconstruction of the amplified I map is considerably better than the order 1 Q and U maps, and this is reflected in the conditioning of the beam-pattern matrix. The conditioning is set to faithfully reconstruct I, however it is such that Q and U remain poorly constrained. The expectation is that if the I map were amplified by many orders of magnitude with respect to the other parameters, these would be extremely hard to determine. Note that as Q and U fields are spin-4, the monopole, dipole and quadrupole for these will be equal to zero.

A different **Gaussian** case is tested including the V field as may be seen in Figure 4.16. Here, the input maps are generated starting from the h_+ and $h_\times$ fields directly, which are generated as spin-2 fields as described in [96] using specific spin-weighted generating functions from the `HealPix` package. These are then composed as in Equations (2.53-2.55) to obtain I, V, Q, U. Note however that this generation method is not consistent with the observation method, as the h_+ and $h_\times$ fields are generated as non-zero maps whereas in reality the ensemble averages of these are null and it is only in the second order moments of the fields in Equation (2.51) that the non-zero Stokes' parameters emerge; to be strictly correct, the maps in input here would then be instantaneous maps and should not be considered as realistic background representations. In any case, this method is sufficient to induce the correct spin characteristics in the maps and test the algorithm. In this case the fields are injected at

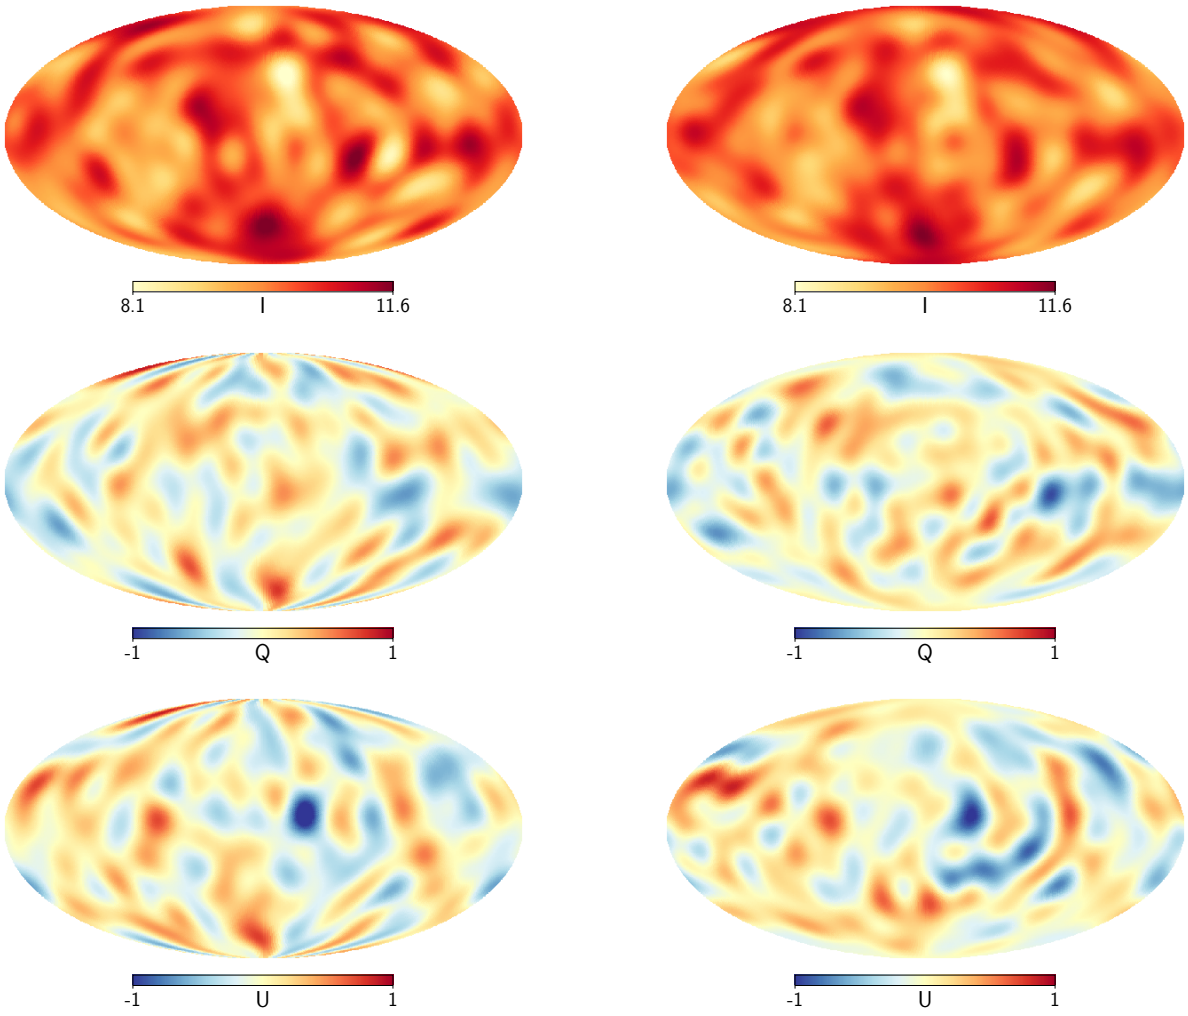


Figure 4.15: Input (left panels) and output (right panels) maps for the I, Q, and U Stokes' parameters. These are all **Gaussian** injections. The I map is the best reconstructed, whereas Q and U are reconstructed less faithfully. This is probably because of the amplitude difference between I and the other parameters. The spin-4 features at the poles in particular are smoothed out of the output maps as these would require a higher resolution than that available through the detector network.

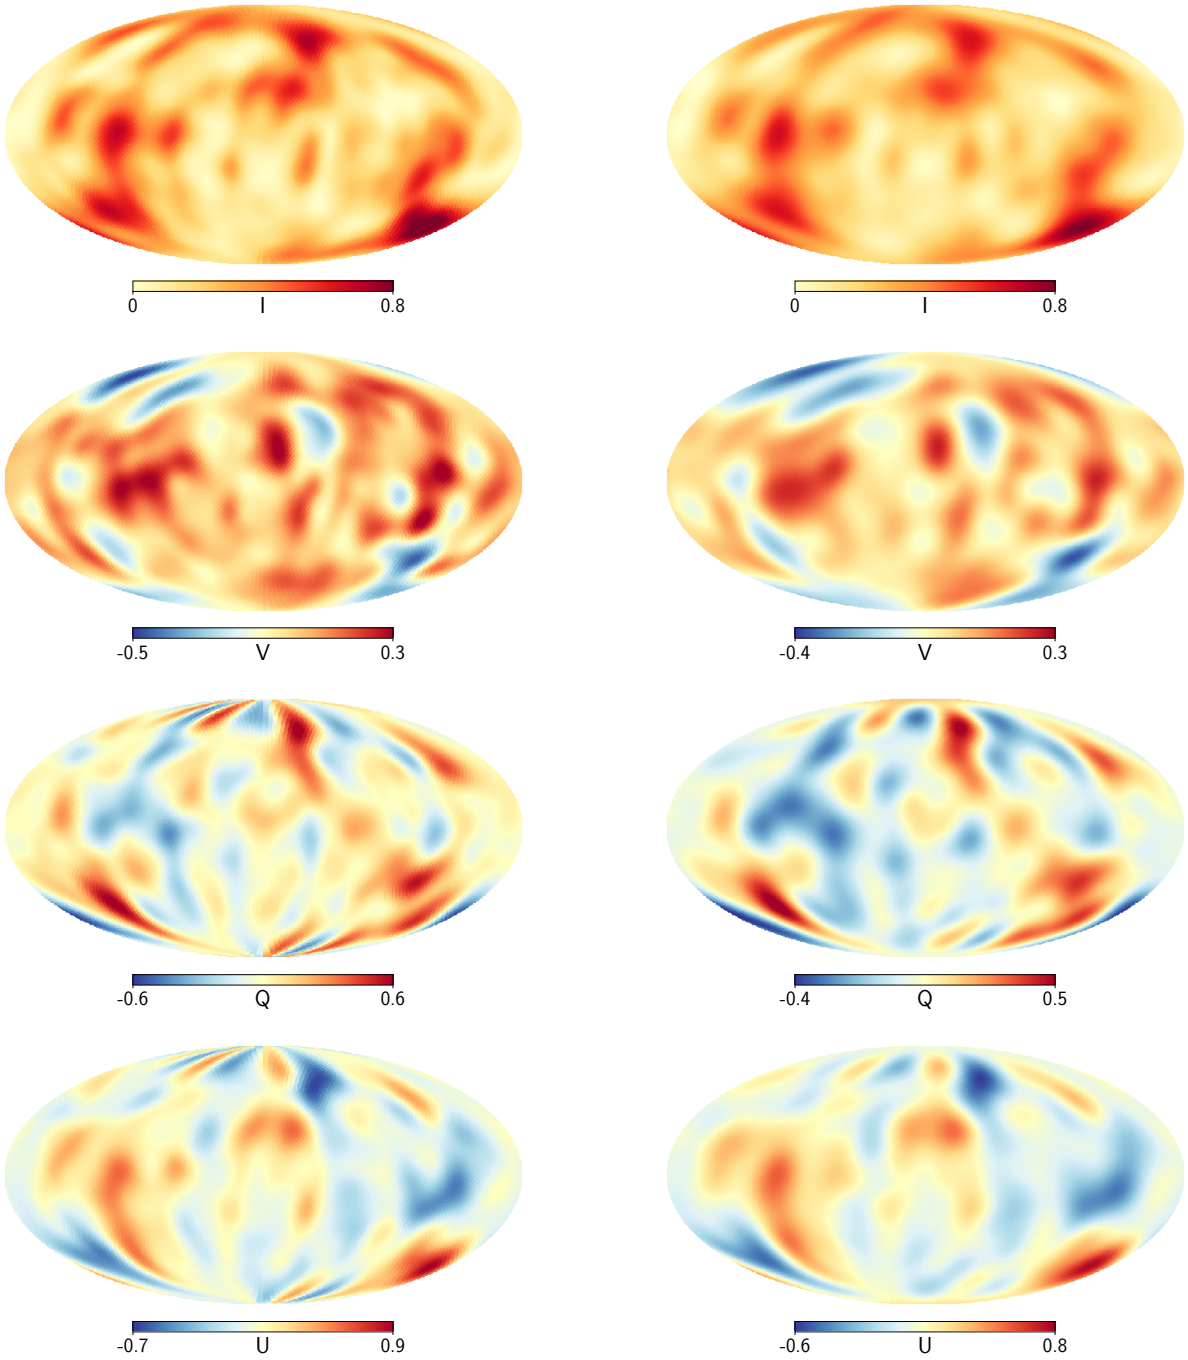


Figure 4.16: Input (left panels) and output (right panels) maps for the I, Q, and U Stokes' parameters. These are all **Gaussian** injections. Note how these are reconstructed equally well, contrary to those shown in Figure 4.15, as the amplitudes of the four fields are comparable. The spin-4 features at the poles of the output maps are still smoothed over but are more visible than in the cases above, as the beam-pattern matrix is better conditioned in this case.

comparable amplitudes, hence the reconstruction is equally good for all four.

The work presented in this section is unpublished to date and represents a new, data-oriented addition to the analysis of anisotropic GW Stokes' parameters. It is clear that given the current state-of-the-art in GW detection, it would be premature to apply the method presented here to real data. Only after a positive detection of the monopole of the GWB, and then the mapping of the intensity I, would it be useful to search for polarised backgrounds in the data.

As previously stated, the majority of stochastic background models don't predict a polarised background. However, in case of a positive detection of a GWB this type of analysis would provide an essential sanity check to test whether the pipeline is correctly set up. As GW detectors are, in fact, polarised and measure the strain directly, any persistent asymmetry in the measured polarisation may indicate an error in detector calibration, modelling, or the interpretation of the timestream.

This analysis presents similarities with the analysis of CMB polarisation, for example presented in [108], in that the Stokes' parameters for the background are extracted from the data through a maximum likelihood solution. However there are some important differences to be pointed out, which also clarify the differences between CMB mapping and GWB mapping discussed in Chapter 3. Firstly, in CMB polarisation mapping the Q and U parameters are combined to form the pure spin-2 gradient and curl modes or E and B *modes*. The equivalent combination for gravitational waves is shown in Equation (2.65) and would return pure spin-4 coefficients for the polarised GWB. While this is not considered here, it is a useful extra step in the analysis as it removes the projection artefacts induced by the spherical coordinate basis and would certainly be part of a polarised SGWB mapping pipeline in the event of a positive detection. Furthermore, the analysis shown here assumes a likelihood for the data of the form shown in Equation (3.4), whereas typically in CMB polarisation data analysis the likelihood is as in Equation (3.19) and it is possible to solve for the power and cross spectra of the temperature, E and B modes directly [108].

4.3 Maps from the Advanced LIGO O1 and O2 data runs

This section illustrates the implementation of the mapping procedure described in Section 4.2 to the first and second Advanced LIGO data sets, O1 and O2. The results and discussion which follow are drawn from [3], where the analysis is limited to O1, and [4] which extends the analysis to O2 and covers new methods as well. These have been developed independently from the LIGO and Virgo collaboration (LVC).

In contrast to similar efforts by the LVC [105, 118, 48] the method used here is based on a pixel domain projection instead of a direct projection into the spherical harmonic domain. This choice is driven by the conditioning of the inversion employed in Equation (4.25), and is backed up by the successful tests illustrated in Section 4.2. At finite resolution, the conditioning is found to be better with this choice due to the complicated mixing of spherical harmonic domain modes implied by the observation of the sky using detectors with non-compact beams [1]. Once obtained, the maximum-likelihood maps can be decomposed into spherical harmonics in order to compare with results in [48] and this comparison is presented in this Section, although it may only be approximate given the very different ways in which the resolution cut-off is applied when projecting directly to different domains. Another distinguishing choice made here is to work in a general sky frame, rather than fixing a specific detector-based frame. Choosing the detector frame may lead to substantial simplifications, see for example [49], however working in sky coordinates allows to set up a general framework for a network of more than two detectors, as described in Section 4.2. The independence and complementarity of this approach with respect to the LVC one will be very useful as the sensitivity levels approach the requirements for detection.

As previously discussed, the choice of spectral dependence $E(f)$ for the intensity on the sky results in strong model dependence of the maximum-likelihood maps. The standard, so far, has been to present results as a function of spectral index α defining a power law assumed shape $E(f) \sim (f/f_0)^{\alpha-3}$ where f_0 is the specific reference frequency at which the estimate is made. This choice is motivated in Chapter 2. In this Section, both a model-dependent, broad-band analysis (4.3.2) and a model independent, narrowband analysis (4.3.3) are presented. In the former, three theoretically motivated shapes are considered: $\alpha = 0$, which corresponds to a flat frequency dependence in Ω_{GW} and thus a template for a background of cosmo-

logical origin [13]; $\alpha = 2/3$, which is the expected dependence for an inspiral dominated background [72]; and $\alpha = 3$ which corresponds to the spectral weight which best describes the the shape of the data noise. For the latter, the signal is broken up into narrowband segments in frequency space and the map-making procedure is applied in each band to obtain a set of independent SNR and noise maps for the intensity of gravitational-wave signal. These are then used to perform an *unbiased* estimate for the spectral shape of $\Omega_{\rm GW}$. This model independent analysis is a new and original approach which was first presented in [4].

The mapping algorithm extracts the maximum-likelihood estimate for the SGWB integrating over filtered timestream data cross-correlated between the LIGO detector pair. Following the recipe in Section 4.2, the noise-weighted projection of the timestream data into map z_p and the beam-pattern matrix $M_{pp'}$ are accumulated over time, where z_p is a projection to inverse noise units. To obtain the correctly normalised, maximum-likelihood map the projection is weighted by the pixel domain noise covariance matrix, which is simply $M^{-1}_{pp'}$. Hence the maximum-likelihood estimate of the sky intensity in each pixel I_p is given by Equation (4.25).

As in the case of the tests presented in the previous Section, the pipeline identifies time-coincident LIGO Hanford and Livingston data blocks which pass a combination of quality flags. The data are divided into 60-second segments and fast Fourier transformed (FFTed) into the frequency domain. As with the mapping tests, a discretisation in frequency indexed by f is assumed in conformity with the standard LIGO down sampled rate of 4096 Hz. The operator $A^{\tau}_{f,p}$ contains all the details of the directional sensitivity of the detector combination on the sky as a function of time and also weights the data with the assumed frequency dependence $E(f)$ for the intensity on the sky. In the convention adopted here pixel indices refer to fixed directions in galactic coordinates. This assumes that the directional sensitivity included in $A^{\tau}_{f,p}$ is rotated from celestial to galactic coordinates accordingly.

As in the previous Section, the maps presented here are recovered at `HEALPix` resolution $N_{\rm side} = 8$ corresponding to $n_{\rm pix} = 768$ pixels on the sky. This corresponds to a pixelisation scale ~ 7 degrees or a Nyquist scale at multipole $\ell \sim 32$. For the purpose of visualising the results the maps are over-smoothed using a Gaussian beam of 10 degrees full width half maximum (FWHM) in order to suppress noise at the resolution scale.

4.3.1 Noise Analysis

In the case of the LIGO data streams, the signal component is expected to be entirely sub-dominant in the frequency range being analysed here and given the duration of the 60 second time segments, the power spectrum of the data timestream itself can be considered an unbiased estimate of the frequency domain noise covariance P_f. A careful characterisation of the noise in each detector is then required to have a handle on the non-stationarity of the noise levels throughout the runs and to optimise the map calculation.

The Advanced LIGO O1 data run spans from 12 September 2015 through to 19 January 2016, while O2 spans from 30 November through to 2016 to August, 2017. The first operation in the pipeline is to select time-coincident sections of the data sets, which reduces the viable O1 timestream to 48.6 days, and the viable O2 timestream to 117.6 days, returning a total of 166.2 days of coincident data. The blocks are then divided into 60 second segments labelled τ, tapered using a narrow cosine window of width 3 seconds to reduce edge effects and fast Fourier transformed (FFT) to the frequency domain. The FFT time length of 60 seconds sets the frequency resolution limit of ~ 0.02 Hz, which is sufficient resolution for this study as the data employed here is the aLIGO data down sampled at the rate of 4096 Hz. Pre-estimated problematic frequencies and known harmonics [124] are carefully notch filtered out of the segments which are then band passed in the frequency range [30.0, 480.0] Hz. The full collection of frequency notches applied to the O1 and O2 data sets may be found in Table 4.2, with the respective widths adopted in this analysis having taken the down sampled frequency resolution into account. In order to cross-correlated consistently, each notch is removed from both Hanford and Livingston timestreams, even though certain lines only appear in one detector. The cut at the high end of the band from 500 Hz to 480 Hz may be considered as a macroscopic notch, as a wandering line in the broad range between 480 Hz and 510 Hz revealed very problematic in O2. This choice mirrors the choice made by the LVC in [48], however the low frequency cut-off is higher than that used in [48], where longer time segments are employed. The cut-off at 30 Hz avoids complications due to very fine frequency combs in the data and is also more robust with respect to non-stationarity in the segments. After the notching routine, a three parameter analytic curve is fitted to the power spectrum density of each segment to obtain a specific weight function, as described in Section 4.2,

Table 4.2: Frequency location f (Hz) and widths σ (Hz) for the inverse Gaussian notching applied to the Hanford and Livingston frequency data. The table covers only the relevant problematic frequencies which may have an impact in frequency range [30, 500] Hz.

O1				O2			
f	σ	f	σ	f	σ	f	σ
34.70	0.5	303.31	1.0	30.0	0.2	240.0	0.1
35.30	0.5	304.99	1.0	31.50	0.2	299.5	1.0
35.90	0.5	306.99	1.0	34.50	0.2	302.0	1.0
36.70	0.5	307.50	1.0	36.20	0.2	303.0	1.0
40.95	0.5	331.90	1.0	40.50	0.01	306.0	1.0
44.69	0.1	453.0	0.05	41.75	0.1	307.0	1.0
47.69	0.1	499.0	2.0	45.50	0.1	315.4	0.2
60.00	0.5	500.0	2.0	46.0	0.1	331.50	0.1
100.0	0.1	510.02	5.0	59.60	0.2	333.33	0.1
120.00	1.0			60.10	0.2	360.0	0.1
179.99	1.0			100.0	0.05	420.0	0.1
299.6	1.0			120.0	0.1	480.0	0.1
302.22	1.0			180.0	0.1	500.25	1.0

and segments which are not consistent with the model are discarded. This procedure cuts a further 15% of the remaining data.

Although the O2 run data constitutes an overall improvement in sensitivity the characterisation of the noise is more complex than in the case of the O1 run, in part due to the increased duration of O2. A number of unidentified noise harmonics are present in the data and these also tend to wander in frequency in specific regions. These lines appear at different frequencies and severity in the two detectors; overall, the Hanford detector is most affected. The noise fitting procedure attempts to minimise the impact of these effects since, although they would not bias the final signal estimate, they would lead to sub-optimal integration of the data.

The noise fitting procedure is as follows. All simultaneous 60 second segment timestreams from each detector are analysed separately. For each detector segment we calculate a windowed power spectrum and perform a three parameter curve fit to $P(f)$, with

$$P(f) = \mathscr{A}^2 \left[\left(\frac{18 f_1}{0.1 + f} \right)^4 + \left(\frac{f}{50 f_2} \right)^2 + (0.07)^2 \right] \times 10^{-44}. \tag{4.49}$$

Here, $\mathscr{A}$ is an overall normalisation factor which depends on the estimation process, and f_1 and f_2 are respectively the low and high knee frequencies of the model. The parameters are re-scaled so as to be approximately centred around 1. This curve is a slightly tweaked version of Equation (4.26) and has proven very stable as it adapts easily to both the O1 and O2 noise spectra, which are considerably different. This may be observed in Figure 4.17, which shows the fractional distribution of the individual detector power spectra over the 60 second segments and the corresponding range in model fits for both data runs. The presence of wandering harmonics in the data is obvious when plotting the spectra in this way as they lead to broader "forests" between 60 and 80 Hz in both data runs. The overall noise level of O1 in the range of [30, 480] Hz is slightly higher than O2, but what makes the real difference are the extremely intense lines in O1, especially around 35 Hz, which give rise to very broad peaks and are necessarily notched out entirely from the data, as they would otherwise dominate the analysis. The O1 noise minus these lines is very stable, whereas the noise in O2 is distributed more finely throughout the frequency range considered here, hence a very careful subtraction of narrow lines is required. The stability of the noise estimation after all these subtractions is reflected in the 95% contours for the scatter in the three parameter fit to the noise in each 60 second segment, which may be seen for both data runs and both LIGO detectors in the four panels of Figure 4.17.

Figure 4.18 shows the one and two dimensional, marginalised distributions for the fitting parameters $\mathscr{A}$, f_1, and f_2 over all viable 60 second segments during the O1 and O2 runs for both detectors. These are not to be confused with parameter posteriors, and are simply the collection of fit parameters implemented in PSD estimation via Equation (4.49) throughout the entirety of the data set. In O1 the parameters are more widely distributed than in O2, and it is possible to identify multiple distinct modes of operation throughout the relatively short run. There is also a significant correlation between the $1/f_2$ knee frequency and the white noise amplitude $\mathscr{A}$, as may be expected from the model in Equation (4.49). Hanford in O1 shows more stability than Livingston, and in particular f_2 has a very peaked distribution. In O2 the distributions are significantly narrower and there is again strong correlation between f_2 and $\mathscr{A}$. It is also possible to see that the Hanford detector operated in two distinct modes with a shift to higher $1/f_2$ knee frequency during the second half of the run. This leads to

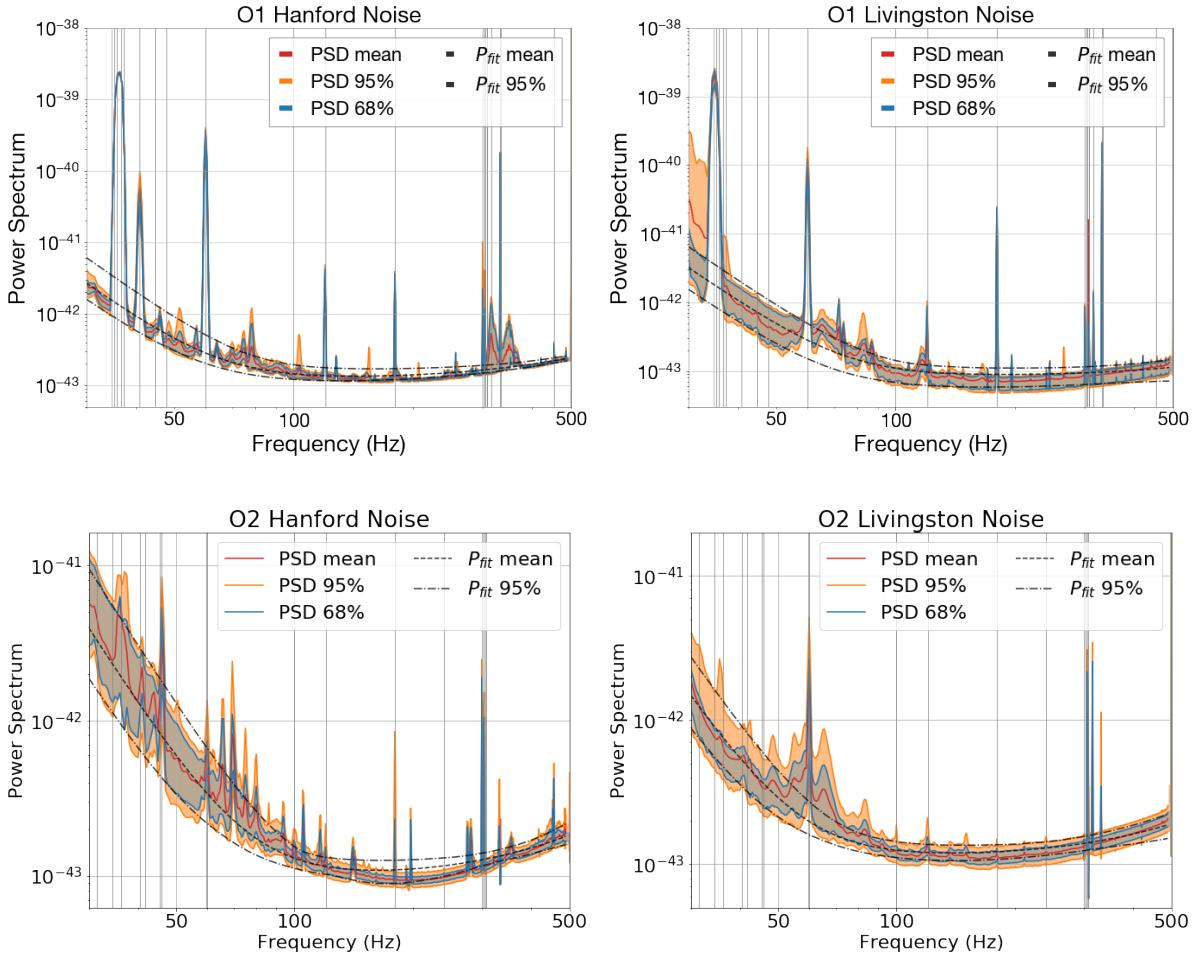


Figure 4.17: Distribution of noise spectra (PSD) for the Advanced LIGO O1 (top panels) and O2 (bottom panels) data runs from the Hanford (left panels) and Livingston (right panels) detectors. The shaded areas represent the 95% and 68% contours of the scatter in 60 second segments. The dashed lines show the mean and 95% contours for the scatter in the three parameter fit to the noise in each 60 second segment. Known systematic harmonics are shown in the vertical lines. Other noise lines are unidentified and shift frequency over time producing wider "forests" in specific regions. The noise fitting procedure masks a number of these regions in order to reduce the bias in the parametrised model.

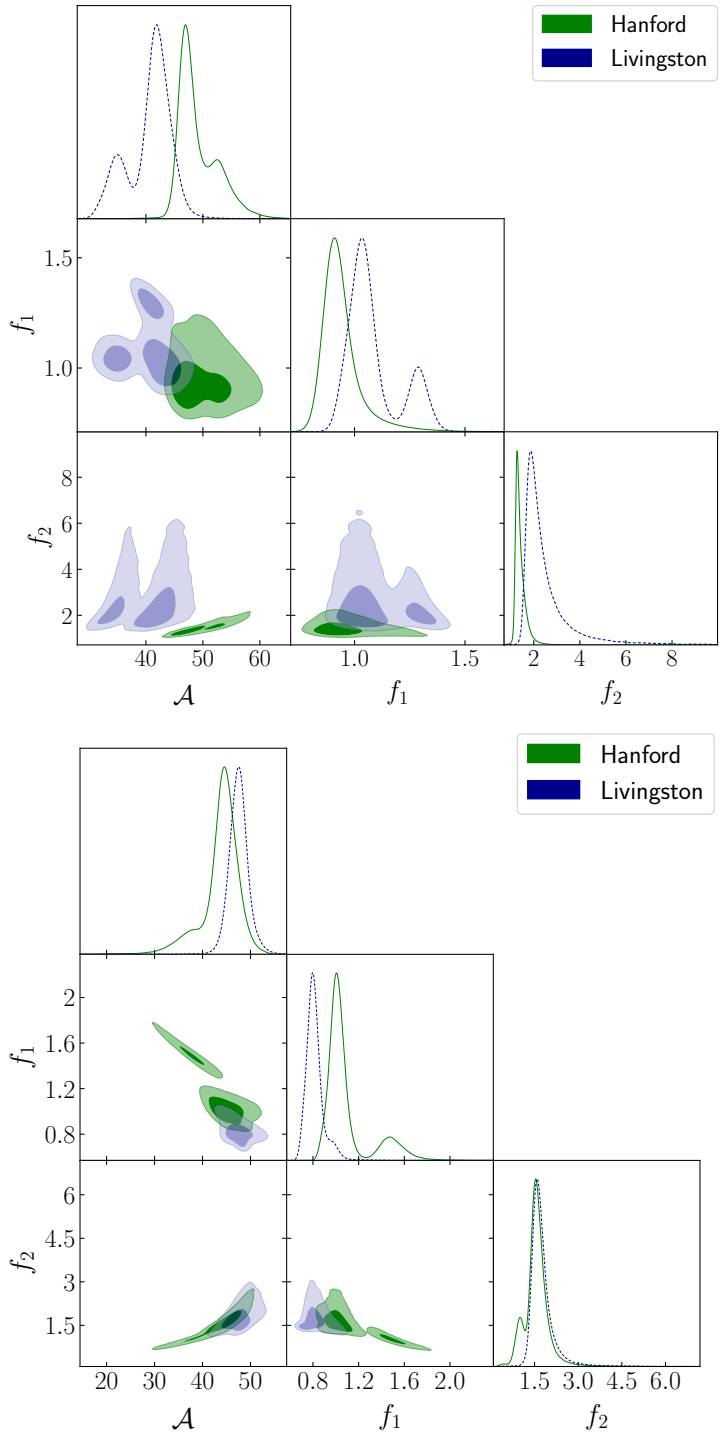


Figure 4.18: One and two dimensional distributions in the noise fitting parameters for the accepted O1 (top panel) and O2 (bottom panel) data segments: white level amplitude $\mathscr{A}$, and re-scaled knee frequencies f_1 and f_2. The distributions are marginalised in the excluded dimension. The triangle plots highlight the multiple sensitivity modes achieved during the runs, characterised by time-correlated splits in the distributions for the three parameters.

an abrupt loss in sensitivity for the lower frequencies, but is limited to the final 20 days of observation. This is correlated to a drop in $\mathscr{A}$, as this parameter is estimated over the whole frequency range and normalisation must be preserved. This splitting was correlated in time with the higher f_1 regime occurring in the last half of the run. On the other hand, Livingston shows extremely good stability compared to O1.

4.3.2 Model-dependent, broadband limits

The results from a broadband integration using three different spectral weightings, $\alpha = 0, 2/3, 3$ are presented below. In this case the entire range of frequencies is integrated over to produce a single map for each choice of α, valid at a reference frequency f_0. As such, the projection map z_p and beam-pattern matrix $M_{pp'}$, in Equations (4.23) and (4.24) respectively, are accumulated over $f = [30, 480]$ Hz and the entirety of the viable observation time. The reference frequencies are chosen within this frequency range and are $f_0 = 50$ for $\alpha = 0, 2/3$, and $f_0 = 100$ for $\alpha = 3$. The former is chosen for both astrophysical and primordial backgrounds as these signals should peak at lower frequencies, hence the chosen reference is at the lower end of the frequency range probed but still above the most problematic frequency lines which may be seen in Figure 4.17. The latter f_0 is chosen for the best-fit case as it is in the most stable frequency segment.

The mapping results of the model-dependent analysis are shown in Figure 4.19. The convention used here is to show SNR and noise maps since the results are only constraints on any signal level. The SNR maps are obtained by normalising the maximum-likelihood maps obtained at the end of the integration by the pixel covariance matrix given by $\mathscr{N}_{pp'} \equiv M^{-1/2}_{pp'}$,

$$S_p = \sum_{p'} \mathscr{N}_{pp'} I_{p'}\,, \tag{4.50}$$

where $\mathscr{N}_{pp'}$ is the inverse of the Hermitian square root of the covariance matrix. This results in a map in units of the expected standard deviation that can be inspected visually for outliers. The off-diagonal pixel to pixel correlations are significant however and therefore an interpretation just based on the SNR maps is only qualitative. The maps are smoothed to a common resolution using a Gaussian filter of 10 degrees FWHM to reduce pixel noise. Note

Table 4.3: Constraints on the isotropic background amplitude for different target spectral indices α. The integration includes frequencies between 30 and 480 Hz.

	α	f_0 [Hz]	$\Omega_{\rm GW}(f_0)$	95% *u.l.*	$\mathscr{P}_{\rm ex}$
O1	0	50	$(-0.3 \pm 1.4) \times 10^{-7}$	2.6×10^{-7}	77%
	2/3	50	$(-7.9 \pm 8.1) \times 10^{-8}$	8.2×10^{-8}	86%
	3	100	$(2.1 \pm 2.0) \times 10^{-7}$	5.9×10^{-7}	41%
O1 + O2	0	50	$(-2.5 \pm 3.9) \times 10^{-8}$	5.5×10^{-8}	19%
	2/3	50	$(-1.6 \pm 3.4) \times 10^{-8}$	5.2×10^{-8}	70%
	3	100	$(1.5 \pm 0.9) \times 10^{-7}$	3.2×10^{-7}	28%

that this smoothing averages over the outliers present in the maps, such that the smoothed maps appear to be within $\sim \pm 2\sigma$ and do not show the 3σ outliers which are present in the original output maps.

The noise maps shown in Figure 4.19 are obtained by plotting the diagonal elements of $\mathscr{N}$. These show how the pixel variance changes across the sky given the full time and frequency integration. The maps also do not carry information about the significant off-diagonal correlations. The variance structure in the maps changes significantly as a function of the spectral index used for the signal weighting in the projectors. To quantify the consistency of the final maps with the expected noise covariances the χ^2 statistics per case are calculated using the full covariance matrices. The probability to exceed statistics $\mathscr{P}_{\rm ex}$ based on these are reported in Table 4.3.

The evolution of the standard deviation in the maximum-likelihood maps as a function of integration time seen in Figure 4.20 is scaling slightly faster than the expected $\sim t^{-1/2}$ behaviour for the O2 data, indicating that there is an improvement in the $1/f$ noise component; this is reflected in the fact that the $\alpha = 0$ and $\alpha = 2/3$ runs are most sensitive to this effect since they are dominated by the lower frequencies.

SGWB Monopole limits

Our estimated intensity maps $I(f_0, \hat{p})$ yield an estimate of $\Omega_{\rm GW}$ at the reference frequency f_0 subject to the broadband integration using an assumed spectral index α. $\Omega_{\rm GW}$ itself is related to the energy density $\rho_{\rm GW}$ of the gravitational waves per logarithmic frequency via Equation (2.9). The monopole, $I(f_0)$, of the maps is calculated using an optimally noise-

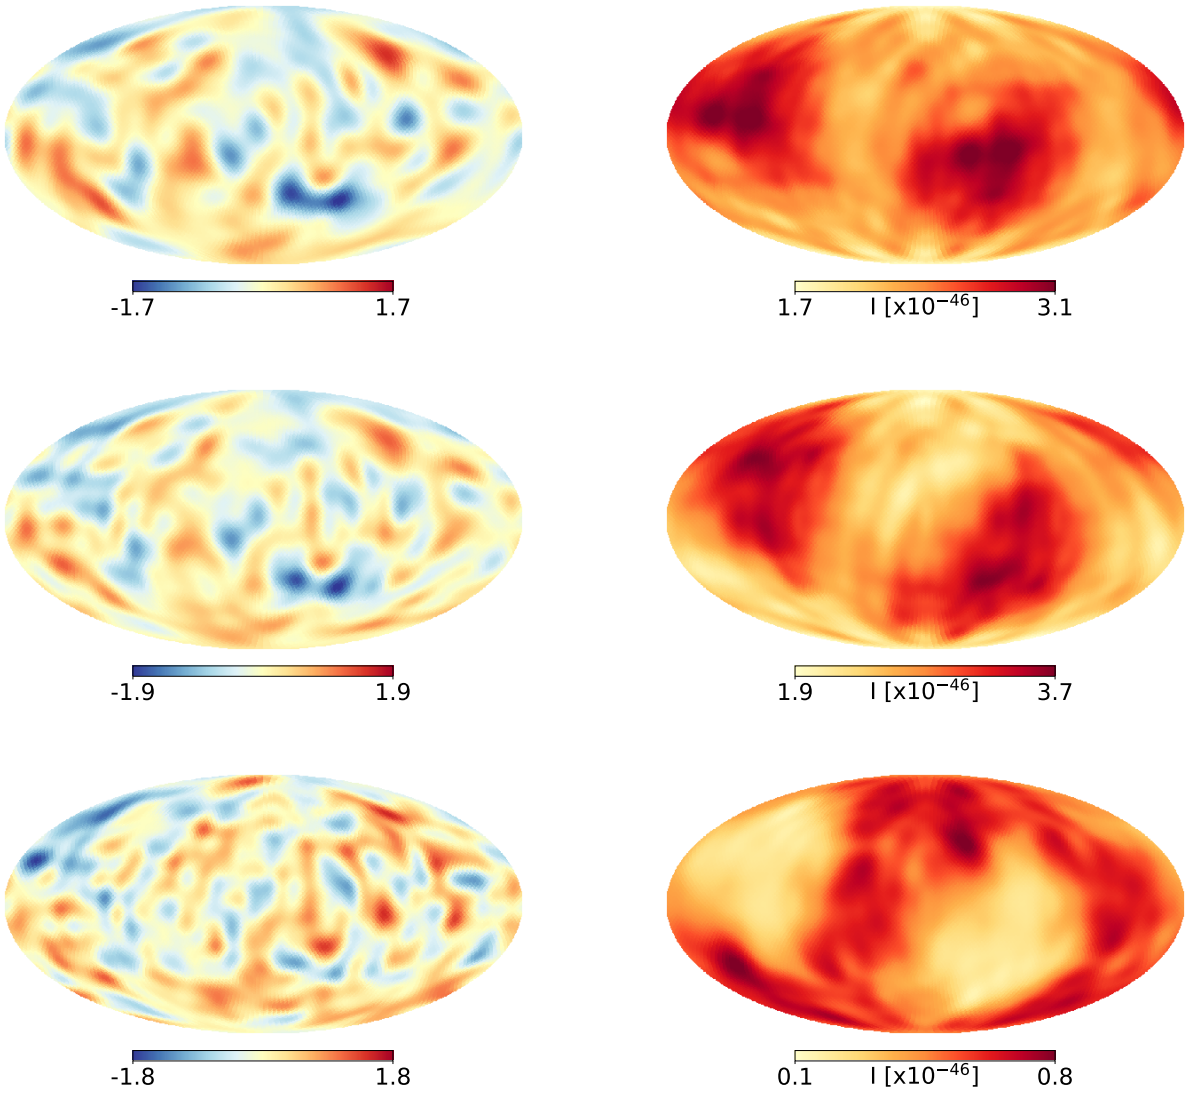


Figure 4.19: SNR (left panels) and noise maps (right panels) of the SGWB with spectral indices, from top to bottom, $\alpha = 0, 2/3, 3$, respectively. All maps have been produced at a `HEALPix` resolution $N_{\rm side} = 8$. For the purpose of visualisation the resulting maps are smoothed with a 10 degree Gaussian FWHM beam. The $\alpha = 0$ and $2/3$ cases are at a reference frequency of $f_0 = 50$ Hz. The $\alpha = 3$ case has $f_0 = 100$ Hz. The noise maps are in units of Hz^{-1}. The SNR maps show that the GWB monopole estimate and its anisotropic features are consistent with zero implying, qualitatively, that the features in the maximum-likelihood solutions are consistent with the noise covariance. The noise distribution maps provide useful insight on the dependence of the sky-localised noise on the frequency weighting. In particular, this dependence seems to excite complementary modes at high and low frequencies respectively.

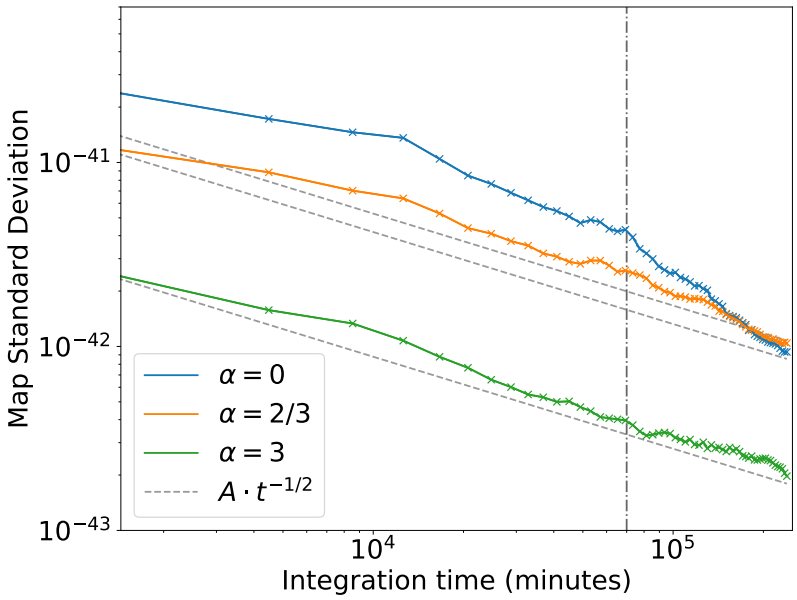


Figure 4.20: Evolution of the standard deviation of the maximum-likelihood maps with integration time for $\alpha = 0, 2/3, 3$ cases. The vertical grey dashed line marks the end of the O1 run and beginning of O2. The constant slopes (short-dashed) are reference $t^{-1/2}$ power laws normalised at the end of the integration. While all three cases show a clear $t^{-1/2}$ trend throughout O1, the $\alpha = 0$ and $2/3$ cases deviate from this constant noise approximation during the O2 run. This indicates that the sensitivity of the O2 run increased as a function of time. Given the spectral weighting of the two cases this also indicates that the improvement was in the low frequency $1/f$ component.

weighted average of the spherical harmonic transform of I_p shown in Equation (4.53) in the next section. Note that these are *not* specific intensity maps and therefore background quantities are obtained by averages over the maps. This can be scaled to yield an estimate of $\Omega_{\rm GW}$ using [93]

$$\Omega_{\rm GW}(f_0) = \frac{4\pi^2}{3H_0^2} f_0^3 I(f_0)\,, \tag{4.51}$$

which is the implementation of Equation (2.66) at the reference frequency f_0. Note that the intensity monopole $I(f_0)$ also includes a normalisation by a factor of $5/8\pi$ due to the normalisation of the detector response functions used in projection operator A_{pf}, as in Equation (4.5). Limits on $\Omega_{\rm GW}$ for the three choices of spectral indices at specific reference frequencies are shown in Table 4.3. None of the choices show any significant excess over the expected noise. The result can be scaled to arbitrary frequencies using the same spectral index assumed in the estimation through the simple scaling relation in Equation (2.68). The limits are many orders of magnitude above the amplitude expected of a primordial background produced during inflation [125] with scale invariant $\Omega_{\rm GW} \sim \mathcal{O}(10^{-14})$. The detection of a

stochastic background of astrophysical origin may not be that far off however. In particular, the $\alpha = 2/3$ spectrum choice is motivated by forecasts of the stochastic signal of binary system inspirals: binary black hole, black hole-neutron star, or neutron star systems. Figure 4.21 shows the upper limit compared to the predictions of [74]. The highest estimates of the stochastic signal are only one order of magnitude below the current upper limit. Hence, with the current detector sensitivities, it would take ~ 100 years of integration time to measure this background, as standard deviation scales as $1/\sqrt{t}$. However the background signal may be within reach of the next planned aLIGO update [14, 126] with integration times similar to O1 and O2, given the projections on sensitivity improvement, as the limits on $\Omega_{\rm GW}$ will scale linearly with individual detector strain sensitivities. The upper limits for $\Omega_{\rm GW}(f_0)$ obtained

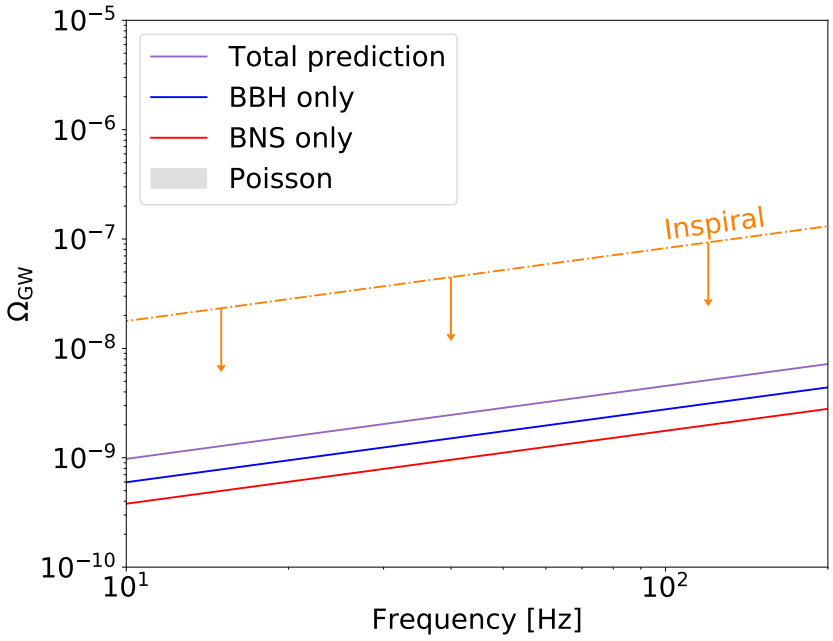


Figure 4.21: The 95% upper limit (dashed) for the broadband integration assuming the spectral index $\alpha = 2/3$ compared to predictions for the stochastic background due to black hole-black hole (BBH) and neutron star-neutron star (BNS) binary system inspirals [74]. The shaded area is an estimate of the uncertainty in the model prediction and the plot is cutoff at 200 Hz as the background is expected to fall off sharply beyond that frequency.

by integrating O1 and O2 data runs presented in Table 4.3 are the most constraining to date. They are in agreement with independent results presented by the LIGO and Virgo collaboration [47], taking into consideration longer integration times used here, as in this analysis fewer time segments are discarded overall. This difference is probably due to the independent quality control pipelines but also due to the shorter segment length adopted for

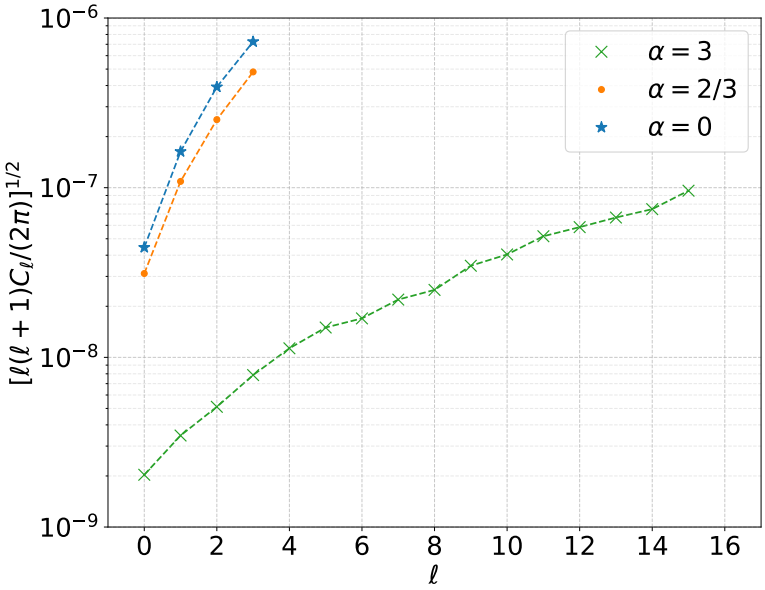


Figure 4.22: 95% upper limits on the signal angular power spectrum C_ℓ's of the SGWB for the the three assumed spectral indices. All values are scaled to a reference frequency $f_0 = 50$ Hz. Depending on the spectral index, the estimate converges up to different maximum multipoles ℓ with the $\alpha = 3$ case giving the highest resolution since it integrates low and high frequency with equal weights. The quantity $\ell(\ell+1)C_\ell/2\pi$ is shown since this is a measure of equal variance per logarithmic interval in ℓ. A scale invariant signal in this measure would be flat.

this analysis. This last point means that, in principle, the stationarity conditions are more easily satisfied in our analysis. Comparing results between runs it is apparent that the $1/f$ component of the noise is significantly better in the O2 run given the improvements in the $\alpha = 0$ case which is the most sensitive to lower frequencies. Despite the fluctuation in noise regimes during the O2 run, the duration and overall sensitivity improvements mean that it dominates the signal-to-noise integration of the combined run.

SGWB Anisotropies

Despite the lack of evidence of any signal it is worth describing how to obtain limits on any anisotropy about the mean background. As anticipated in Chapter 2, anisotropies are most often parametrised in terms of the coefficients $a_{\ell m}$ of the spherical harmonic expansion of the sky map

$$a_{\ell m}(f_0) = \int d\Omega_{\hat{\boldsymbol{n}}}\, I(f_0, \hat{\boldsymbol{n}})\, Y^{\star}_{\ell m}(\hat{\boldsymbol{n}})\,, \tag{4.52}$$

where the $Y_{\ell m}(\hat{\boldsymbol{n}})$ are the spherical harmonic basis functions. In this case, since the maximum-likelihood solution is obtained in the pixel domain, both the sky maps and the noise covariances need to be expanded in spherical harmonics in order to obtain optimal estimates of the $a_{\ell m}$. Since the map and $a_{\ell m}$'s are related by a *linear projection* it is simple to show that the maximum-likelihood estimate of the vector of spherical harmonic coefficients a given a map I_p with covariance M is given by

$$a_{\ell m} = \left(\mathscr{Y}^{\dagger}_{\ell m,p} M_{pp'} \mathscr{Y}_{p',\ell' m'}\right)^{-1} \mathscr{Y}^{\dagger}_{\ell' m',p'} M_{p'p''} I_{p''}\,, \tag{4.53}$$

where the linear operator $\mathscr{Y}_{\ell m,p} \equiv Y_{\ell m}(\hat{\boldsymbol{p}})$ and similar for its adjoint. $\hat{\boldsymbol{p}}$ is the unit vector associated to the pixel p, in the appropriate coordinate basis. Einstein summation over the pixel and spherical harmonic indices is assumed. If the signal is expected to be statistically isotropic (i.e. no preferred direction) then the coefficients can be compressed to an angular power spectrum using the estimator

$$C_\ell(f_0) = \frac{1}{2\ell+1}\sum_m |a_{\ell m}(f_0)|^2 - \tilde{\mathscr{N}}_\ell\,, \tag{4.54}$$

where the isotropised noise bias $\tilde{\mathscr{N}}_\ell$ can be computed from an expansion of the inverse pixel noise covariance M,

$$\tilde{\mathscr{N}}_\ell = \frac{1}{2\ell+1}\sum_m \left|\left(\mathscr{Y}_{\ell m,p} M_{pp'} \mathscr{Y}^{\dagger}_{p',\ell m}\right)^{-1}\right|. \tag{4.55}$$

This estimator is *unbiased* under the assumption of Gaussian, isotropic noise. This is similar to the approach used in CMB power spectrum estimation, for example in the MASTER method [127]. Noise is typically assumed to be isotropic in GWB anisotropy estimations (see e.g. [6]), however note that the isotropisation carried out in Equation (4.55) is effectively averaging over the directional structure seen in the noise maps in Figure 4.19, on top of ignoring the off-diagonal correlations.

Constraints on the C_ℓ's, scaled to $\Omega_{\rm GW}$ for the three separate spectral indices, are shown in Figure 4.22. The choice of maximum multipole ℓ in Equations (4.53) and (4.55) is important; a value that is too low will not exploit the available resolution in the maps, a value that is too high will lead to ill-conditioning which will be manifested as aliasing in the final spectra. Here, the maximum ℓ is increased until the spectra have converged at low multipoles, typically $\ell \sim 16$

and truncated at an ℓ above which the spectra are clearly aliased. The resulting truncation depends on the choice of spectral index with $\alpha = 3$ containing the highest resolution. This is in agreement with the structure shown in Figure 4.19 and is a consequence of the spectral weighting being flat in intensity.

The comparison with the LVC C_ℓ upper limits presented in [48] is possible, noting that the scaling is slightly different and that these are calculated at different reference frequencies. Overall, the results are found to be in agreement, given also the difference in approach. Specifically, the truncation at low ℓ carried out *a priori* in the LVC results flattens out the C_ℓ trend over ℓs, whereas in our results the upper limits rise significantly as a function of increasing ℓ due to the noise dominating the higher modes. Also note the different treatment of the time segments between this analysis and that of the LVC: in the latter case the timestream is parsed into longer, 192 second segments which are Hann-windowed and subsequently analysed with a 50% overlap to minimise the bias due to the wide window. The cross correlation of coincident segments at both detectors is computed in the down-sampled 4096 Hz frequency domain, and then coarse-grained to a frequency resolution of 1/32 Hz. This will necessarily change the effects of notches, and high- and low- pass filtering.

Comparing the C_ℓ constraints here with the models presented in the literature, e.g. in Figure 2.2, is not immediately straightforward as the upper limits in Figure 4.22 are calculated from $\Omega_{\rm GW}$, whereas the theoretical predictions refer to the fractional quantity $\delta\Omega_{\rm GW}/\Omega_{\rm GW}$ as highlighted in Chapter 2. Also, the upper limits here are reference frequency dependent. A useful, preliminary comparison between the anisotropy predictions and the noise power spectrum of the aLIGO baseline is discussed in [6], where it is found that the anisotropies remain a few orders of magnitude out of reach of the current detector configuration. However, planned improvements and the addition of 3G detectors will considerably improve the sensitivity of the ground-based detector network to GWB anisotropies – this is discussed further in Chapter 6.

The SNR scaling in the maps in Figure 4.19 and the time evolution of the standard deviation of the sky strain intensity plotted in Figure 4.20 only show the noise distribution and the diagonal pixel-pixel correlation on the sky. As such it is not meaningful to compare them with those presented in [48] from the LIGO and Virgo collaboration, as these distributions depend

substantially on the different methods used in the integration of the data. A quantitative comparison of the output of the two independent pipelines will become more important as detection becomes within reach. The small scales present in the maps in Figs. 4.19 and 4.24 are then simply the scales at which the noise fluctuates for a given frequency weighting and pixel size, and are not to be mistaken with the angular resolution of the detectors.

Despite the results being presented here are upper limits it is useful to note two important contributions of this work. The first is the consistent evaluation of both the monopole of the background *and* its anisotropies in generalised sky coordinates. This means this method is ideally suited for the future regime where any number of detectors will be networked to provide a number of cross-correlation timestreams, as explored in Section 4.2. This method is applicable to both ground or space-based detectors, or indeed, a mixture of the two as long as frequency bands can be overlapped with sufficient synchronous precision, and providing the conditions for the χ^2 solution in Equation 3.10 are satisfied. Alternatively, this method can be used to explore the efficiency of future network configurations in mapping the background. Forecasting the effective mode coverage of a network of gravitational wave interferometers is complicated by the non-compact beams of each detector. It is a harder task than, for example, in radio interferometry where the primary beam is compact on the sky and the flat-sky approximation can be used along with assumptions of finite resolution in uv-space.

4.3.3 Model-independent spectral limits

An alternative approach to signal estimation is to reduce, to a minimum, the assumptions made about the spectral dependence of the stochastic background. This approach is motivated if there is a possibility that the background is populated by distinct sources with different spectral shapes – this would lead to breaks in the power-law for the background or even more complicated features such as well-defined peaks [61, 52] – see Figure 2.1 in Chapter 2. In fact, a model-independent spectral estimation as proposed here could be part of a method aimed at separating the contributions of different sources based on their spectral dependence. Combined with the angular separation provided by the mapping itself, this

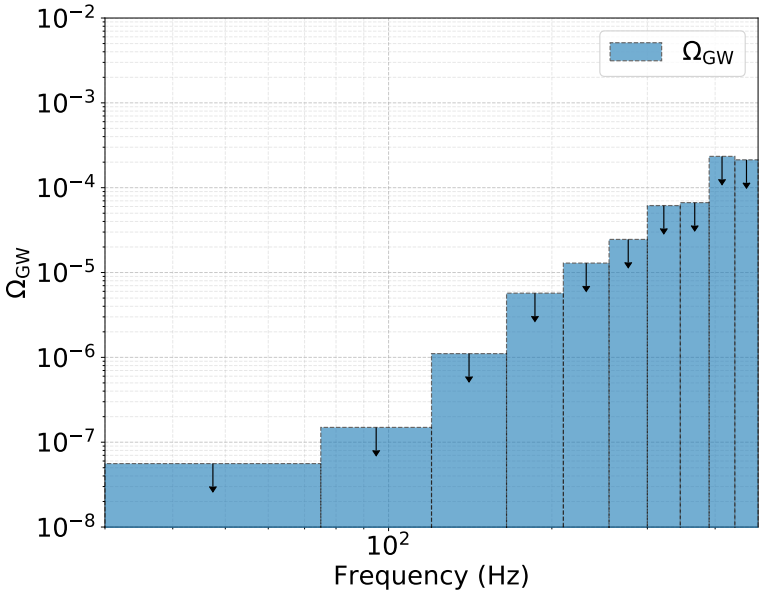


Figure 4.23: 95% upper limits for the SGWB monopole for each of the maps obtained in the ten separate frequency bins used in the model-independent spectral analysis. The method assumes a scale invariant shape for the signal in each spectral bin. The spectral shape is found to be consistent with a noise dominated estimate.

constitutes a prototype separation method for future observations where more than one family of sources will be visible in the target frequency range.

To test this method the algorithm described above is applied on ten separate frequency bins, accumulating the projection operators only over frequencies assigned to each bin. The spectral dependence assumed within each bin is scale-invariant with respect to the dimensionless background density, i.e. $E(f) \sim 1/f^3$ with $\alpha = 0$. This is a form of "least informative" prior on the spectral dependence but other choices can be made in order for the compression to be optimal with respect to a specific spectral shape [121]. For example, it is also possible to use a "broken power law" approach assuming different spectral shapes in different bins, enabling the reconstruction of a spectral dependence of the background beyond a simple power law and which may be due to multiple backgrounds of different origin contributing to the timestream.

The bins are defined as ten equally spaced intervals of width 45 Hz between 30 and 480 Hz. Ten maps are produced in this exercise, one for each frequency bin, with associated noise covariance matrices. As described above in the broad-band case, it is possible to estimate

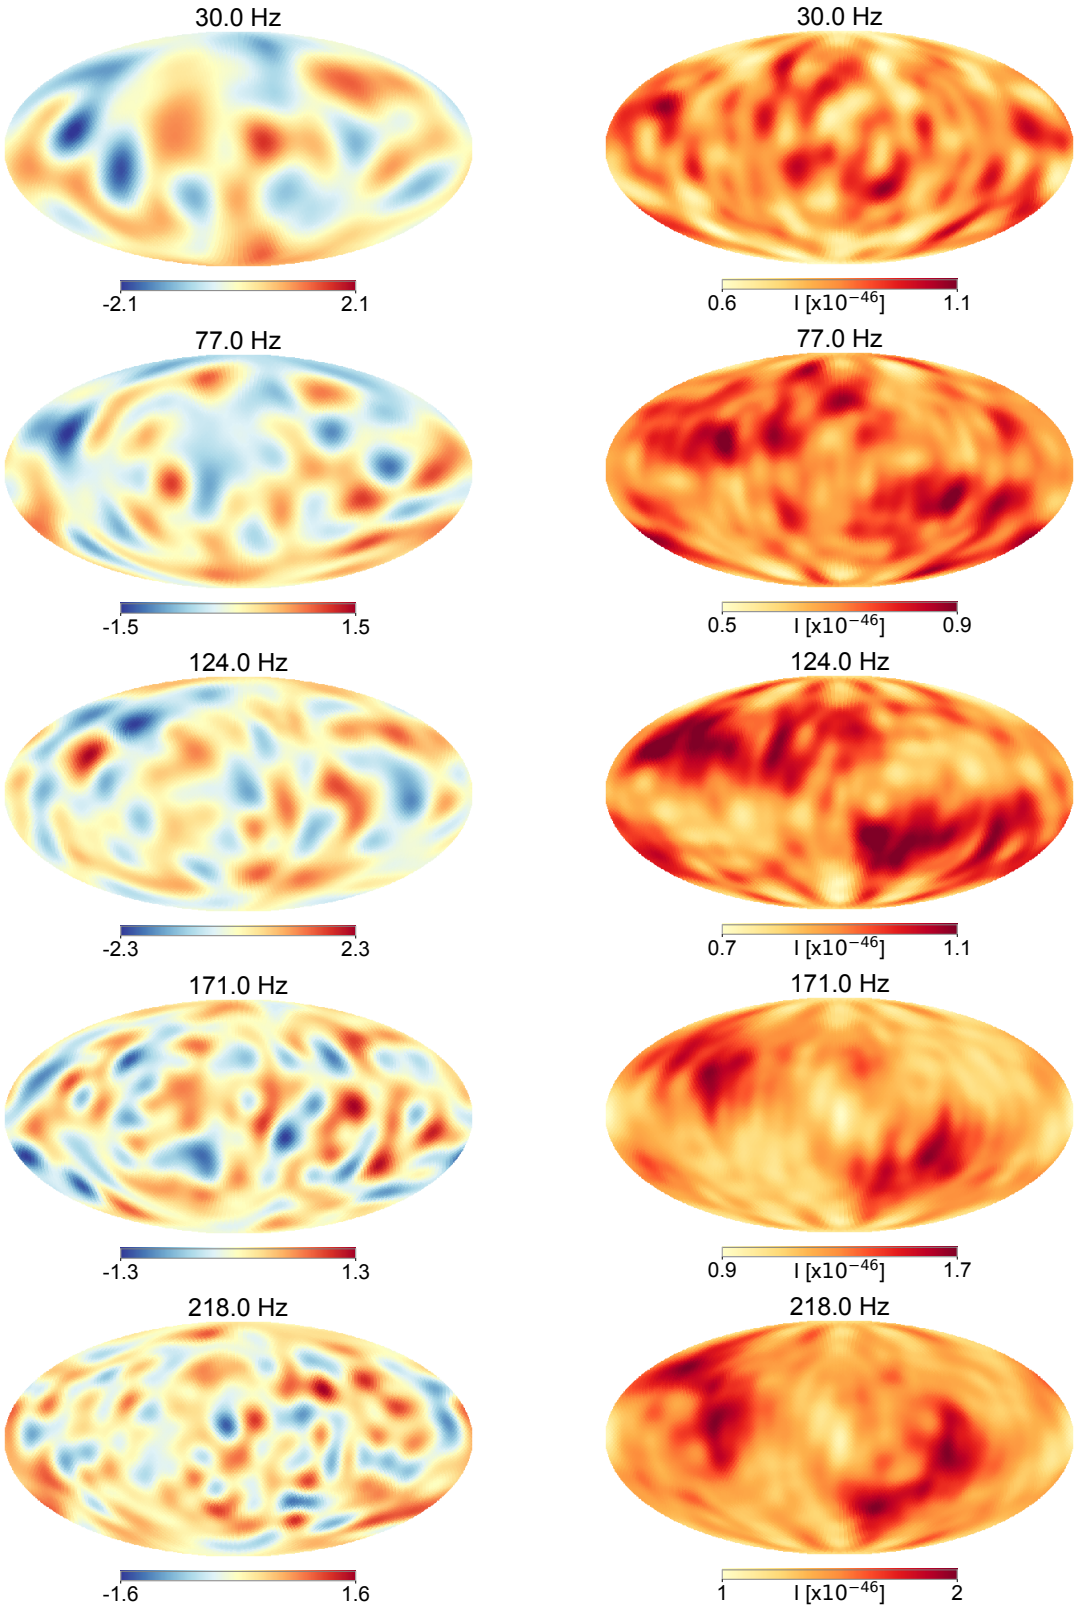


Figure 4.24: SNR (left panels) and noise maps (right panels) for the frequency bins used in the model-independent limits. Each set is labelled by the central frequency in the corresponding bin.

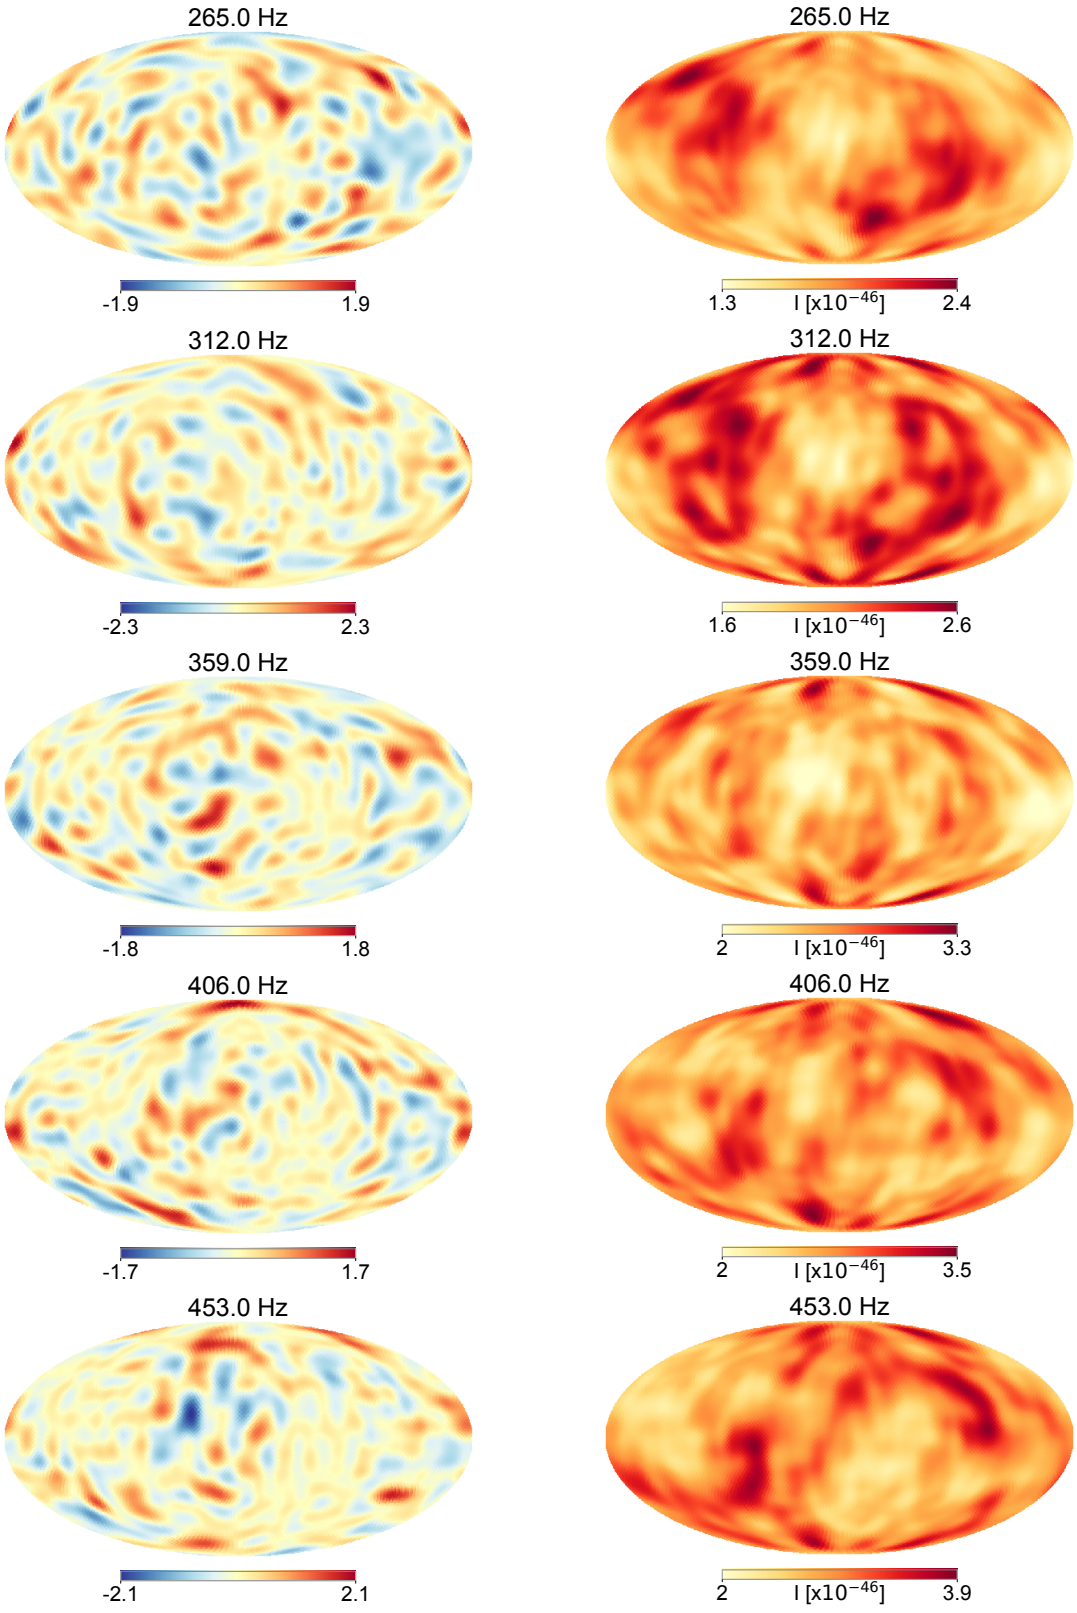


Figure 4.25: SNR (left panels) and noise maps (right panels) for the frequency bins used in the model-independent limits. Each set is labelled by the central frequency in the corresponding bin.

both a Ω_{GW} and its anisotropies from each map. The results for the 95% upper limit in the monopole of Ω_{GW} are shown in Figure 4.23. The constraints are higher than the signal in the equivalent single, broadband estimate, since the information is split into ten separate estimates but it is useful to note that the spectral dependence is consistent with a noise dominated estimate with increasing power as a function of frequency.

Figures 4.24 and 4.25 show the SNR and noise maps for the ten 45 Hz–wide frequency bins. Given the noise level in this analysis we do not extract upper limits on any measures of the anisotropy but it is useful to note how the mode structure on the sky changes as a function of frequency. This effect, evident in the left-hand columns of Figures 4.24 and 4.25, becomes very apparent in this kind of spectral analysis and can be understood through the frequency-to-sky mode coupling encoded in the projection operator $A^{\tau}_{f,p}$. The right-hand columns reveal the sky modulation of the noise as a function of frequency, which appears smoothed over the whole weighted frequency range in the noise maps in Figure 4.19. This modulation drives the apparent structure in the SNR maps, which is thus purely noise dominated. The hot spots in the noise maps are most likely particularly noisy frequencies in the narrow frequency window considered, which couple to a specific sky direction.

This is the first all-sky, frequency-binned analysis of the aLIGO O1 and O2 data sets. Although these continue to be upper limits of a signal which is expected to be out of the reach of present detectors, this method will prove very useful once GW science enters the high signal-to-noise regime when multiple backgrounds may be detected simultaneously. This method is a basis for the *combined* angular/frequency based separation of multiple backgrounds.

Chapter 5

Mapping the Gravitational-Wave Background with the Laser Interferometer Space Antenna

The Laser Interferometer Space Antenna (LISA) is an ESA led mission in collaboration with NASA planned to launch in the mid 2030s which will allow us to tune into a new range of the GW sky. The undisputed success of the LISA Pathfinder mission, which served as proof-of-concept for LISA's feasibility, removed the final roadblock to LISA, securing funding for research and development [128]. At time of launch, LISA is expected to be the most sensitive GW detector ever built to Ω_{GW}. The LISA sensitivity curve spans roughly 4 orders of magnitude in frequency, from 10^{-5} Hz to 10^{-1} Hz [21], with maximum sensitivity around 10^{-3} Hz, as may be seen in Figure 1.1 in the Introduction. It will probe the Universe with unprecedented depth, receiving even the faintest GW signals from the galaxy, extra-galactic stellar-mass binary inspirals and mergers, all the way to loud massive compact binary events out to redshift z = 6 and beyond [129]. Given the variety of signals contributing to this band and the expected noise curve [130], the timestream will be densely populated by waveforms with respective SNRs ranging from a few hundred to a few. This scenario is radically different from that in LIGO and other ground-based interferometers, such that data analysis will require careful component separation to disentangle the individual coherent sources, minimising bias [131]. There will also be the need for tools to analyse the stochastic signals which build

up incoherently in the timestream, and which may well have high SNRs [132].
There are several GW sources which will contribute to a stochastic background in the LISA band. These include galactic, extra-galactic or even primordial compact binaries, too distant or faint to be resolved [12], or a relic background from a burst of inflation or a phase transition at early times [13, 59–63]. All backgrounds mentioned here are the focus of science objectives of the LISA mission [21].

This Section presents the first maximum-likelihood method to extract the stochastic signal in the LISA data, and how to solve for its directionality, effectively mapping it on the sky. This will allow to assess the angular resolution of the LISA detector for a distribution of extended, unresolved sources, as a function of SNR. The aim is to deliver the framework of the LISA SGWB map-making pipeline which will be progressively refined and updated until data become available. The mapper described here is based on the optimal quadratic estimator in the case of correlated data described in Chapter 3 which solves for maps in the pixel domain. The approach is influenced by work done previously with LIGO data, both as a part of the LIGO Virgo Collaboration [105, 118, 48] and as independent efforts [3, 4] which are presented in Chapter 4. Much work has been done to assess the ability of LISA to detect and characterise GWBs [35, 50–53], and also study the sky response of LISA for map-making purposes [36, 43]. Some studies have focused on how to distinguish detector noise from the galactic stochastic background, often referred to as a foreground when compared to other signals, which is expected to imprint a seasonal modulation on the time stream [133, 134]. These are particularly relevant here as the ability of LISA to map the galactic foreground will be probed in Section 5.2.2 by analysing an idealised signal with a Milky Way–like sky distribution. Map-making with space-based detectors is discussed extensively in [54, 37], both in the low frequency limit and at high frequencies where a higher resolution may be achieved. However, these approaches reduce the mapping problem to a least-squares solution and solve for the spherical harmonic coefficients of the signal by looking at the individual cross-correlations between data-streams.

This Chapter is drawn from [2] by this author and presents the first all-sky analysis based on the full likelihood of the data given a GWB intensity signal. First, the detector response

is introduced and analysed both in the pixel and spherical harmonic domain. Then, the optimal map solution is tailored to the LISA set-up. To test the map-making procedure different input signals with a variety of input maps are injected in the detector in order to verify the reconstruction of a known sky, with simple spectral dependence. The inputs consist in stationary maps of strain intensity of varying amplitudes and with anisotropies given by a Gaussian random field. These are analogous in spirit to the tests carried out in Chapter 4 and are not to be regarded as realistic simulations of any particular GWB.

5.1 Detector response

The space antenna is comprised of three spacecrafts which will be positioned on three heliocentric yearly orbits in such a way that they will maintain an equilateral triangle configuration throughout the whole duration of the mission (4+ years) [21]. The spacecrafts will specify a detector plane at a 60° angle with the ecliptic consistently throughout the motion of the constellation, which will appear to rotate on the detector plane with period 1 year. Each pair of spacecrafts will share two laser *links*, which may be represented as two opposite vectors along the virtual detector arm. Effectively, the raw LISA data will be the response of each of these links, however it is necessary to combine the link data to measure gravitational waves, and not all link combinations are equally sensitive to the GW signal, as explained below.

The arm-length will in reality be time dependent, however in this work equal, stationary arms are assumed so the links are equal and opposite along each arm. The trajectories for the three spacecrafts used here are defined in the internal LISA document [135], and correspond to analytical elliptical orbits in the Solar System Barycentre (SSB) frame. These are

$$x_i = \cos\theta + e\big(\sin\theta\cos\theta\sin\beta_i - \big(1+\sin^2\theta\big)\cos\beta_i\big)\,, \tag{5.1}$$

$$y_i = \cos\theta + e\big(\sin\theta\cos\theta\cos\beta_i - \big(1+\cos^2\theta\big)\sin\beta_i\big)\,, \tag{5.2}$$

$$z_i = -\sqrt{3}e\cos\big(\alpha-\beta_i\big)\,, \tag{5.3}$$

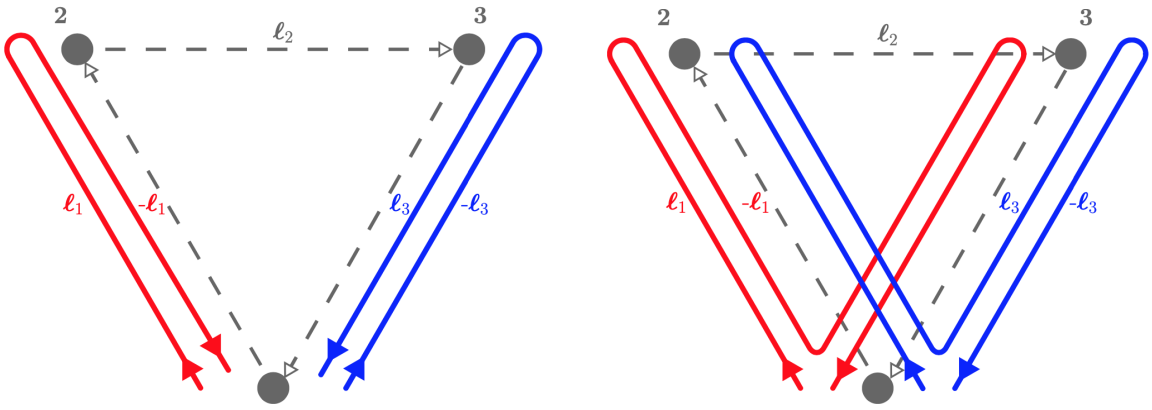


Figure 5.1: Illustrations of the TDI 1 (left panel) and TDI 1.5 configuration X (right panel). The three numbered grey circles represent the spacecrafts arranged in an equilateral triangle. The links are labelled $\boldsymbol{\ell}_i$ in the left panel for clarity, and are the same as the ones in the right panel.

and each spacecraft position vector labelled by $i = 1, 2, 3$ is $\boldsymbol{p}_i = 1\text{AU} \cdot (x, y, z)_i$. The phases

$$\beta_i = (i-1)\frac{2\pi}{3} + \lambda\,, \tag{5.4}$$

$$\theta = \frac{2\pi}{1\text{year}}t + \kappa\,, \tag{5.5}$$

are such that the starting position at zero time is defined by $\lambda = 0$, $\kappa = 0$ and is chosen such that the second arm labelled as $\boldsymbol{\ell}_2$ in Figure 5.1 is aligned with the negative y axis of the frame, and the two remaining arms are determined accordingly. e is the eccentricity which is the same for each orbit as it is computed based on the arm length, $e = L/(1\text{AU} \cdot 2\sqrt{3})$.

GW detection with LISA will rely on time-delay interferometry (TDI), which involves time-shifting and linearly combining the independent Doppler measurements made with the the laser links to obtain combinations sensitive to gravitational waves [136]. This will occur in post-processing and will require accurate modelling of the response of each link. The study of the optimal TDI combinations, or channels, for GW detection is ongoing [137]. Given 6 links it is possible to construct three correlated channels, referred to as X, Y, Z, or two independent channels, typically referred to as A, E channels, plus a *Sagnac* channel T where the signal component cancels out at low frequencies [138]. The independent channel basis is found simply by diagonalising the correlated basis and rotating into the *diagonal* TDI space. As the measurements are time-delayed, this rotation will never be perfect and adds a degree of complication which may be hard to constrain, hence it is preferable to work in the fully correlated basis, with the caveats it entails. Specifically, the rotation will only cancel out the

correlations assuming the three arm-lengths are equal, and that the noise in each spacecraft is identical. While these assumptions are made in this work, the method developed here is general and may be easily extended to accommodate a more complete noise model and detector response.

The simple Michelson-Morley-style combination is referred to as TDI 1 and features three channels, each of which is centred around a spacecraft such that X is centred around spacecraft 1, Y around spacecraft 2, and Z around 3. A similar, slightly more sophisticated combination designed to minimise the spacecraft breathing noise is referred to as TDI 1.5; both are illustrated in Figure 5.1. The TDI 1.5 configuration is employed in the work presented below.

The timestream $s_C(t, \boldsymbol{p}_i)$ measured by a single 1.5 TDI channel C at time t in the ith spacecraft at position $\boldsymbol{p}_i$ can be Fourier expanded between t and $t + \Delta t$ to yield the signal as a function of frequency,

$$\tilde{s}_C^{\tau}(f, \boldsymbol{p}_i) = \sum_P \int_{S^2} d\hat{\boldsymbol{n}}\, R_C^P(f, \hat{\boldsymbol{n}}; \tau, \boldsymbol{p}_i)\, h^P(f, \hat{\boldsymbol{n}})\,, \tag{5.6}$$

where τ is the time segment label, R_C^P is the polarisation response function for the TDI channel C and $h^P(f, \hat{\boldsymbol{n}})$ is the incoming GW strain decomposed into its polarisation components. The τ label is now dropped for simplicity. The response functions for TDIs 1.5, $\{X, Y, Z\}$ are derived from the internal document [135]. In the case of TDI 1.5 X the measured strain may be written as

$$\tilde{s}_X(f) = -2i \sin(2\pi f L)\, e^{-i2\pi f L} \left[e^{-i2\pi f L} \left(y_{\boldsymbol{\ell}_1} - y_{-\boldsymbol{\ell}_3} \right) + y_{-\boldsymbol{\ell}_1} - y_{\boldsymbol{\ell}_3} \right], \tag{5.7}$$

where the $y_{\boldsymbol{\ell}_i}$ terms correspond to the strain measured along the oriented link $\boldsymbol{\ell}_i$ and the exponential terms correspond to the phase shifts in Fourier space required to build the TDI 1.5 channel. The strain along each link is

$$y_{\boldsymbol{\ell}_i} = -i\pi f L \sum_P \int_{S^2} d\hat{\boldsymbol{n}}\, Q^P_{\hat{\boldsymbol{\ell}}} h^P\, e^{-i2\pi f \hat{\boldsymbol{n}}\cdot\boldsymbol{p}_i} \operatorname{sinc}\left(\pi f L(1 - \hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}}_i)\right) e^{-i\pi f L(1-\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}}_i)}\,, \tag{5.8}$$

where the response of a single arm $\boldsymbol{\ell}$ to P- polarised modes, $Q^P_{\hat{\boldsymbol{\ell}}}$, is simply the contraction of the arm tensor with the polarisation basis element,

$$Q^P_{\hat{\boldsymbol{\ell}}} = \epsilon^P : \hat{\boldsymbol{\ell}} \otimes \hat{\boldsymbol{\ell}} \,. \tag{5.9}$$

Plugging Equation (5.8) into Equation (5.7) the sky response R^P_X is recovered,

$$R^P_X(\hat{\boldsymbol{n}}) = i\pi f L \left(1 - e^{-4\pi i f L}\right) e^{-2\pi i f \hat{\boldsymbol{n}}\cdot\boldsymbol{p}_1} \left[Q^P_{\hat{\boldsymbol{\ell}}_1} \mathcal{T}(+\hat{\boldsymbol{\ell}}_1) - Q^P_{\hat{\boldsymbol{\ell}}_3} \mathcal{T}(-\hat{\boldsymbol{\ell}}_3) \right], \tag{5.10}$$

where the transfer function $\mathcal{T}$ to each sky direction $\hat{\boldsymbol{n}}$ is

$$\mathcal{T}(+\hat{\boldsymbol{\ell}}) = \mathrm{sinc}\left(\pi f L(1 - \hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}})\right) e^{-\pi i f L(3+\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}})} + \mathrm{sinc}\left(\pi f L(1 + \hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}})\right) e^{-\pi i f L(1+\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\ell}})}, \tag{5.11}$$

and the equal arm assumption implies that $\boldsymbol{p}_i = \boldsymbol{p}_{i-1} + \boldsymbol{\ell}_{i-1}$ where i obeys cyclic permutations. R^P_Y, R^P_Z are then derived by permuting the arms,

$$R^P_Y(\hat{\boldsymbol{n}}) = i\pi f L \left(1 - e^{-4\pi i f L}\right) e^{-2\pi i f \hat{\boldsymbol{n}}\cdot\boldsymbol{p}_2} \left[Q^P_{\hat{\boldsymbol{\ell}}_2} \mathcal{T}(+\hat{\boldsymbol{\ell}}_2) - Q^P_{\hat{\boldsymbol{\ell}}_1} \mathcal{T}(-\hat{\boldsymbol{\ell}}_1) \right], \tag{5.12}$$

$$R^P_Z(\hat{\boldsymbol{n}}) = i\pi f L \left(1 - e^{-4\pi i f L}\right) e^{-2\pi i f \hat{\boldsymbol{n}}\cdot\boldsymbol{p}_3} \left[Q^P_{\hat{\boldsymbol{\ell}}_3} \mathcal{T}(+\hat{\boldsymbol{\ell}}_3) - Q^P_{\hat{\boldsymbol{\ell}}_2} \mathcal{T}(-\hat{\boldsymbol{\ell}}_2) \right]. \tag{5.13}$$

A signal composed of an incoherent superposition of many components, as in the case under examination here, will vanish when averaged over time. To observe an incoherent GWB it is therefore necessary to consider the integration of the square of the signal. Making use of all three channels the TDI strain vector may be written as $\boldsymbol{s} = (s_X, s_Y, s_Z)$ and it is possible to construct the quadratic strain tensor $\boldsymbol{S}(f)$,

$$\boldsymbol{S}(f) = s(f) \otimes s^\star(f) \,. \tag{5.14}$$

The ensemble average is expected to yield

$$\langle \boldsymbol{S}(f) \rangle = \int_{S^2} d\hat{\boldsymbol{n}}\, \boldsymbol{A}(f, \hat{\boldsymbol{n}})\, I(f, \hat{\boldsymbol{n}}) \,, \tag{5.15}$$

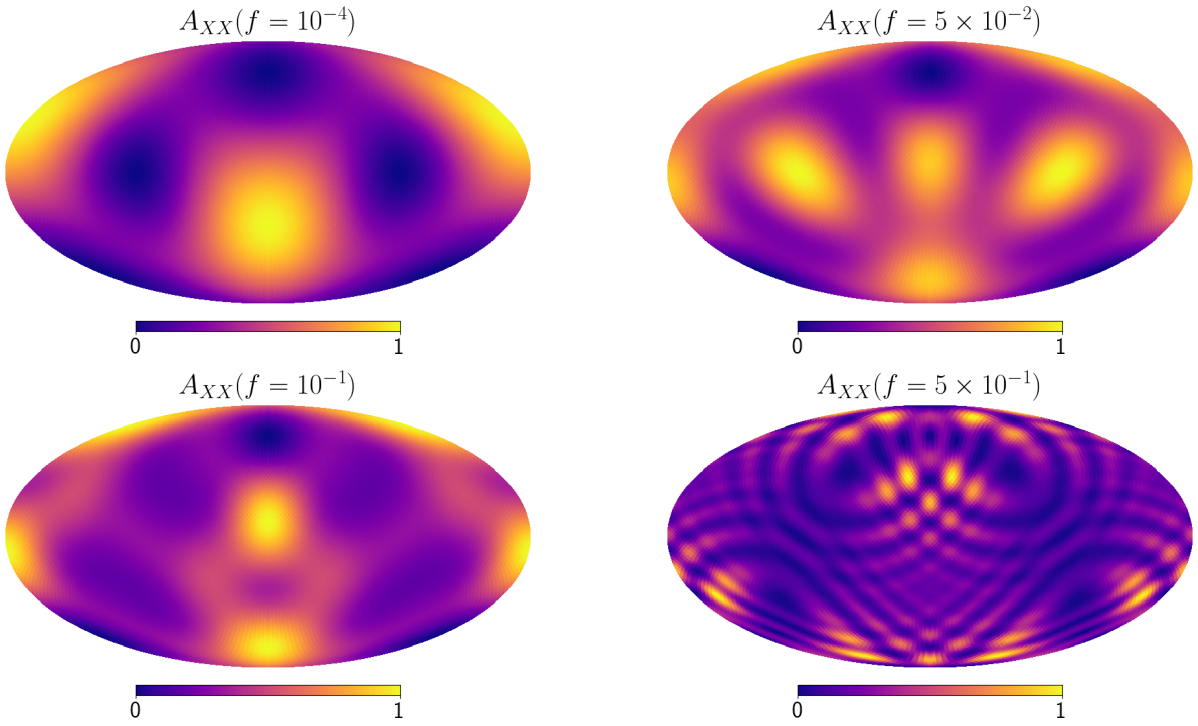


Figure 5.2: Normalised autocorrelated response of TDI channel X, A_{XX}, at time $t = 0$ and in the SSB reference frame, at frequencies $f = 10^{-4}$ Hz, $f = 5 \times 10^{-2}$ Hz, $f = 10^{-1}$ Hz, $f = 5 \times 10^{-1}$ Hz as illustrated in the panels.

having expanded in the polarisation bases and inserting the relation between the second order moments of the strain and the Stokes parameters as previously done in the case of ground-based measurements in Chapter 4. Note that the GWB considered here is stochastic also in polarisation; as such the Q, U, and V Stokes parameters average to zero. The quadratic response tensor $\boldsymbol{A}$ is simply constructed with the linear response TDI vectors $\boldsymbol{R}^P = (R_X, R_Y, R_Z)^P$,

$$\boldsymbol{A}^\tau = \boldsymbol{R}^+ \otimes \boldsymbol{R}^{+\star} + \boldsymbol{R}^\times \otimes \boldsymbol{R}^{\times\star} \,. \tag{5.16}$$

The τ label is recovered here to point out the time-dependence of the detector response, which is implicitly expressed in Equations (5.7-5.16) through the arm and positions vectors. Note also the directional and frequency dependence of $\boldsymbol{A}$, a sample of which is given in Figure 5.2. Here the continuous sky response has been discretised to produce the Mollweide projection in pixel space, $\boldsymbol{A}(\hat{\boldsymbol{n}}) \to \boldsymbol{A}_p$, where p labels a pixel on the sky, and is shown in the SSB reference frame. This is analogous to the approach used for ground-based detectors in Chapter 4, hence the notation is maintained. As in the case of LIGO, the sky response

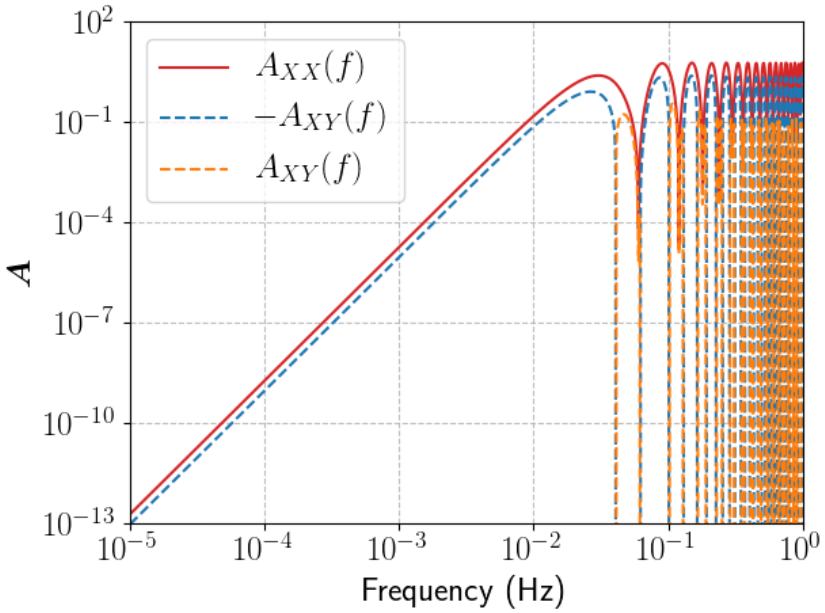


Figure 5.3: Sky-integrated auto-correlated and cross-correlated responses, A_{XX} and A_{XY}, across the ferquency spectrum. Note the trend in the low frequency limit is $\propto f^4$, which follows from taking the square of Equation (5.10).

gives rise to an inhomogeneous and non-compact antenna pattern. However, with the LISA TDI channels this pattern is also significantly dependent on the frequency probed, as may be observed in the different panels of Figure 5.2. It is straightforward to show that in the low frequency limit the spectral dependence of the quadratic response tensor $\boldsymbol{A}$ goes as f^4 [52], then starts deviating above 10^{-2} Hz, as may be observed in Figure 5.3. The oscillating behaviour seen here above $f \sim 0.05$ Hz is due to the fact that the probed GW wavelength becomes comparable to the arm length of the detector, or equivalently the GW frequency is comparable to the characteristic frequency of the detector; hence the maxima in A_{XX} occur at frequencies which are integer multiples of the characteristic frequency, and the minima correspond to the half integer multiples. The constant pattern and flat trend will set the limit to LISA's resolution at low frequency, while this substantial deviation and the high-ℓ pattern will prove crucial to obtain higher resolution at high frequency.

Although the map-maker developed here solves for a map directly in the pixel domain, it can be used to assess the ability of the cross-correlated TDI channels to reconstruct the

angular power spectrum of a background defined as

$$C_\ell^{\rm GW} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\left|\int_{S^2}\frac{d\hat{\boldsymbol{n}}}{4\pi}\,Y_{\ell m}(\hat{\boldsymbol{n}})\,I(\hat{\boldsymbol{n}})\right|^2. \tag{5.17}$$

The assessment of the sensitivity to $C_\ell^{\rm GW}$ induced by the response in spherical harmonic space,

$$\boldsymbol{A}_\ell = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\left|\int_{S^2}\frac{d\hat{\boldsymbol{n}}}{4\pi}\,Y_{\ell m}(\hat{\boldsymbol{n}})\,\boldsymbol{A}(\hat{\boldsymbol{n}})\right|^2, \tag{5.18}$$

is present in the literature. Studying the relative sensitivity of a single TDI channel to angular scales ℓ on the sky, the angular resolution of the LISA TDI 1.5 configuration is predicted to be constrained as $\ell \lesssim 10$ [35, 53]. This resolution will strongly depend on the frequency interval examined, such that for lower frequencies the cutoff will be much more severe. To see this in more detail, it is possible to follow the steps presented for ground-based measurements in Chapter 4 and recover the *full* spherical harmonic response, which does not require isotropisation. Firstly, the intensity field is expanded over basis functions $Y_{\ell m}(\hat{\boldsymbol{n}})$ in terms of coefficients $a^I_{\ell m}$,

$$I(f,\hat{\boldsymbol{n}}) = E(f)\sum_{lm} a^I_{\ell m}\,Y^\star_{\ell m}(\hat{\boldsymbol{n}}), \tag{5.19}$$

where the spectral dependence of the GWB intensity, $E(f)$, is explicitly separated out and is seen in Equation (2.69): $E(f) = (f/f_0)^{\alpha-3}$. The assumption here is the same made in Chapter 4 in the model-dependent analysis, i.e. that all of the frequency dependence of the sky signal is the same in each direction.

In the low frequency limit the spectral dependence of the quadratic response tensor $\boldsymbol{A}$ goes as f^4 [52], thus $\boldsymbol{A}$ may be expanded in spherical harmonics as

$$\boldsymbol{A}(f,\hat{\boldsymbol{n}}) = f^4\sum_{lm} A_{\ell m}\,Y^\star_{\ell m}(\hat{\boldsymbol{n}}) \tag{5.20}$$

Inserting this into Equation (5.15) and making use of the spherical harmonic expansion of the plane wave

$$e^{i\boldsymbol{k}\cdot\boldsymbol{x}} = 4\pi\sum_{\ell m} i^\ell\,j_\ell(kx)\,Y_{\ell m}(\hat{\boldsymbol{k}})\,Y^\star_{\ell m}(\hat{\boldsymbol{x}}), \tag{5.21}$$

where $j_\ell(x)$ are spherical Bessel functions, one obtains an expression for the observed harmonic coefficient $d^{AB}_{\ell m}$ by a pair of TDI channels A, B, defined via

$$d_{AB}(f) = f^4 E(f) \sum_{LM} d^{LM}_{AB}\, Y^\star_{LM}(\hat{\boldsymbol{b}}_{AB})\,, \tag{5.22}$$

as

$$d^{LM}_{AB} = i^\ell 4\pi \sum_{\ell m \ell' m'} a^I_{\ell m} \mathcal{K}^{LM}_{\ell m, \ell' m'}\,, \tag{5.23}$$

where $\boldsymbol{b}_{AB} = \boldsymbol{x}_A - \boldsymbol{x}_B$ is the arm connecting the spacecrafts related to TDI channels A and B, which has module $|\boldsymbol{b}_{AB}| = L$ in the equal arm approximation. The coupling kernel $\mathcal{K}^{LM}_{\ell m, \ell' m'}$ incorporates the Gaunt integral

$$\begin{aligned}
\mathcal{K}^{LM}_{\ell m, \ell' m'} &= j_\ell(2\pi f L) A^{LM}_{AB} \int_{S^2} d\hat{\boldsymbol{n}}\, Y_{LM}(\hat{\boldsymbol{n}})\, Y^\star_{\ell m}(\hat{\boldsymbol{n}})\, Y^\star_{\ell' m'}(\hat{\boldsymbol{n}})\,, \\
&= (-1)^{m+m'} j_\ell(2\pi f L) A^{LM}_{AB} \sqrt{\frac{(2L+1)(2\ell+1)(2\ell'+1)}{4\pi}} \times \\
&\qquad \begin{pmatrix} L & \ell & \ell' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} L & \ell & \ell' \\ M & -m & -m' \end{pmatrix},
\end{aligned} \tag{5.24}$$

where the terms in brackets indicate $3j$ coefficients and A^{LM}_{AB} are the harmonic coefficients for each entry of $\boldsymbol{A}$. The coupling kernel may be used to predict the resolution of the LISA TDI 1.5 configurations, isotropising but still taking into account the couplings imposed by Equation (5.24) and difference in resolution at high frequencies given by the structure such as may be seen in the panels of Figure 5.2. From $\mathcal{K}^{LM}_{\ell m, \ell' m'}$ it is possible to recover the C_ℓ coupling kernel $\mathcal{K}_{\ell \ell'}$ by integrating over the m phases and taking the spectrum of the response coefficients A^{LM}_{AB} and of the Bessel function j_ℓ. The normalised, eigen-weighted sum per mode of the eigenvectors of the kernel provides an estimate for the relative sensitivity to each mode, as may be seen in Figure 5.4. This estimate includes the frequency-integrated Bessel function, which couples the frequencies to the modes on the sky. It is clear that the resolution achievable with the LISA TDI channels is strongly dependent on the frequencies probed, and that for frequencies below 10^{-2} Hz the resolution cutoff will be $\ell < 5$ and thus extremely limiting. Higher frequencies are more promising, potentially reaching $\ell \sim 15$ in the 10^{-1} Hz range, however any frequency higher than that would be extremely challenging to

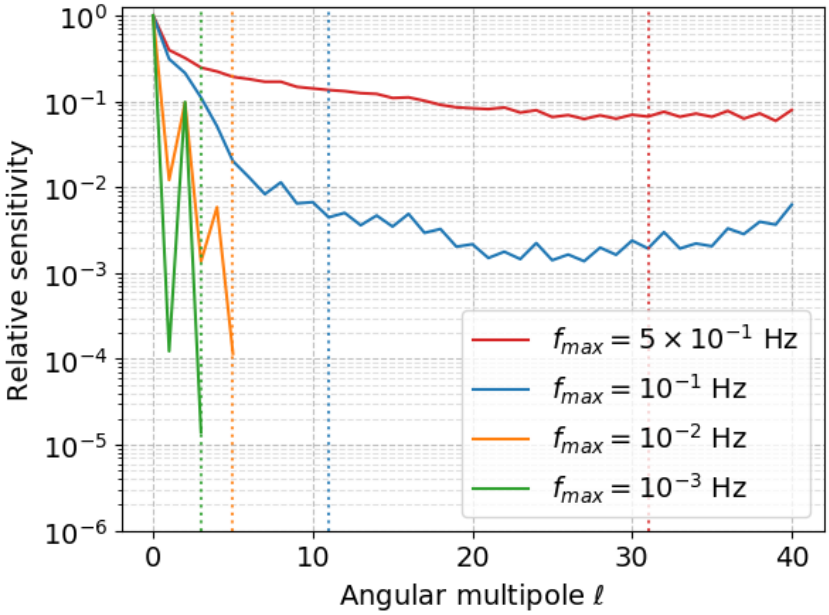


Figure 5.4: Normalised and cumulatively eigen-weighted eigenvectors per angular scale ℓ of the C_ℓ coupling kernel $\mathcal{K}_{\ell\ell'}$, frequency-integrated from 10^{-5} Hz up to different frequency cutoffs $f_{\rm max}$, for the autocorrelated channel X. The rough estimates obtained with the integrated Bessel function are shown as dashed lines, with the same colour coding.

probe given the expected noise curves (see discussion below). The j_ℓ itself can also be used for a back-of-the-envelope estimate of the angular resolution, as done in the previous Chapter for LIGO, however this does not take into account the structure in the antenna patterns which is substantially different with LISA. The values of the resolution cutoff $\ell_{\rm cut}$ where the broad-band integrated Bessel function goes to zero, with different frequency cutoffs, are shown as dashed lines in Figure 5.4 and are in good agreement with the $\mathcal{K}_{\ell\ell'}$ estimates at low frequencies, where the structure of $\boldsymbol{A}$ does not have a big impact, however they are considerably more pessimistic at higher frequencies, as is to be expected.

Just as in the case of ground-based interferometers, the non-uniform weighting of the sky introduces a coupling between ℓ, ℓ', and L in the spherical harmonic expansion, and the projection of plane waves on a sphere breaks the degeneracy of the odd-ℓ modes. The reconstruction of the GWB spherical modes should result in a well-conditioned inversion problem, providing sufficient angular coverage of the sky. LISA's phase coverage will be provided by the detector's cartwheeling motion around the Sun, thus a significant improvement in the conditioning is to be expected when integrating over at least six months. However, the inhomogeneity of the integrated beam is still strong, also due to the symmetry of the

coverage as the LISA plane remains at a constant angle with respect to the ecliptic throughout the year. As such, the conditioning of the problem remains an issue, and could only be solved with additional detectors on different orbits/planes.

For the simulations and study that follows the coordinate system and polarisation basis are fixed to the SSB frame as it is convenient for the description of both the spacecraft trajectories and the injected signal's directional dependence.

5.2 Optimal map solutions for LISA

The data output of the three TDI 1.5 channels $\{X, Y, Z\}$: $\boldsymbol{d} = (d_X, d_Y, d_Z)$ will be correlated. Furthermore, as argued in the previous Chapters, in the case of GWs both signal and noise components in the time-domain vanish in the ensemble mean. Transforming to the Fourier domain does not alleviate this problem unless the background itself is phase-coherent or if the signal is a single point source on the sky. Hence no residual can be defined since both $\langle \boldsymbol{d} - \langle \boldsymbol{d} \rangle \rangle = 0$ and $\langle |\boldsymbol{d}|^2 - \langle |\boldsymbol{d}|^2 \rangle \rangle = 0$, and the correct maximum-likelihood map solution is obtained by minimizing the full likelihood in Equation (3.19), which is detailed in Chapter 3. The data TDI vector is decomposed as

$$\boldsymbol{d}_{\tau,f} = \boldsymbol{s}_{\tau,f} + \boldsymbol{n}_{\tau,f} = \boldsymbol{R}_{\tau,f} h_f + \boldsymbol{n}_{\tau,f} \,, \tag{5.25}$$

where $\boldsymbol{n}$ is the noise TDI vector, $\boldsymbol{R}$ is the linear response TDI vector and there is an implicit sum over polarisations on the right-hand side. Assuming the noise is zero–mean and Gaussian with covariance $\boldsymbol{N}_f = \langle \boldsymbol{n} \otimes \boldsymbol{n}^\star \rangle$ and the signal component is also Gaussian with covariance $\boldsymbol{C}_{\tau,f} = \boldsymbol{A}\tilde{I} + \boldsymbol{N}$, the solution in Equation (3.27) is adapted to the case considered here as

$$\tilde{I}_p = \mathscr{F}^{-1}_{pp'} \cdot \frac{1}{2} \mathrm{Tr} \left[\sum_{\tau,f} \boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{I}_{p'}} \boldsymbol{C}^{-1}_{\tau,f} (\boldsymbol{D}_{\tau,f} - \boldsymbol{N}_f) \right] , \tag{5.26}$$

$$\mathscr{F}_{pp'} = \frac{1}{2} \mathrm{Tr} \left[\boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{I}_{p'}} \boldsymbol{C}^{-1}_{\tau,f} \frac{\partial \boldsymbol{C}_{\tau,f}}{\partial \tilde{I}_{p'}} \right] , \tag{5.27}$$

where formally $\boldsymbol{D}_{\tau,f} = \boldsymbol{d}_{\tau,f} \otimes \boldsymbol{d}_{\tau,f}{}^{\star}$. In this application, the trace is taken over the three TDI channels, and the sum is over all frequencies in the FFT and all observation times, to maximise sky coverage. One can also consider integrating over short frequency intervals or single frequencies to probe the frequency dependence of the GWB, within the limits imposed by the conditioning of the Fisher matrix. An iterative scheme is employed to reach the maximum likelihood estimate for I, starting with an initial guess $\tilde{I}_{\rm in}$ to plug into $\boldsymbol{C}$, obtaining a first estimate of $\tilde{I}_p$, then repeating until convergence. In this analysis the noise model $\boldsymbol{N}$ is assumed to be perfectly known, and the expectation is that when this method will be applied to real LISA data it will be possible to rely on an independent noise estimation which informs the gradient term. However, it is possible yet expensive to extend this estimator to include the noise as a free parameter, and solve for both noise and signal components directly. In this case also the noise would be estimated iteratively and it would be necessary to define convergence carefully in order not to bias the results.

It is important to note that in case the signal is non-Gaussian, the maximum-likelihood maps derived below remain unchanged, however the interpretation of their covariance would be affected by the presence of higher order terms in the underlying probability densities. Effectively, in this case the maps would cease to be optimal estimates as they would not correspond to the solution which minimises the variance. See e.g. [139] for a preliminary analysis of detection methods of non-Gaussianity in a GWB induced by short-duration signals, and [140] for a study of how to use higher-order cumulants to characterize its properties.

5.2.1 Data generation

The data is generated directly in the Fourier domain as the projection of an input map $I_{\rm in}$ plus a noise component, $\boldsymbol{D} = \boldsymbol{A}I_{\rm in} + \boldsymbol{N}_{\rm in}$, assuming the fully stochastic limit applies as in Equation (2.51). This is not a realistic realisation of each individual data segment but the assumption is that it will be valid once the full integral over time is performed. $\boldsymbol{D}$ is generated over the period of one year two-day segments at a time for different frequency bands. Each segment is heavily down-sampled in frequency to $\mathtt{N}_f = 200$ samples in order to lighten the computational load as much as possible without losing too much resolution in frequency, which would lead to a loss of information on the sky due to the coupling between f and $\hat{\boldsymbol{n}}$

via $\boldsymbol{A}$. $\boldsymbol{A}I_{\rm in}$ is the signal component of the data contracted over pixels and $\boldsymbol{N}_{\rm in}$ is a Gaussian realisation of the noise model $\boldsymbol{N}$. The resulting data streams for the three channels are then plugged into Equation (3.26), and the Fisher matrix is iteratively computed, where the covariance matrix $\boldsymbol{C}$ is the full covariance of the data and simply $\frac{\partial \boldsymbol{C}}{\partial I_p} = \boldsymbol{A}_p$. All the terms in Equations (5.26) and (5.27) are complex matrices and the reality condition imposed by the data is recovered when integrating over frequencies.

The nominal SNR of the signal component of the data is estimated by integrating over observation time and frequency range as follows:

$$\mathrm{SNR} = \sqrt{\Delta T_{\rm obs} \int_{f_{\rm min}}^{f_{\rm max}} df \left(\frac{AI_{\rm in}}{N}\right)^2}, \tag{5.28}$$

where A is the instantaneous response of channel X and N is the appropriate noise model.

The signal component in the data $\boldsymbol{A}I_{\rm in}$ is generated by simply scanning an input map for the GWB intensity on the sky $I_{\rm in}$ with the quadratic response of our TDI configuration. This is a completely deterministic calculation which relies on the assumptions that the signal is truly stochastic and obeys Equation (2.69), and that the overall integration over time and frequencies satisfies the ensemble average limit, such that Equation (2.51) holds.

Except where explicitly stated, the $I_{\rm in}$ maps scanned in the simulations below are random Gaussian realisations of an $\ell^2 C_\ell$ flat power spectrum with $\ell_{\rm max} = 20$, chosen uniquely for testing purposes with a high angular structure to fully scan the potential for reconstructing spatial anisotropies. The assumption is then that the monopole C_0 sets the size of the anisotropies which scale as a fixed power law $C_\ell = C_0/\ell^2$. In a more realistic modelling of the signal, one should rather introduce two separate scales: a monopole C_0, which would set the size of the SNR as in Equation (5.28), and a typical scale of anisotropies C_1. These two scales are simply assumed to be the same here. The ability of LISA to reconstruct these input maps is explored by changing the frequency integration range, varying the spectral parameter α as defined in Equation (2.69), and toggling the sky-integrated amplitude $I(f_0)$ with respect to the noise level set by N.

LISA TDI noise is an active area of research, and there are ongoing studies on which TDI configuration will yield the lowest noise measurements. However these go beyond the scope of this project, as these model the orbits and breathing modes of the constellation more

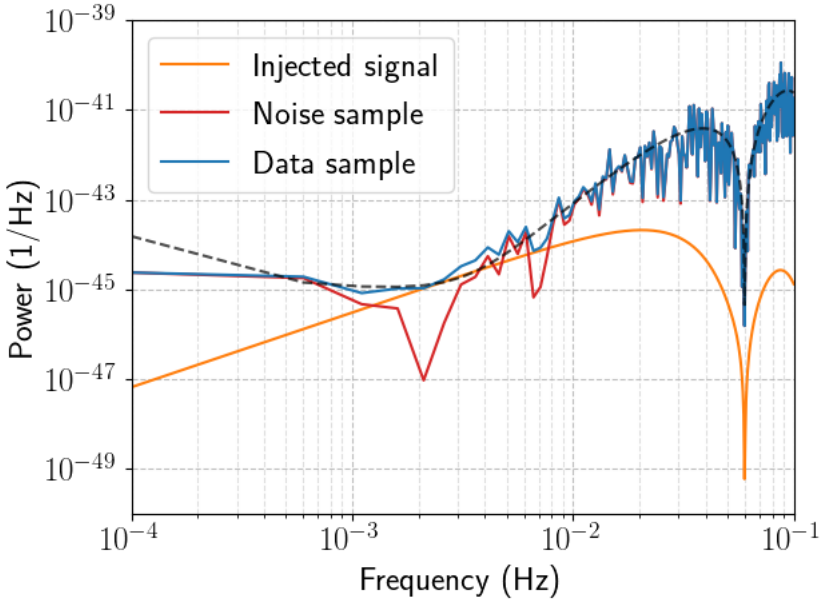


Figure 5.5: Data power sample from the autocorrelated X channel, decomposed into injected signal (orange line) and simulated noise (red line) components. The Injected signal has and SNR of 0.6, and is evidently buried in the noise. The sample represents the data accumulated over a 2-day period and then FFTed. The dashed grey line is the noise model which informs the noise generation, renormalised by a factor which accounts for the lower sampling rate.

realistically. In the simplified equal-arm scenario considered here, to good approximation the noise in each auto-correlation of TDI channels may be described by the same power spectrum, S_{CC}, and similarly the noise in each cross-correlation may be described by S_{CD}. The 3×3 model TDI noise matrix $\boldsymbol{N}$ is then completely described by S_{CC} on the diagonal and S_{CD} on the off-diagonal terms. The expressions for these power spectra used in this work are

$$S_{CC} = 16 \sin^2 a \, (S_{\text{int}} + (3 + \cos 2a) S_{\text{acc}}) \,, \tag{5.29}$$

$$S_{CD} = -8 \sin^2 a \cos a \, (S_{\text{int}} + 4 S_{\text{acc}}) \,, \tag{5.30}$$

as presented in [135], where S_{int} and S_{acc} are the interferometer and acceleration noise-components respectively.

The noise realisation $\boldsymbol{N}_{\text{in}}$ must respect the correlations imposed by $\boldsymbol{N}$, hence the noise is first generated linearly in the noise-diagonal space, and then rotated back into correlated noise space. This is achieved by generating a random 3-vector in noise-diagonal space and rotating it into the TDI noise vector $\boldsymbol{n}_{\text{in}} = (n_X, n_Y, n_Z)_{\text{in}}$ using the eigenvector matrix of $\boldsymbol{N}$.

$\boldsymbol{N}_{\rm in}$ is then simply the outer product

$$N_{\rm in} = \boldsymbol{n}_{\rm in} \otimes \boldsymbol{n}_{\rm in}^{\star}\,. \tag{5.31}$$

An example data segment, decomposed into the noise and signal components, is provided in Figure 5.5.

5.2.2 Mapping tests

This Section covers the map reconstruction tests based on the maximum-likelihood solution in Equations (5.26) and (5.27). These serve as a proof of concept for the recipe, which is the first stepping stone towards testing it on mock LISA timestreams. For these tests, convergence is defined by the fluctuation of the monopole value of the map, $\delta \bar{I}_{\rm out} < 0.1\%$. This is quite a loose definition of convergence and it may be preferable to set up a pixel-by-pixel test, however it is found to be sufficient here for these preliminary studies.

A variety of simulations has been carried out to extensively probe the ability of the TDI configurations $\{X, Y, Z\}$ to reconstruct an anisotropic GWB, both in high and low signal-to-noise ratio (SNR) scenarios. For all the maps discussed below the inversion problem is strongly ill-conditioned, hence the Fisher matrix is inverted using the singular value decomposition technique as in similar work done with LIGO data [3, 4]. Specifically, the condition number of the Fisher matrices in the cases presented below, which are all of dimension $N_{\rm pix}^{\rm out} \times N_{\rm pix}^{\rm out}$, is of order 10^{18}. In the case of high SNR the output maps mildly depend on the conditioning imposed at pseudo-inversion as the signal will dominate the information, whereas in the low SNR scenario the output map is extremely dependent on the conditioning. In these tests, the choice made was to monitor the monopole level of the output map and gauge the conditioning such that it would match the input. This technique may be employed also with real data, where the measurement of the monopole can be done independently of the inversion problem.

To test the limits of the geometric set-up of LISA, in the first set of simulations the generated data contains anisotropic, flat spectrum backgrounds ($\alpha = 3$) of effectively infinite SNR and the output maps are reconstructed truncating the frequency integration at different

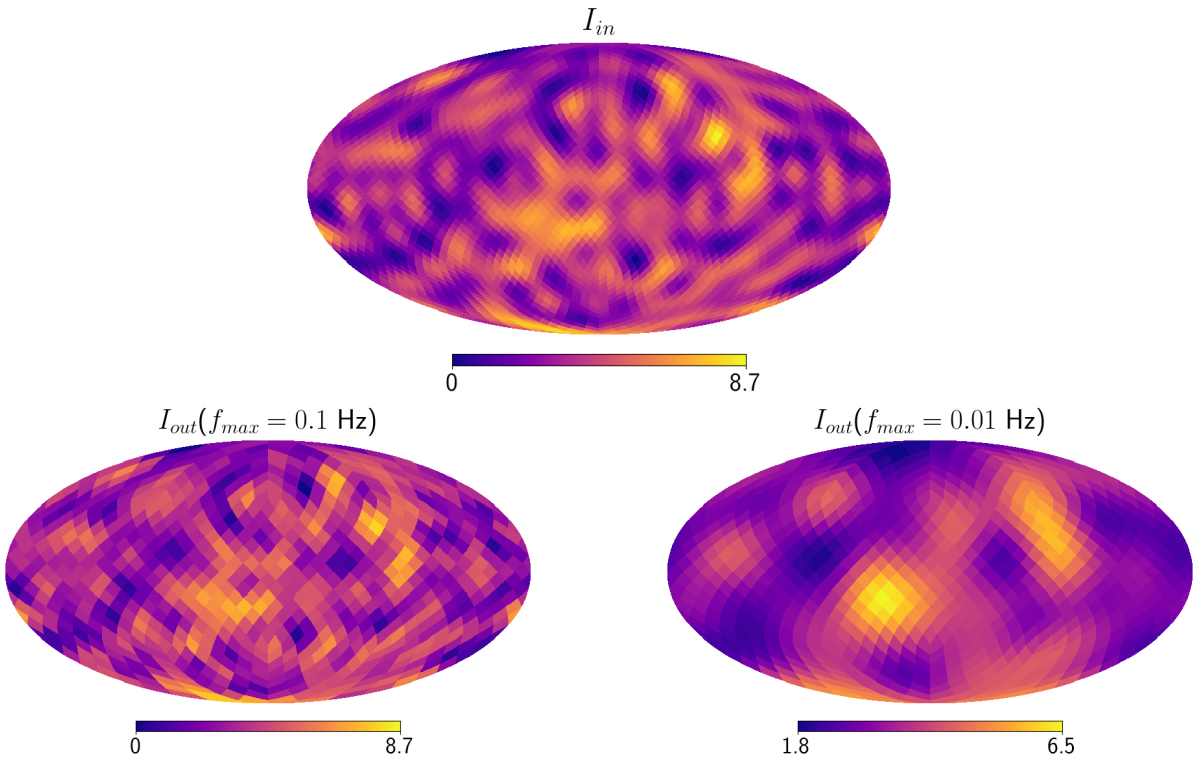


Figure 5.6: An example input map from the simulations (top panel) to be compared to the final output maps obtained integrating with different frequency cutoffs, $f_{\rm max} = 0.1$ Hz (bottom-left panel) and $f_{\rm max} = 0.01$ Hz (bottom-right panel). These highlight the different resolutions the LISA channels have in different ranges of frequency.

values of $f_{\rm max}$. This value of α is chosen as it best describes the shape of LISA noise in the relevant frequency range and thus optimally down-weights the noise in the signal search. In this extremely high signal scenario, the solution converges after a single iteration. As may be seen in the example maps in Figure 5.6, the input map is reconstructed remarkably differently in the case of $f_{\rm max} = 10^{-1}$ Hz and $f_{\rm max} = 10^{-2}$ Hz. The higher angular modes are well preserved when allowing the estimator to integrate up to higher frequencies, where there is finer structure in the response pattern, whereas they are aliased into lower modes when integrating only over the lower frequencies.

To study this trend four sets of 50 analogous simulations have been run with $f_{\rm max} = 10^{-3}$ Hz, 10^{-2} Hz, 5×10^{-2} Hz, 10^{-1} Hz respectively and the average output C_ℓs in each set has been calculated. The transfer functions

$$T_\ell = \frac{C_\ell}{C_\ell^{\rm in}} \tag{5.32}$$

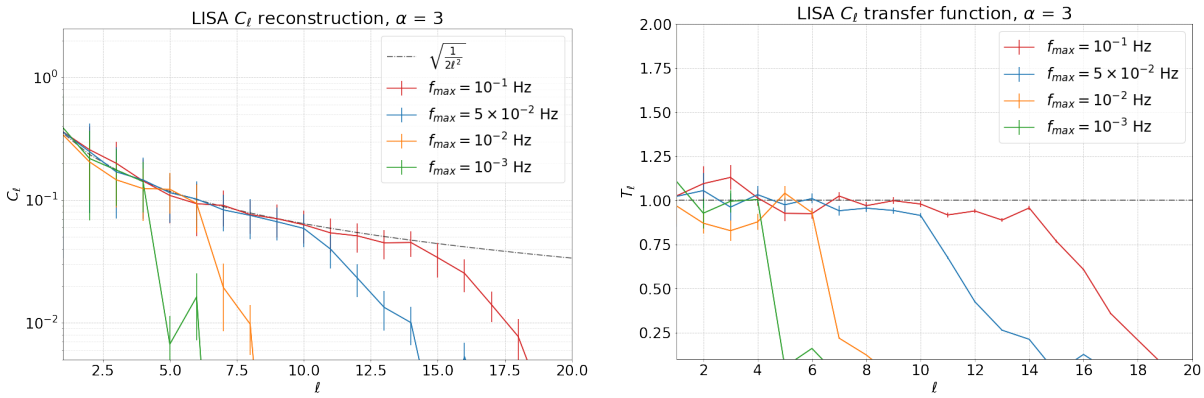


Figure 5.7: Average reconstructed C_ℓ trends (left) and resulting transfer functions T_ℓ (right) obtained with different frequency cutoffs. Each simulation set consists of 50 maps, each a different realisation of the same C_ℓ input, shown as the grey dashed line in the left panel. There appears to be a clear one-to-one relation between the resolution of the instrument and the frequency cutoff.

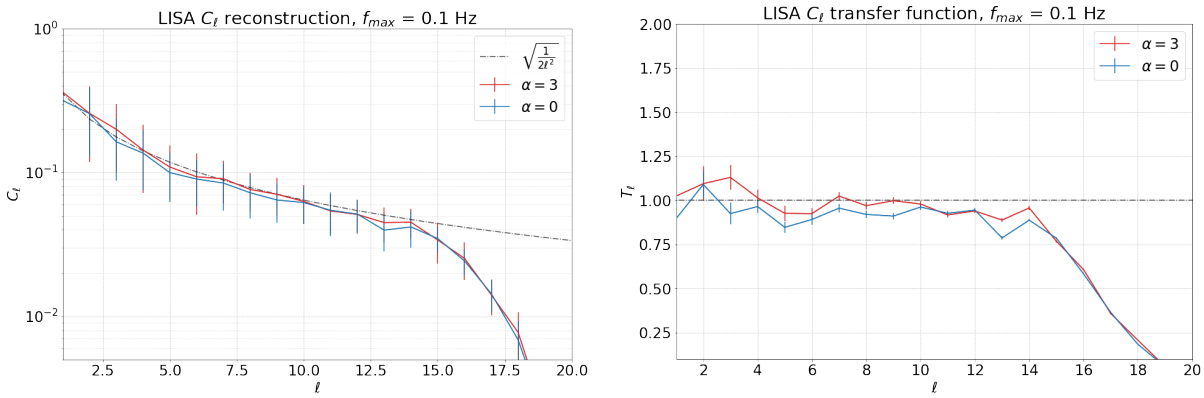


Figure 5.8: Average reconstructed C_ℓ trends (left) and resulting transfer functions T_ℓ (right) obtained with different spectral shapes, $\alpha = 3$ and $\alpha = 0$, both in the high frequency case $f_{\text{max}} = 0.1$ Hz. Each simulation set consists of 50 maps, each a different realisation of the same C_ℓ input, shown as the grey dashed line in the left panel. There is an average difference of 5% between the two transfer functions, however this does not affect the resolution cutoff.

are also computed for each set. The C_ℓ reconstruction and comparison between transfer functions with different frequency cutoffs may be seen in Figure 5.7.

Different spectral shapes are reconstructed slightly differently, as may be seen in the right panel of Figure 5.8 where the high frequency $\alpha = 3$ set is compared with a high frequency $\alpha = 0$ set of simulations. The C_ℓ reconstruction and resulting transfer functions differ at certain modes, on average by 5%, but the cutoff ℓ_{cut} appears to be the same, hence it appears the frequency weighting in the signal and reconstruction do not have a strong impact on the resolution of the reconstructed maps.

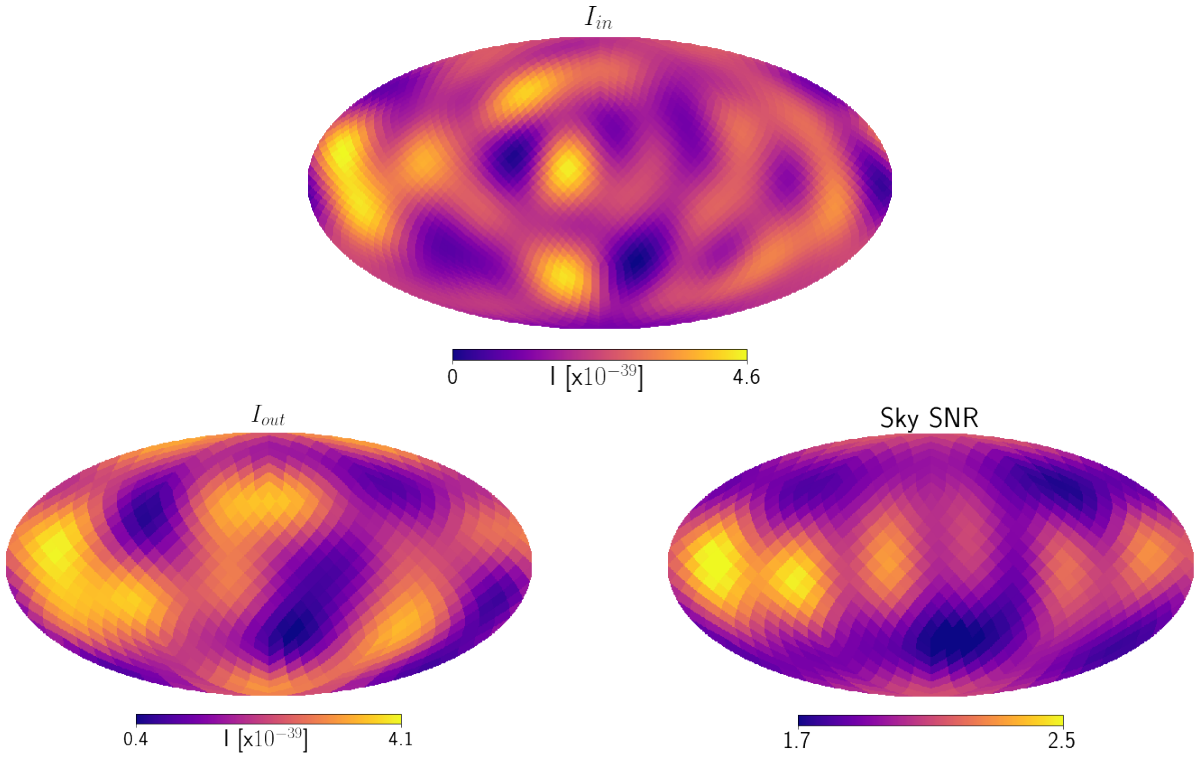


Figure 5.9: Input (left panel) and reconstructed output (central panel) maps in a low SNR case. The input map is an $\ell^2 C_\ell$ flat gaussian realisation with $\ell_{\rm max} = 10$. The output map is a heavily smoothed version of the input, with some loss of power due to the strong conditioning required. The SNR of this particular reconstruction is shown in the right panel.

An example of map reconstruction in the presence of high noise is shown in Figure 5.9. The map in input here is a ℓC_ℓ flat Gaussian realisation with maximum angular scale $\ell_{\rm max} = 10$ to make the visual comparison easier. The spectral shape of the signal is $\alpha = 2/3$, which together with the power spectrum is typically associated with an astrophysical inspiral-dominated GWB [72] which traces the large scale structure [87]. The signal component has an SNR of 0.6 and is thus buried under the noise, as may be seen in the specific data sample in Figure 5.5. The spectral shape of the signal is $\alpha = 2/3$, which is typically associated with an astrophysical inspiral-dominated GWB [72], and the maximum angular scale of the input realisation is $\ell_{\rm max} = 10$ to make the visual comparison easier. As may be observed in Figure 5.9, the presence of loud noise has a similar effect as having a lower frequency cutoff, as the noise dominates at higher frequencies. Additionally, severe conditioning is required to cut out the noisier high modes, and multiple iterations are necessary to reach convergence. Specifically, the final conditioning cuts 97% of the total eigenvalues of the Fisher information matrix. This implies that only 3% of the modes in the Fisher matrix actually carry signal, whereas the

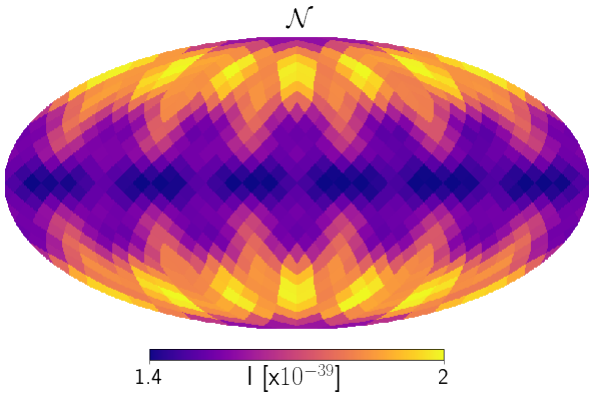


Figure 5.10: Sky distribution of the noise in the SSB frame. Note the imprinted 6-fold symmetry of the orbit, given by the three concentric orbits of the spacecrafts.

remaining modes are pure noise. Because noise dominates the higher frequency modes, and these are associated to smaller angular scales via the response (see e.g. the coupling kernel in Equation (5.24)), this loss of information directly translates into loss of resolution on the sky. The value of the monopole of the reconstructed map however does not depend strongly on the conditioning and is in good agreement with that in input. Thus the resulting heavily conditioned output map appears to be a low-ℓ smoothed version of the initial map with some loss of power. The noise distribution on the sky, $\mathcal{N} = \mathscr{F}^{-1/2}$, is also recovered as the diagonal of the inverse Hermitian square root of the converged Fisher matrix and is shown in Figure 5.10. This is an incomplete representation of the noise as the off-diagonal correlations are ignored, however it already shows that there is more noise off the ecliptic, as expected.

Finally, an example of reconstruction of a galaxy-like sky signal both with and without noise is shown in Figure 5.11. The input signal has spectral shape $\alpha = 2/3$ in the frequency range $[10^{-4}, 5 \times 10^{-3}]$ Hz, as suggested in [141], and the directional dependence is that of the Milky Way's white dwarf binary (WDB) population from [142]. The high SNR case is extremely unrealistic but provides a useful visual test of what the best resolution of LISA can be for a confused GWB at WDB frequencies. By relaxing the conditioning and including the higher modes in the Fisher matrix before inverting, the resolution can be pushed slightly higher however the pixels appear noisier, as may be observed in the top right panel of Figure 5.11. The bottom panel shows a more realistic SNR case, calculated based on the model of the white dwarf galaxy used in the LDC Radler data. This test proves, as in the simulation discussed above, that strong conditioning is required to recover the input distribution and that this

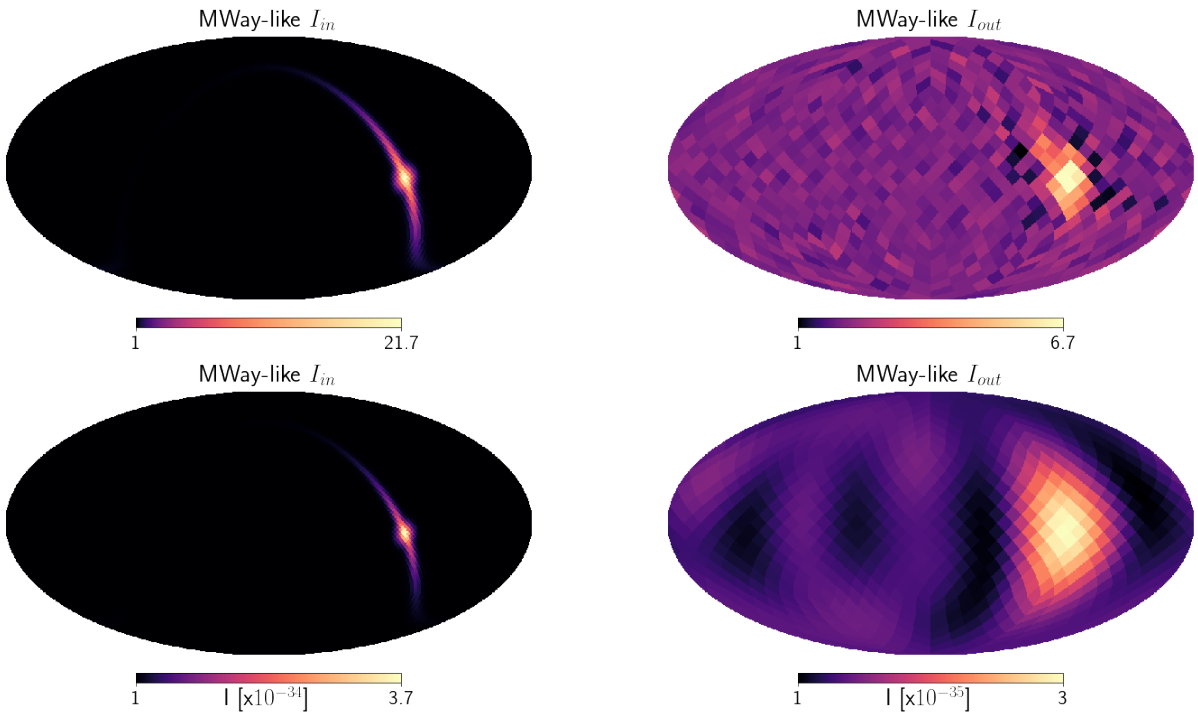


Figure 5.11: Sky inputs (left panels) and reconstructed outputs (right panels) for a Milky Way - like GWB distribution. The signal is injected in the range $[10^{-4}, 5\times10^{-3}]$ Hz. The top row is an extremely high SNR case, whereas the bottom row is a more realistic SNR = 52, hence it emerges above the noise around 10^{-3} Hz. Note that the value of the SNR itself can not be translated directly into a reliable measure of the resolution as it is an integral over a large frequency range as per Equation (5.28). These Mollweide projections are presented in log scale.

leads to the exclusion of higher modes in the Fisher matrix. The level of the noise in this case is the same as in the one above, however as the signal is lower in particular at the higher end of the frequency range the resolution of the galactic bulge image is very low. The monopole of the signal is very well recovered in both cases.

The mapping algorithm for the LISA GW detector presented here is the first of its kind, based on the maximisation of the full likelihood and involving all three correlated TDI channels. The tests discussed above clearly show there is a one-to-one relationship between the resolution of the instrument and the frequency cutoff in the trace, and this is not particularly dependent on the input signal's spectral shape nor the spectral weighting assumed in the reconstruction. This frequency cutoff will naturally depend on both the nature of the signal and the level of the noise. As may be read off of Figure 5.7, the maximum angular resolution for LISA TDI 1.5 channels is $\ell \lesssim 15$ in the best case scenario when the signal is loud and

clear over the noise up to $f_{\rm max} = 10^{-1}$ Hz. A possible candidate for this type of signal is a stochastic background of astrophysical origin from stellar mass black hole and neutron star binaries [132]. The outlook is a lot less optimistic for lower frequency signals which are strong at $f < 10^{-2}$, for which the resolution is $\ell \lesssim 7$. An intermediate signal, which should peak around $f \sim 5 \times 10^{-2}$ [141], is the galactic background from binary white dwarfs. Considering this type of signal to be described by a simple power law in frequency, we have tested its reconstruction on the sky and have found that the strongest limit is set by the strong noise modes at high frequency. The reconstruction of low SNR signals is similarly limited, even if they are present at high frequencies, due to the shape and nature of the noise.

In the tests presented above the size of the monopole sets both the SNR and the anisotropy level of the injected signal. However a more realistic modelling would see the introduction of two separate scales, a monopole C_0 and a typical anisotropy scale C_1. This would allow to test separately the detectability of background components with similar monopole values but different levels of anisotropy. For example, astrophysical extra-galactic backgrounds are expected to have a level of anisotropy $\delta\Omega/\bar{\Omega} = 10^{-3}$ whereas cosmological backgrounds should present CMB-like ansotropies of $\delta \sim 10^{-5}$ [29, 87]. Hence, cosmological backgrounds from the early Universe are expected to have a low SNR and a highly suppressed scale invariant anisotropy spectrum. These include inflationary backgrounds which seed primordial black holes [77], post-inflationary backgrounds due to strong first order phase transitions [13] or cosmic defect networks [143]. Map reconstruction for these background components will be quite challenging. The most promising candidate remains an astrophysical galactic background, hence the preliminary study carried out here. The LDC has produced mock time-domain data for a galactic background made up of thousands of superposed white dwarf binary inspirals which will be sent through this pipeline soon.

The results obtained in this paper have been derived under the ideal assumption of a perfectly known noise PSD model. A natural expansion of this method would be to parametrise a given uncertainty in the noise model and perform parameter estimation on the noise and signal parameters simultaneously. This can happen by fitting noise parameters for the acceleration and interferometer components for each cross-correlated channel or in each individual TDI channel. A technique for simultaneous reconstruction of the monopole component of

the signal and a simplified mode of the noise power spectra, parametrised by acceleration and optical metrology system contributions only, has been implemented in [133, 134], and more recently in [51, 52].

Another significant limit of this study is the assumption that the GWB intensity $I(f, \hat{\boldsymbol{n}})$ can be reduced to two independent components: a simple spectral shape and a constant pattern on the sky. The validity of this assumption will be explored in future work, studying specific GWB signals in the LISA band and analysing the time-domain generated LDC data containing anisotropic backgrounds. The results of this mapping algorithm may also be used to assess the detectability of different types of anisotropic backgrounds in the LISA band by studying the effective sky sensitivity of the detector and comparing it to the expected levels of anisotropy. This will be investigated further in [144].

Chapter 6

Conclusions and Future Prospects

This thesis has explored the topic of GW map-making, offering different methods to obtain maps of the GWB and applying these to both simulated and real data from different types of detectors. It is found that currently it is only possible to obtain directional upper limits for a GWB from the Advanced LIGO data, and that the expected levels of anisotropy of astrophysical and cosmological backgrounds remain, for now, out of reach. Furthermore, although not investigated in depth here the stochastic searches performed with PTAs continue to yield upper limits and no positive detection – see for example results from the Nanograv collaboration [15]. Even though the tools developed here have not yet shown their fruits, there is reason to believe detection may be just around the corner, and beyond that there are exciting new projects ahead which may be sensitive to stochastic backgrounds, presenting a new set of detection challenges.

On a short time-scale, the second generation (2G) Advanced LIGO and Virgo interferometers will be updated to reach *design* sensitivity, with an improvement of up to an order of magnitude in the noise threshold, such that the predicted SGWB from extra-galactic compact binary systems involving neutron stars and black holes will be within reach of detection [48, 145]. Further improvements of the lasers and mirrors are planned at the LIGO facilities, which should reach the *A*+ class by 2024 almost doubling Advanced LIGO's sensitivity. As models for the astrophysical GWB anisotropies differ, it remains unclear whether these would be measured even in the case of a positive detection of the background monopole, and what information they would convey given the strong local and shot-noise effects [101, 123].

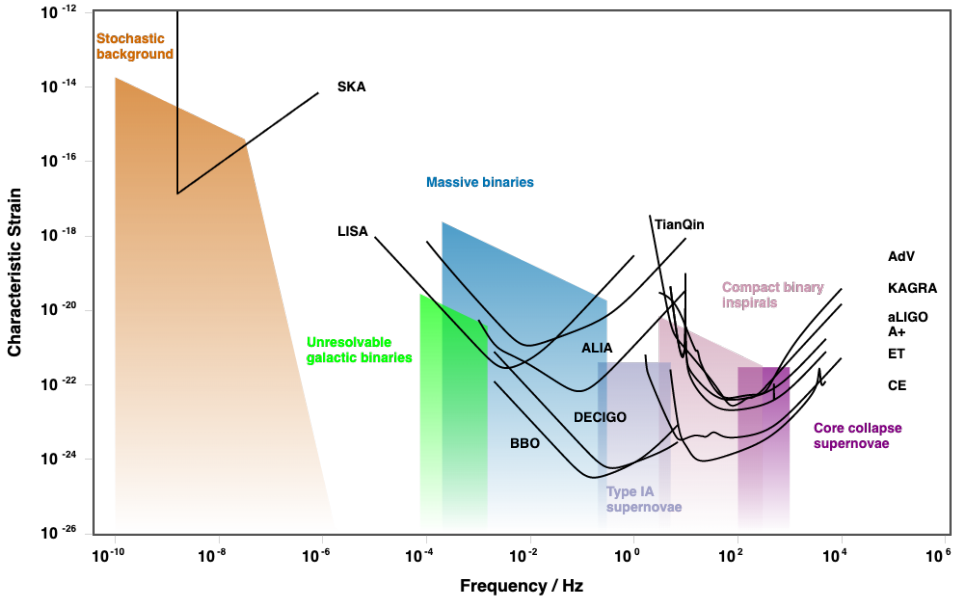


Figure 6.1: Expected gravitational-wave detector sensitivities and characteristic signals in their relative frequency bands from [19]. From left to right: sensitivity of the Square Kilometre Array (SKA) to millisecond pulsar timing residuals; LISA, TianQin, ALIA, DECIGO and BBO sensitivity curves probing the mid-range; the Advanced Virgo (AdV), KAGRA, Advanced design LIGO (aLIGO), A+ LIGO (A+), ET and CE curves all probing the higher end of the frequency range.

The addition of KAGRA and LIGO-India by 2023 [146] will certainly improve sensitivity and sky localisation for resolved GW sources, however these are not expected to have an immediate impact on stochastic directional searches as they will enter the network with a lower sensitivity than other detectors.

The big change for observations of stochastic signals from the ground is expected to arrive with 3G detectors. These should outperform 2G by two orders of magnitude in the strain sensitivity [147], potentially resolving the vast majority of black hole binary coalescences in the observable Universe [28] and scanning the astrophysical stochastic background extensively, albeit with limited resolution in the case of a single detector. There are currently two independently proposed 3G detectors, ET and CE, which have been considered both as stand-alone observatories and as part of an *extended* ground-based network; see [6] for a discussion on the noise power spectrum and angular resolution in this case. Primordial signals may be accessible to 3G detectors and if present in the timestream will be buried under the astrophysical foregrounds and incoherent background, hence to detect them it is paramount that all higher SNR signals are well understood. The work on SGWB detection

presented here is readily extended to 3G detectors, and provides an efficient and versatile estimator of astrophysical backgrounds in the 2G and 3G frequency band. It will also form the basis of a "component separation method" which will be essential for distinguishing between different backgrounds and attempting detection of any primordial component. The simultaneous estimation of frequency and angular dependence of any signal, such as the preliminary method laid out in Section 4.3.3, will be central to future component separation efforts.

Space observatories of GWs are also an exciting prospect for stochastic measurements, and present similar challenges to ground-based 3G detectors. As highlighted in this thesis, measuring the GWB with LISA requires a different approach than what has been used for the LIGO network, mainly as the LISA TDI channels may not be considered as uncorrelated detectors. Furthermore, resolution with a single LISA constellation is very limited at low frequencies especially, and conversely measurements at higher frequencies are obstructed by the expected noise levels [6]. TianQin, a Chinese mission analogous to LISA, is also expected to launch in the 2030s and is expected to yield similar results [148]. Other, unconfirmed space-based detector proposals include the Deci-hertz Interferometer Gravitational wave Observatory (DECIGO), a Fabry-Perot cavity interferometer which would bridge the frequency gap between Earth and Space probing the range around 1 Hz, and the Advanced Laser Interferometer Antenna (ALIA) considered in [149], a LISA-like constellation presenting a shorter arm-length than LISA and hence probing higher frequencies ($\sim 10^{-1}$ Hz). The ideal set-up to measure anisotropies in the GWB from space would be to have a *space network* comprised of at least two constellations, as detailed in [5] where a case is made for a new ESA mission. Space-based networks preliminarily studied in the literature are the LISA-TianQin network; ALIA in Stereo (ALIAS), comprised of a pair of ALIA detectors, one ahead and one following the Earth in its orbit; and the Big Bang Observer (BBO) first proposed as a NASA mission in [58] and studied for example in [150]. BBO is the most ambitious project considered to date, and would be comprised of four ALIA-like constellations all centred in the same orbit: two arranged in a six-pointed *star* constellation, and the other two *outrigger* constellations preceding and following the star by $2\pi/3$, respectively [149]. This intricate design and shorter arm length would make it most sensitive to the frequency band 0.03−3 Hz,

which is unaffected by the galactic white dwarf binaries that will plague the LISA band. As argued in [150], BBO should directly resolve all the compact object coalescence signals out to high redshift (at least z$\sim$ 5) which may then be removed from the timestream and stochastic techniques such as those explored in this thesis may be employed to estimate the monopole and anisotropies of a primordial signal. The outrigger constellations give BBO its angular resolution, which is useful for characterizing and removing the merging compact binary foreground: each source's angular position is triangulated on the sky using the differences in arrival times of emitted GWs at the antipodal constellations. A summary plot with the sensitivity curves of all the detectors listed above is provided in Figure 6.1, including typical GW signals in each range.

Map-making with GWs, given the variety of sources observed and their scatter in redshift space, will also become the cornerstone of multi-messenger cosmography. While multi-messenger observations with GWs are typically of resolved sources and require an electromagnetic counterpart to determine sky localisation [151], there are also interesting prospects for multi-messenger stochastic GW measurements. By LISA mission launch astrophysical and cosmological surveys carried out with innovative instruments such as Euclid, Gaia, the Large Synoptic Survey Telescope (LSST) and the Wide Field Infrared Survey Telescope (WFIRST) will be available to the community, and will be both instructive and complementary to science with LISA [152, 153]. The LISA-Gaia-LSST synergy will be particularly fruitful for Milky Way observations, as these will monitor a large set of white dwarf binary systems and it will be possible to measure system parameters with unprecedented precision [142]. Cross-correlation between large scale structure surveys and GW surveys is already being considered [154], and precision GW map-making is an essential step in this direction. Also cross-correlation with the CMB is a new topic of study, for example in [155] an analysis is laid out for the case where non-trivial correlations in the primordial matter field generate non-Gaussian effects which imprint specific anisotropies on both the CMB and the primordial SGWB. Another prospective background to analyse in conjunction with the astrophysical GWB is the diffuse supernova neutrino background [156], which may gradually be detected in the new *Gadolinium phase* of the Super-Kamiokande and Hyper-Kamiokande experiments [157]. This prospect is particularly exciting in the case of 3G detectors, as the expected astrophysical

stochastic background may very well include a significant supernova signal, if all nearby binary and inspiral signals are resolved and removed from the timestream.

It is clear that this is just he beginning of GWB analysis and stochastic searches, and the future holds plenty of opportunities to apply and refine the tools developed in this thesis. GW astronomy and cosmology has only just been born, and the best is yet to come.

Acknowledgements

I would like to acknowledge my supervisor, Carlo Contaldi, for all his guidance and support, and for the brilliant ideas shared. I would also like to acknowledge the Imperial College Theory Group and Astrophysics Group students and staff for many useful discussions and endless tea breaks. This PhD was funded by an Imperial College Schrodinger Fellowship. 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), 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.

I would like to thank my parents and scientific beacons, Laura Greggio and Alvio Renzini, for their advice, endorsement, and their bottomless repository of constructive criticism.

Finally, I would like to thank my life partner, Tommaso Boschi, for his patience, kindness, and boundless computing skills.

Bibliography

[1] A. I. Renzini and C. R. Contaldi, "Mapping Incoherent Gravitational Wave Backgrounds", *Monthly Notices of the Royal Astronomical Society* **481** no. 4, (2018) 4650–4661.

[2] C. R. Contaldi *et al.*, "Maximum likelihood map-making with the Laser Interferometer Space Antenna", arXiv:2006.03313 [astro-ph.CO].

[3] A. I. Renzini and C. R. Contaldi, "Gravitational Wave Background Sky Maps from Advanced LIGO O1 Data", *Physical Review Letters* **122** no. 8, (2019) 081102.

[4] A. Renzini and C. Contaldi, "Improved limits on a stochastic gravitational-wave background and its anisotropies from Advanced LIGO O1 and O2 runs", *Physical Review* **D100** no. 6, (2019) 063527.

[5] J. Baker *et al.*, "High angular resolution gravitational wave astronomy", arXiv:arXiv:1908.11410 [astro-ph.HE]. *ESA Voyage 2050 White Paper*.

[6] D. Alonso *et al.*, "The N_ℓ of gravitational wave background experiments", arXiv:2005.03001 [astro-ph.CO].

[7] **LIGO Scientific Collaboration and Virgo Collaboration**, B. Abbott *et al.*, "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs", *Physical Review X* **9** no. 3, (2019) 031040, arXiv:1811.12907 [astro-ph.HE].

[8] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "GW150914: The Advanced LIGO Detectors in the Era of First Discoveries", *Physical Review Letters* **116** no. 13, (2016) 131103, arXiv:1602.03838 [gr-qc].

[9] T. Accadia *et al.*, "Status of the Virgo project", *Classical and Quantum Gravity* **28** no. 11, (June, 2011) 114002.

[10] **LIGO Scientific Collaboration and Virgo Collaboration**, R. Abbott *et al.*, "GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses", arXiv:2004.08342 [astro-ph.HE].

[11] **LIGO Scientific Collaboration and Virgo Collaboration**, B. Abbott *et al.*, "GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_\odot$", *Astrophysical Journal Letters* **892** (2020) L3, arXiv:2001.01761 [astro-ph.HE].

[12] T. Regimbau, "The astrophysical gravitational wave stochastic background", *Research in Astronomy and Astrophysics* **11** no. 4, (2011) 369.

[13] C. Caprini and D. G. Figueroa, "Cosmological backgrounds of gravitational waves", *Classical and Quantum Gravity* **35** no. 16, (Jul, 2018) 163001. http://dx.doi.org/10.1088/1361-6382/aac608.

[14] M. E. L. Barsotti, P. Fritschel and S. Gras, "Advanced ligo *design* sensitivity curve", https://dcc.ligo.org/LIGO-T1800044/public.

[15] S. R. Taylor *et al.*, "Are We There Yet? Time To Detection of Nanohertz Gravitational Waves Based on Pulsar-Timing Array Limits", *The Astrophysical Journal* **819** no. 1, (2016) L6, arXiv:1511.05564.

[16] S. Burke-Spolaor *et al.*, "The Astrophysics of Nanohertz Gravitational Waves", *Astron. AstroPhysical Review* **27** no. 1, (2019) 5, arXiv:1811.08826 [astro-ph.HE].

[17] G. B. Hobbs *et al.*, "Gravitational-wave detection using pulsars: Status of the Parkes Pulsar Timing Array project", *Publications of the Astronomical Society of Australia* **26** no. 2, (2009) 103–109, arXiv:0812.2721.

[18] C. M. F. Mingarelli *et al.*, "Characterizing gravitational wave stochastic background anisotropy with pulsar timing arrays", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **88** no. 6, (2013) 1–18, arXiv:1306.5394.

[19] C. Moore, R. Cole, and C. Berry, "Gravitational-wave sensitivity curves", *Classical Quantum Gravity* **32** no. 1, (2015) 015014, arXiv:1408.0740 [gr-qc].

[20] J. Cruz *et al.*, "Laser Interferometer Space Antenna Simulator", *Laser Physics* **15** no. 7, (2017) 1056–1061, arXiv:1702.00786.

[21] P. Amaro-Seoane *et al.*, "Laser interferometer space antenna", arXiv:1702.00786 [astro-ph.IM].

[22] **Planck Collaboration**, R. Adam *et al.*, "Planck 2015 results. I. Overview of products and scientific results", *Astronomy and Astrophysics* **594** (2016) A1, arXiv:1502.01582 [astro-ph.CO].

[23] M. White *et al.*, "Interferometric Observation of Cosmic Microwave Background Anisotropies", *Astrophysical Journal* **514** (Mar., 1999) 12–24, arXiv:9712195 [astro-ph].

[24] N. W. Halverson *et al.*, "DASI first results: A Measurement of the cosmic microwave background angular power spectrum", *Astrophysical Journal* **568** (2002) 38–45, arXiv:0104489 [astro-ph].

[25] S. T. Myers *et al.*, "A fast gridded method for the estimation of the power spectrum of the CMB from interferometer data with application to the cosmic background imager", *Astrophysical Journal* **591** (2003) 575–598, arXiv:0205385 [astro-ph].

[26] B. Sathyaprakash *et al.*, "Scientific Potential of Einstein Telescope", in *46th Rencontres de Moriond on Gravitational Waves and Experimental Gravity*, pp. 127–136. 8, 2011. arXiv:1108.1423 [gr-qc].

[27] **LIGO Scientific Collaboration**, B. P. Abbott *et al.*, "Exploring the Sensitivity of Next Generation Gravitational Wave Detectors", *Classical Quantum Gravity* **34** no. 4, (2017) 044001, arXiv:1607.08697 [astro-ph.IM].

[28] T. Regimbau *et al.*, "Digging deeper: Observing primordial gravitational waves below the binary black hole produced stochastic background", *Physical Review Letters* **118** no. 15, (2017) 151105, arXiv:1611.08943 [astro-ph.CO].

[29] C. R. Contaldi, "Anisotropies of Gravitational Wave Backgrounds: A Line Of Sight Approach", *Physical Letters* **B771** (2017) 9–12, arXiv:1609.08168 [astro-ph.CO].

[30] G. Cusin, C. Pitrou, and J.-P. Uzan, "Anisotropy of the astrophysical gravitational wave background: Analytic expression of the angular power spectrum and correlation with cosmological observations", *Physical Review* **D96** no. 10, (2017) 103019, arXiv:1704.06184 [astro-ph.CO].

[31] A. C. Jenkins and M. Sakellariadou, "Anisotropies in the stochastic gravitational-wave background: Formalism and the cosmic string case", *Physical Review D* **98** no. 6, (2018) 063509, arXiv:1802.06046 [astro-ph.CO].

[32] S. Mukherjee and J. Silk, "Time-dependence of the astrophysical stochastic gravitational wave background", *Monthly Notices of the Royal Astronomical Society* **491** no. 4, (2020) 4690–4701, arXiv:1912.07657 [gr-qc].

[33] N. J. Cornish and J. D. Romano, "When is a gravitational-wave signal stochastic?", *Physical Review D* **92** no. 4, (2015) 042001, arXiv:1505.08084 [gr-qc].

[34] N. Christensen, "Measuring the stochastic gravitational radiation background with laser interferometric antennas", *Physical Review D* **46** (1992) 5250–5266.

[35] N. J. Cornish, "Mapping the gravitational wave background", *Classical Quantum Gravity* (2001) , arXiv:0105374 [astro-ph.CO].

[36] H. Kudoh and A. Taruya, "Probing anisotropies of gravitational-wave backgrounds with a space-based interferometer: Geometric properties of antenna patterns and their angular power", *Physical Review D* **71** (2005) 024025, arXiv:0411017 [gr-qc].

[37] A. Taruya, "Probing anisotropies of gravitational-wave backgrounds with a space-based interferometer III: Reconstruction of a high-frequency skymap", *Physical Review D* **74** (2006) 104022, arXiv:0607080 [gr-qc].

[38] S. W. Ballmer, "A radiometer for stochastic gravitational waves", *Classical and Quantum Gravity* **23** no. 8, (Apr, 2006) S179–S185.

[39] S. Mitra *et al.*, "Gravitational wave radiometry: Mapping a stochastic gravitational wave background", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **77** no. 4, (2008) 1–24, arXiv:0708.2728.

[40] E. Thrane *et al.*, "Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **80** no. 12, (2009) 1–21, arXiv:0910.0858.

[41] J. Gair *et al.*, "Mapping gravitational-wave backgrounds using methods from CMB analysis: Application to pulsar timing arrays", *Physical Review* **D90** no. 8, (2014) 082001, arXiv:1406.4664 [gr-qc].

[42] J. D. Romano *et al.*, "Phase-coherent mapping of gravitational-wave backgrounds using ground-based laser interferometers", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **92** no. 4, (2015) 1–22, arXiv:1505.07179.

[43] J. D. Romano and N. J. Cornish, "Detection methods for stochastic gravitational-wave backgrounds: A unified treatment", *Living Reviews in Relativity* **20** no. 1, (2017) 1–223, arXiv:1608.06889.

[44] B. P. Abbott *et al.*, "An upper limit on the stochastic gravitational-wave background of cosmological origin", *Nature* **460** no. 7258, (2009) 990–994, arXiv:0910.5772.

[45] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run", *Physical Review Letters* **118** no. 12, (2017) 121101.

[46] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence", *Astrophysical Journal* **851** no. 2, (2017) L35, arXiv:1711.05578 [astro-ph.HE].

[47] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "A search for the isotropic stochastic background using data from Advanced LIGO's second observing run", arXiv:1903.02886.

[48] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "Directional limits on persistent gravitational waves using data from Advanced LIGO's first two observing runs", arXiv:1903.08844.

[49] A. Ain, P. Dalvi, and S. Mitra, "Fast gravitational wave radiometry using data folding", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **92** no. 2, (2015) 1–11, arXiv:1504.01714.

[50] N. Cornish and T. Robson, "Galactic binary science with the new LISA design", *Journal of Physics: Conference Series* **840** no. 1, (2017) 012024, arXiv:1703.09858 [astro-ph.IM].

[51] N. Karnesis, A. Petiteau, and M. Lilley, "A template-free approach for detecting a gravitational wave stochastic background with LISA", arXiv:1906.09027 [astro-ph.IM].

[52] C. Caprini *et al.*, "Reconstructing the spectral shape of a stochastic gravitational wave background with LISA", arXiv:1906.09244.

[53] T. L. Smith and R. R. Caldwell, "LISA for cosmologists: Calculating the signal-to-noise ratio for stochastic and deterministic sources", *Physical Review D* **100** (Nov, 2019) 104055.

[54] A. Taruya and H. Kudoh, "Probing anisotropies of gravitational-wave backgrounds with a space-based interferometer. II. Perturbative reconstruction of a low-frequency skymap", *Physical Review D* **72** (2005) 104015, arXiv:0507114 [gr-gc].

[55] A. Margalit, C. R. Contaldi, and M. Pieroni, "Phase Decoherence of Gravitational Wave Backgrounds", arXiv:2004.01727 [astro-ph.CO].

[56] R. M. Wald, *General Relativity*. University of Chicago Press, 1984.

[57] B. Allen and J. D. Romano, "Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities", *Physical Review D* **59** no. 10, (1999) 102001, arXiv:9710117 [gr-qc].

[58] E. Phinney, "A Practical theorem on gravitational wave backgrounds", arXiv:0108028 [astro-ph].

[59] L. Grishchuk, "Amplification of gravitational waves in an isotropic universe", *Soviet Journal of Experimental and Theoretical Physics* **40** no. 3, (1975) 409.

[60] M. Maggiore, "Gravitational wave experiments and early universe cosmology", *Physical Reptorts* **331** (2000) 283–367, arXiv:gr-qc/9909001 [gr-qc].

[61] C. J. Hogan, "Gravitational radiation from cosmological phase transitions", *Monthly Notices of the Royal Astronomical Society* **218** no. 4, (1986) 629–636. http://dx.doi.org/10.1093/mnras/218.4.629.

[62] R. A. Battye, R. R. Caldwell, and E. P. S. Shellard, "Gravitational waves from cosmic strings", in *Topological defects in cosmology*, pp. 11–31. 1997. arXiv:9706013 [astro-ph].

[63] A. Vilenkin and E. S. Shellard, *Cosmic Strings and Other Topological Defects*. Cambridge University Press, 7, 2000.

[64] N. Bartolo *et al.*, "Primordial Black Hole Dark Matter: LISA Serendipity", *Physical Review Letters* **122** no. 21, (2019) 211301, arXiv:1810.12218 [astro-ph.CO].

[65] N. Bartolo *et al.*, "Gravitational wave anisotropies from primordial black holes", *Journal of Cosmology and Astroparticle Physics* **02** (2020) 028, arXiv:1909.12619 [astro-ph.CO].

[66] A. Celotti, J. C. Miller, and D. W. Sciama, "Astrophysical evidence for the existence of black holes", *Classical Quantum Gravity* **16** (2000) A3, arXiv:9912186 [astro-ph].

[67] T. Callister *et al.*, "Shouts and Murmurs: Combining Individual Gravitational-Wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers", arXiv:2003.12152 [astro-ph.HE].

[68] M. Fishbach, D. E. Holz, and W. M. Farr, "Does the Black Hole Merger Rate Evolve with Redshift?", *Astrophysical Journal Letters* **863** no. 2, (2018) L41, arXiv:1805.10270 [astro-ph.HE].

[69] P. Madau and M. Dickinson, "Cosmic Star Formation History", *Annual Review of Astronomy and Astrophysics* **52** (2014) 415–486, arXiv:1403.0007 [astro-ph.CO].

[70] A. R. Cooray, "Gravitational wave background of neutron star-white dwarf binaries", *Monthly Notices of the Royal Astronomical Society* **354** (2004) 25–30, arXiv:0406467 [astro-ph].

[71] T. Callister *et al.*, "The limits of astrophysics with gravitational-wave backgrounds", *Physical Review X* **6** no. 3, (2016) 031018, arXiv:1604.02513 [gr-qc].

[72] A. Sesana, A. Vecchio, and C. N. Colacino, "The stochastic gravitational-wave background from massive black hole binary systems: implications for observations with pulsar timing arrays", *Monthly Notices of the Royal Astronomical Society* **390** no. 1, (2008) 192–209.

[73] K. Thorne, S. Hawking, and W. Israel, *Three hundred years of Gravitation*. Cambridge University Press, 1987.

[74] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "Gw170817: Implications for the stochastic gravitational-wave background from compact binary coalescences", *Physical Review Letters* **120** (Feb, 2018) 091101.

[75] C. R. Contaldi and J. Magueijo, "Unsqueezing of standing waves due to inflationary domain structure", *Physical Review D* **98** no. 4, (2018) 043523, arXiv:1803.03649 [astro-ph.CO].

[76] M. Kamionkowski, A. Kosowsky, and A. Stebbins, "A Probe of primordial gravity waves and vorticity", *Physical Review Letters* **78** (1997) 2058–2061, arXiv:9609132 [astro-ph].

[77] N. Bartolo *et al.*, "Science with the space-based interferometer LISA. IV: Probing inflation with gravitational waves", *Journal of Cosmology and Astroparticle Physics* **12** (2016) 026, arXiv:1610.06481 [astro-ph.CO].

[78] S. Alexander, A. Marciano, and L. Smolin, "Gravitational origin of the weak interaction's chirality", *Physical Review D* **89** no. 6, (2014) 065017, arXiv:1212.5246 [hep-th].

[79] L. Freidel, D. Minic, and T. Takeuchi, "Quantum gravity, torsion, parity violation and all that", *Physical Review D* **72** (2005) 104002, arXiv:0507253 [hep-th].

[80] S. H. Alexander, "Isogravity: Toward an Electroweak and Gravitational Unification", arXiv:0706.4481 [hep-th].

[81] L. Smolin, "The Plebanski action extended to a unification of gravity and Yang-Mills theory", *Physical Review D* **80** (2009) 124017, arXiv:0712.0977 [hep-th].

[82] A. C. Jenkins *et al.*, "Anisotropies in the astrophysical gravitational-wave background: The impact of black hole distributions", *Physical Review Letters* **122** no. 11, (2019) 111101, arXiv:1810.13435 [astro-ph.CO].

[83] R. Durrer, *The Cosmic Microwave Background*. Cambridge University Press, 2008.

[84] P. J. E. Peebles, *The Large-Scale Structure of the Universe*. Princeton University Press, 2020.

[85] R. Sachs and A. Wolfe, "Perturbations of a cosmological model and angular variations of the microwave background", *Astrophysical Journal* **147** (1967) 73–90.

[86] N. Bartolo *et al.*, "Anisotropies and non-Gaussianity of the Cosmological Gravitational Wave Background", *Physical Review D* **100** no. 12, (2019) 121501, arXiv:1908.00527 [astro-ph.CO].

[87] G. Cusin *et al.*, "First predictions of the angular power spectrum of the astrophysical gravitational wave background", *Physical Review Letters* **120** (2018) 231101, arXiv:1803.03236 [astro-ph.CO].

[88] A. C. Jenkins *et al.*, "Anisotropies in the astrophysical gravitational-wave background: Predictions for the detection of compact binaries by LIGO and Virgo", *Physical Review D* **98** no. 6, (2018) 063501, arXiv:1806.01718 [astro-ph.CO].

[89] M. Geller *et al.*, "Primordial Anisotropies in the Gravitational Wave Background from Cosmological Phase Transitions", *Physical Review Letters* **121** no. 20, (2018) 201303, arXiv:1803.10780 [hep-ph].

[90] D. Bertacca *et al.*, "Projection effects on the observed angular spectrum of the astrophysical stochastic gravitational wave background", *Physical Review D* **101** no. 10, (2020) 103513, arXiv:1909.11627 [astro-ph.CO].

[91] E. E. Flanagan and S. A. Hughes, "The Basics of gravitational wave theory", *New J. Phys.* **7** (2005) 204, arXiv:0501041 [gr-qc].

[92] C. P. Berry and J. R. Gair, "Linearized f(R) Gravity: Gravitational Radiation and Solar System Tests", *Physical Review D* **83** (2011) 104022, arXiv:1104.0819 [gr-qc]. [Erratum: Phys.Rev.D 85, 089906 (2012)].

[93] B. Allen and A. C. Ottewill, "Detection of Anisotropies in the Gravitational-Wave Stochastic Background", *Physical Review D* **56** no. 2, (1996) 545–563.

[94] N. Seto and A. Taruya, "Measuring a parity-violation signature in the early universe via ground-based laser interferometers", *Physical Review Letters* **99** no. 12, (2007) 1–6, arXiv:0707.0535.

[95] N. Seto and A. Taruya, "Polarization analysis of gravitational-wave backgrounds from the correlation signals of ground-based interferometers: Measuring a circular-polarization mode", *Physical Review D - Particles, Fields, Gravitation and Cosmology* **77** no. 10, (2008) 1–27, arXiv:0801.4185.

[96] C. Conneely, A. H. Jaffe, and C. M. Mingarelli, "On the Amplitude and Stokes Parameters of a Stochastic Gravitational-Wave Background", *Monthly Notices of the Royal Astronomical Society* **487** no. 1, (2019) 562–579, arXiv:1808.05920 [astro-ph.CO].

[97] G. Gubitosi and J. Magueijo, "Correlation between opposite-helicity gravitons: Imprints on gravity-wave and microwave backgrounds", *Physical Review D* **95** no. 2, (2017) 023520, arXiv:1610.05702 [gr-qc].

[98] P. A. Rosado, A. Sesana, and J. Gair, "Expected properties of the first gravitational wave signal detected with pulsar timing arrays", *Monthly Notices of the Royal Astronomical Society* **451** no. 3, (2015) 2417–2433, arXiv:1503.04803 [astro-ph.HE].

[99] I. Dvorkin *et al.*, "Exploring stellar evolution with gravitational-wave observations", *Monthly Notices of the Royal Astronomical Society* **479** no. 1, (2018) 121–129, arXiv:1709.09197 [astro-ph.HE].

[100] D. Meacher, E. Thrane, and T. Regimbau, "Statistical properties of astrophysical gravitational-wave backgrounds", *Physical Review D* **89** no. 8, (2014) 084063, arXiv:1402.6231 [astro-ph.CO].

[101] A. C. Jenkins and M. Sakellariadou, "Shot noise in the astrophysical gravitational-wave background", *Physical Review D* **100** no. 6, (2019) 063508, arXiv:1902.07719 [astro-ph.CO].

[102] J. Borrill, "Madcap - the microwave anisotropy dataset computational analysis package", arXiv:9911389 [astro-ph].

[103] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "Searching for a Stochastic Background of Gravitational Waves with LIGO", *Astrophysical Journal* **659** (2007) 918–930.

[104] **LIGO Scientific Collaboration and Virgo Collaboration**, B. Abbott *et al.*, "An Upper Limit on the Stochastic Gravitational-Wave Background of Cosmological Origin", *Nature* **460** (2009) 990, arXiv:0910.5772 [astro-ph.CO].

[105] **LIGO Scientific Collaboration**, J. Abadie *et al.*, "Directional limits on persistent gravitational waves using LIGO S5 science data", *Physical Review Letters* **107** (2011) 271102, arXiv:1109.1809 [astro-ph.CO].

[106] J. R. Bond and R. G. Crittenden, "CMB analysis", *NATO Science Series C* **565** (2001) 241–280, arXiv:0108204 [astro-ph].

[107] J. R. Bond, A. H. Jaffe, and L. Knox, "Estimating the power spectrum of the cosmic microwave background", *Physical Review* **D57** (1998) 2117–2137, arXiv:9708203 [astro-ph].

[108] G. Rocha *et al.*, "Application of XFaster power spectrum and likelihood estimator to Planck", *Monthly Notices of the Royal Astronomical Society* **414** (2011) 823–846, arXiv:0912.4059 [astro-ph.CO].

[109] M. Cerdonio *et al.*, "The Ultracryogenic gravitational wave detector AURIGA", *Classical Quantum Gravity* **14** (1997) 1491–1494.

[110] C. Affeldt *et al.*, "Advanced techniques in GEO 600", *Classical Quantum Gravity* **31** no. 22, (2014) 224002.

[111] B. C. Barish and R. Weiss, "LIGO and the Detection of Gravitational Waves", *Physics Today* **52** no. 10, (1999) 44–50.

[112] R. C. Hilborn, "GW170814: Gravitational wave polarization analysis", arXiv:1802.01193 [gr-qc].

[113] B. P. Abbott *et al.*, "Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy", *Physical Review D* **93** no. 11, (2016) 112004, arXiv:1604.00439 [astro-ph.IM]. [Addendum: Phys.Rev.D 97, 059901 (2018)].

[114] Y. Himemoto and A. Taruya, "Impact of correlated magnetic noise on the detection of stochastic gravitational waves: Estimation based on a simple analytical model", *Physical Review D* **96** (Jul, 2017) 022004. https://link.aps.org/doi/10.1103/PhysRevD.96.022004.

[115] Y. Himemoto and A. Taruya, "Correlated magnetic noise from anisotropic lightning sources and the detection of stochastic gravitational waves", *Physical Review D* **100** no. 8, (2019) 082001, arXiv:1908.10635 [astro-ph.IM].

[116] B. F. Burke and F. Graham-Smith, *An Introduction to Radio Astronomy*. Cambridge University Press, 3 ed., 2009.

[117] **LIGO Scientific Collaboration and Virgo Collaboration**, R. Abbott *et al.*, "Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo", arXiv:1912.11716 [gr-qc].

[118] **LIGO Scientific Collaboration and Virgo Collaboration**, B. P. Abbott *et al.*, "Directional Limits on Persistent Gravitational Waves from Advanced LIGO's First Observing Run", *Physical Review Letters* **118** no. 12, (2017) 121102.

[119] K. M. Gorski *et al.*, "HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere", *Astrophysical Journal* **622** (2005) 759–771, arXiv:0409513 [astro-ph].

[120] A. S. Rahlin, *The First Flight of the SPIDER Balloon-Borne Telescope*. PhD thesis, Princeton University, Mudd Manuscript Library, 7, 2016. http://arks.princeton.edu/ark:/88435/dsp01tm70mx65g.

[121] J. R. Bond, A. H. Jaffe, and L. Knox, "Radical compression of cosmic microwave background data", *Astrophysical Journal* **533** no. 1, (Apr, 2000) 19–37.

[122] **Planck Collaboration**, R. Adam *et al.*, "Planck 2015 results. X. Diffuse component separation: Foreground maps", *Astronomy and Astrophysics* **594** (2016) A10, arXiv:1502.01588 [astro-ph.CO].

[123] A. C. Jenkins, J. D. Romano, and M. Sakellariadou, "Estimating the angular power spectrum of the gravitational-wave background in the presence of shot noise", *Physical Review D* **100** no. 8, (2019) 083501, arXiv:1907.06642 [astro-ph.CO].

[124] P. B. Covas *et al.*, "Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO", *Physical Review D* **97** no. 8, (2018) .

[125] L. P. Grishchuk, "Relic gravitational waves and limits on inflation", *Physical Review D* **48** (Oct, 1993) 3513–3516.

[126] N. Christensen, "Stochastic Gravitational Wave Backgrounds", *Reports on Progress in Physics* **82** (2019) 016903.

[127] E. Hivon *et al.*, "Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets", *Astrophysical Journal* **567** (2002) 2, arXiv:astro-ph/0105302.

[128] M. Armano *et al.*, "*LISA Pathfinder*", arXiv:1903.08924.

[129] N. Tamanini *et al.*, "Science with the space-based interferometer eLISA. III: Probing the expansion of the Universe using gravitational wave standard sirens", *Journal of Cosmology and Astroparticle Physics* **1604** no. 04, (2016) 002.

[130] C. J. Moore, R. H. Cole, and C. P. L. Berry, "Gravitational-wave sensitivity curves", *Classical Quantum Gravity* **32** no. 1, (Dec, 2014) 015014.

[131] T. Robson and N. Cornish, "Impact of galactic foreground characterization on a global analysis for the LISA gravitational wave observatory", *Classical Quantum Gravity* **34** no. 24, (Nov, 2017) 244002.

[132] Z.-C. Chen, F. Huang, and Q.-G. Huang, "Stochastic Gravitational-wave Background from Binary Black Holes and Binary Neutron Stars and Implications for LISA", *Astrophysical Journal* **871** no. 1, (2019) 97, arXiv:1809.10360 [gr-qc].

[133] M. R. Adams and N. J. Cornish, "Discriminating between a stochastic gravitational wave background and instrument noise", *Physical Review D* **82** (Jul, 2010) 022002.

[134] M. R. Adams and N. J. Cornish, "Detecting a stochastic gravitational wave background in the presence of a galactic foreground and instrument noise", *Physical Review D* **89** (Jan, 2014) 022001.

[135] LISA documents, https://www.cosmos.esa.int/web/lisa/lisa-documents.

[136] M. Tinto and S. V. Dhurandhar, "Time-delay interferometry", *Living Review of Relativity* **8** (2005) 4, arXiv:gr-qc/0409034.

[137] M. Muratore, D. Vetrugno, and S. Vitale, "Revisitation of time delay interferometry combinations that suppress laser noise in LISA", arXiv:2001.11221 [astro-ph.IM].

[138] T. A. Prince *et al.*, "The LISA optimal sensitivity", *Physical Review D* **66** (2002) 122002, arXiv:0209039 [gr-qc].

[139] S. Drasco and E. E. Flanagan, "Detection methods for nonGaussian gravitational wave stochastic backgrounds", *Physical Review* **D67** (2003) 082003, arXiv:gr-qc/0210032 [gr-qc].

[140] N. Seto, "Non-Gaussianity analysis of GW background made by short-duration burst signals", *Physical Review* **D80** (2009) 043003, arXiv:0908.0228 [gr-qc].

[141] G. Nelemans, L. Yungelson, and S. F. Portegies Zwart, "The gravitational wave signal from the galactic disk population of binaries containing two compact objects", *Astronomy and Astrophysics* **375** (2001) 890–898, arXiv:0105221 [astro-ph].

[142] V. Korol *et al.*, "Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA", *Monthly Notices of the Royal Astronomical Society* **470** no. 2, (2017) 1894–1910, arXiv:1703.02555 [astro-ph.HE].

[143] P. Auclair *et al.*, "Probing the gravitational wave background from cosmic strings with LISA", *Journal of Cosmology and Astroparticle Physics* **04** (2020) 034, arXiv:1909.00819 [astro-ph.CO].

[144] Characterization and detectability of stochastic gravitational wave background anisotropies with LISA, Cosmology Working Group project, 2020. In preparation.

[145] **LIGO Scientific Collaboration and Virgo and KAGRA Collaboration**, B. Abbott *et al.*, "Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA", *Living Review of Relativity* **21** no. 1, (2018) 3, arXiv:1304.0670 [gr-qc].

[146] **LIGO Scientific Collaboration**, D. Shoemaker, "Gravitational wave astronomy with LIGO and similar detectors in the next decade", arXiv:1904.03187 [gr-qc].

[147] Gravitational Wave International Committee, https://gwic.ligo.org/3Gsubcomm/documents/3G-observatory-science-case.pdf, 2019.

[148] **TianQin Collaboration**, J. Luo *et al.*, "TianQin: a space-borne gravitational wave detector", *Classical Quantum Gravity* **33** no. 3, (2016) 035010, arXiv:1512.02076 [astro-ph.IM].

[149] J. Crowder and N. J. Cornish, "Beyond LISA: Exploring future gravitational wave missions", *Physical Review* **D72** (2005) 083005.

[150] C. Cutler and D. E. Holz, "Ultra-high precision cosmology from gravitational waves", *Physical Review D* **80** (2009) 104009, arXiv:0906.3752 [astro-ph.CO].

[151] M. Branchesi, "Multi-messenger astronomy: gravitational waves, neutrinos, photons, and cosmic rays", *Journal of Physics Conference Series* **718** no. 2, (2016) 022004.

[152] N. Gehrels and D. Spergel, "Wide-Field InfraRed Survey Telescope (WFIRST) Mission and Synergies with LISA and LIGO-Virgo", *Journal of Physics: Conference Series* **610** (May, 2015) 012007.

[153] V. Korol, E. M. Rossi, and E. Barausse, "A multimessenger study of the Milky Way's stellar disc and bulge with LISA, Gaia, and LSST", *Monthly Notices of the Royal Astronomical Society* **483** no. 4, (Dec, 2018) 5518–5533.

[154] G. Canas-Herrera, O. Contigiani, and V. Vardanyan, "Cross-correlation of the astrophysical gravitational-wave background with galaxy clustering", arXiv:1910.08353 [astro-ph.CO].

[155] P. Adshead *et al.*, "Multimessenger Cosmology: correlating CMB and SGWB measurements", arXiv:2004.06619 [astro-ph.CO].

[156] J. F. Beacom, "The Diffuse Supernova Neutrino Background", *Annual Review of Nuclear Particle Science* **60** (2010) 439–462, arXiv:1004.3311 [astro-ph.HE].

[157] J. H. Davis and M. Fairbairn, "A "nu" look at gravitational waves: The black hole birth rate from neutrinos combined with the merger rate from LIGO", *Journal of Cosmology and Astroparticle Physics* **07** (2017) 052, arXiv:1704.05073 [astro-ph.HE].