Topics in Bayesian Population Inference for Gravitational Wave Astronomy

Author:
Riccardo Buscicchio

Supervisor:
Prof. Alberto Vecchio

A thesis submitted to the University of Birmingham for the degree of
Doctor of Philosophy

School of Physics & Astronomy
Institute for Gravitational Wave Astronomy
College of Engineering and Physical Sciences
University of Birmingham

January 1, 2022

Declaration of Authorship

I, Riccardo Buscicchio, declare that this thesis titled, “Topics in Bayesian Population inference for Gravitational Wave Astronomy” and the work presented in it are my own. I confirm that:

- This work was done wholly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date: January 1, 2022

"Do not go gentle into that good night,
Old age should burn and rave at close of day;
Rage, rage against the dying of the light.

Though wise men at their end know dark is right,
Because their words had forked no lighting they
Do not go gentle into that good night."

D. M. Thomas

UNIVERSITY OF BIRMINGHAM

Abstract

College of Engineering and Physical Sciences
School of Physics & Astronomy

Doctor of Philosophy

Topics in Bayesian Population inference for Gravitational Wave Astronomy

by Riccardo Buscicchio

The first detection of a gravitational wave by LIGO and Virgo is a milestone for the study of compact objects in the Universe. Since it took place in 2015, a few tens of detections have been confirmed by the LIGO Virgo Collaboration (LVC). Afterwards, other collaborations have confirmed and extended such catalogues with independent analysis on the same datasets. Through exquisite experimental devices, sophisticated data analysis algorithms, and an accurate interpretational effort, these observations constitutes now an invaluable body of knowledge, and are a key piece of observational evidence for our understanding of the astrophysical population of binary systems. The science case is set to continuously expand as detections increase, and space-based gravitational–wave observatories are going to complement current ones with an otherwise inaccessible window to the “gravitational-wave sky”.

In this thesis I report on my contributions on a few scientific investigations, including: (i) the development of statistical tools for the simultaneous inference on multiple sources and multiple populations of compact binaries; (ii) the development of a framework for parameter estimation on space-based detector observations, focusing on stellar mass binary black–holes and binary white dwarfs systems; (iii) the predictions of yet unobserved phenomena (e.g. gravitational lensing of gravitational waves) or specific signals (e.g. the stochastic foreground of gravitational waves) and their mutual connections;

While developing the tools above I have had the opportunity to provide some insight on: (i) the astrophysical population of binary black hole masses and spins, and their distribution across redshift, therefore providing observational evidence in support of different formation channels; (ii) the future detectability of binary white dwarfs in satellites galaxies of the Milky Way with space-based detectors, with implications on the assembly history of satellite galaxies.

Acknowledgements

My deepest thanks to my supervisor, prof. Alberto Vecchio, for his roaring, fierce, caring enthusiasm. To prof. Will Farr, for his fresh thoughts. To prof. Szabolcs Marka, for the restoring walks along the Hudson river. To Matteo and David for the most precious laughters. To Laura, for the hugs and the kind words. To Marta, for the constant awakening. To Antoine, for pulling the meaning out of things. To Valeriya, for keeping my burning passions alive. To Chris, for the patience. To Davide, for challenging me constantly. To Irene, for being there on the backstage. To Gianmarco, for listening. To Amit, for the rituals.

To the countless coaches along the path: Giancarlo, Walter, Alberto, Zsuzsa, Andreas, Conor, Christopher P.L., Patricia, Geraint, Silvia, Otto, Stas, Antoine P.. To the countless enemies, fellow gladiators in the arena. To the countless friends along the journey: Chiara, Farrah, Janna, Elinore, Sebastian, Sam, Aaron, George, Phil, Sam R., Siyuan, Coen, Eliot, Lucy. And to the patience of the old ones: Francesco T., Nicola D.M., Francesco L., Luigi, Gloria, Giulia, Elisa, Lorenzo, Fabio, Federico, Antonio, Paolo, Nicola L., Iljr, Damian, Salah. And to those who are not there anymore. To my family, myself, and that one person who knows me better than I do.

To the fox and the glacier.

Contents

List of Figures

List of Tables

Physical Constants and Units

Speed of Light	$c = 2.997\,924\,58 \times 10^{8}\,\mathrm{m\,s^{-1}}$ (exact)
Gravitational constant	$G = 6.674\,30 \times 10^{-11}\,\mathrm{m\,kg^{-1}\,s^{-2}}$
Solar Mass	$M_{\odot} = 1.988\,47 \times 10^{30}\,\mathrm{kg}$
Parsec	$pc = 3.085\,67 \times 10^{16}\,\mathrm{m}$

To the meritless silences,
which set fire to my anger.

Contributions overview

In Chapter 1 I provide the foundational material for subsequent Chapters. A basic summary of *gravitational waves* (GWs) emission in the context of *general relativity* (GR) is followed by a brief introduction to the coupling of metric perturbations to ground-based (e.g. *Laser Interferometer Gravitational-Wave Observatory* (LIGO) and *Virgo interferometer* (Virgo)) and space-based (e.g. *Laser Interferometer Space Antenna* (LISA)) detectors.

In Chapter 2 I introduce some relevant concepts for an accurate understanding of BBHs, BNSs, and *double white dwarfs* (DWDs) GW signals. This leads naturally to the description of a hierarchical inference framework using catalogue of events. Following, I summarize the current knowledge on the population of individual sources presented above. All the above is established material in scientific literature. Here I provide an overview for context, insofar it serves as a background for the original work presented in the following Chapters.

In Chapter 3 I describe the identifiability constraint for multimodal likelihoods in two different scenarios: simultaneous parameter estimation of multiple DWDs with LISA; non-parametric hierarchical inference on BBH component masses with LIGO-Virgo Collaboration (LVC) catalogues. I've led both studies, in collaboration with the co-authors of the paper enclosed [1]. The population inference code is entirely developed by me, under the guidance of Dr. Chris Moore. The code for the parameter estimation of LISA sources is part of an ongoing group work carried out by members of the University of Birmingham Institute for Gravitational Wave Astronomy. I've made substantial contributions to its development as part of my PhD research activity.

In Chapter 4 I present a study on the detectability of DWDs through GWs in *Milky Way* (MW) satellite galaxies. I've co-authored the paper [2], lead a significant fraction of the parameter estimation campaign, and contributed to the statistical analysis to assess the discovery potential of new satellites through LISA. The code used for this campaign is the same of Chapter 3.

In Chapter 5 I present a study on the detectability of *stellar-mass binary black Holes* (SmBBHs) through GWs with gravitational wave detectors in space. I've lead the study presented in the paper [3], and the parameter estimation campaign associated. The code used for the campaign is a significantly evolved version of the one used in Chapter 3, with major contributions of Dr. Antoine Klein.

In Chapter 6 I introduce the gravitational lensing of GWs, and enclose two short-author papers addressing the implications of a *stochastic gravitational wave background* (SGWB) non-detection to the expected rate of lensed observations of BBHs (Section 6.1) and BNSs (Section 6.2), respectively. I've led both studies, in collaboration with the co-authors of the papers enclosed [4, 5]. Subsequently, I have performed a similar analysis within the LVC, providing up-to-date constraints using public data from the *LVC third observing run* (O3) [6].

In Chapter 7 we draw conclusions and future prospects of my research activity,

According to the University regulations, a more detailed contribution summary is also prepended at the start of each Chapter.

Chapter 1

Introduction

Chapter 1 is a review Chapter (following closely [8] in Section 1.1), and no original work is presented.

1.1 Linearized gravity

Gravitational waves can be thought of as travelling waves of space-time perturbations. They are a straightforward consequence of the existence a speed limit for the propagation of physical influences, when their geometric interpretation is contextualized in GR. Starting from the metric $g_{\mu\nu}$, it must be a solution of Einstein's fields equation [9]

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu} \tag{1.1}$$

where Ricci tensor $R_{\mu\nu}$ and scalar curvature R involve up to second order derivatives of $g_{\mu\nu}$. The energy-momentum tensor $T_{\mu\nu}$ is associated to matter and radiation distribution of the system. The non-linear nature of such equations makes it challenging to find exact solutions. However, we are interested in a perturbative solution of (1.1), with respect to the one in absence of matter and radiation (i.e. $T_{\mu\nu} = 0$), the flat space-time metric $\eta_{\mu\nu}$

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{1.2}$$

$$|h_{\mu\nu}| \ll 1 \tag{1.3}$$

As a consequence, GR invariance under diffeomorphisms F

$$x^{\mu} \to x'^{\mu} = F^{\mu}(x) \tag{1.4}$$

is restricted to an appropriate set of reference frames, everyone exhibiting *small* $h_{\mu\nu}$. The result is the Poincarè group, with the addition of *small* diffeomorphisms

$$x^{\mu} \to x'^{\mu} = x^{\mu} + \xi^{\mu}(x) \tag{1.5}$$

The resulting classical field theory degree of freedom is $h_{\mu\nu}$, and its gauge invariance is

$$h'_{\mu\nu} = h_{\mu\nu} - (\partial_{\nu}\xi_{\mu} + \partial_{\mu}\xi_{\nu}) \tag{1.6}$$

The differential operator $G_{\mu\nu}$ is linearized through the expansion for $g_{\mu\nu}$, and (1.1) becomes a gauge invariant equations of motion for $h_{\mu\nu}$

$$\Box\bar{h}_{\mu\nu} + \eta_{\mu\nu}\partial^{\rho}\partial^{\sigma}\bar{h}_{\rho\sigma} - \partial^{\rho}\partial_{\nu}\bar{h}_{\mu\rho} - \partial^{\rho}\partial_{\mu}\bar{h}_{\nu\rho} = -\frac{16\pi G}{c^4}T_{\mu\nu} \tag{1.7}$$

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}h^{\alpha}_{\alpha} \tag{1.8}$$

In the so-called *Lorenz gauge*

$$\partial^{\mu}\bar{h}_{\mu\nu} = 0 \tag{1.9}$$

the three terms on the *left-hand side* (LHS) in (1.7) vanish, and $\Box\bar{h}_{\mu\nu}$ satisfies a 4-dimensional tensor wave equation

$$\Box\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu} \tag{1.10}$$

By separating the background metric from the freely propagating waves, we recover the energy-momentum tensor conservation, as a gauge-consistency condition

$$\partial^\mu T_{\mu\nu} = -\frac{c^4}{16\pi G} \partial^\alpha \partial_\alpha \partial^\mu \bar{h}_{\mu\nu} = 0 \tag{1.11}$$

By contrast, in full GR we have a non-conserved energy-momentum tensor, as shown with the introduction of covariant derivatives

$$D^\mu T_{\mu\nu} = \partial^\mu T_{\mu\nu} - \Gamma^\lambda_{\mu\nu} T^\mu_\lambda + \Gamma^\mu_{\mu\lambda} T^\lambda_\nu = 0 \tag{1.12}$$

This is because in non-linear regime matter and radiation exchange energy and momentum with the gravitational field, too.

In summary, the linearization in (1.2) —and subsequent conservation in (1.11)— prescribes that GW sources interact and evolve in a reference spacetime $\eta_{\mu\nu}$ through a well-defined and conserved energy-momentum tensor. It is a known result in literature (see, e.g., [8] or [10] for extensive discussions on the topic) that the background metric does not have to be necessarily flat. Analogue results to the above hold in the presence of additional large-scale low-frequency background gravitational fields, too. They effectively decouple in the linearized theory, and act purely as a background metric which $h_{\mu\nu}$ propagates through. This is an important feature of this framework, since both cosmological expansion and gravitational-wave lensing (relevant mechanisms for the following sections) can be described as such.

Far from the emitting sources, test masses are affected by the metric perturbation $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$ satisfying

$$\Box \bar{h}_{\mu\nu} = 0 \tag{1.13}$$

whose solutions are free waves propagating at light speed. They originate as prescribed by the integral of the inhomogeneous (1.10) over the source volume $\mathcal{V}$

$$\bar{h}_{\mu\nu}(\boldsymbol{x}, t) = \frac{4G}{c^4} \int_{\mathcal{V}} \frac{1}{|\boldsymbol{x} - \boldsymbol{x}'|} T_{\mu\nu} \left(\boldsymbol{x}', t - \frac{|\boldsymbol{x} - \boldsymbol{x}'|}{c} \right) d^3x' \tag{1.14}$$

in that values on the domain boundary $\partial\mathcal{V}$ fix the propagating fluctuations. Since Lorenz gauge in (1.9) is only a partial gauge fixing, it is usually convenient to remove the remaining degrees of freedom. Under a small diffeomorphism (1.5), the tensor $h_{\mu\nu}$ transforms as follows

$$\begin{aligned} \bar{h}_{\mu\nu} \to \bar{h}'_{\mu\nu} &= \bar{h}_{\mu\nu} - (\partial_\nu \xi_\mu + \partial_\mu \xi_\nu - \eta_{\mu\nu} \partial_\rho \xi^\rho) \\ &\equiv \bar{h}_{\mu\nu} - \mathcal{D}_{\mu\nu\rho} \xi^\rho \end{aligned} \tag{1.15}$$

and we can use the four independent arbitrary fields ξ^μ to rearrange the physical content of a GW into the $h_{\mu\nu}$ components, by means of (1.15). The most suitable for our purposes is the *transverse traceless* (TT) gauge. They are implicitly defined by a set of equations for the resulting h^{TT} tensor $h^{TT}_{\mu\nu}$

$$h^{TT}_{0\mu} = 0 \quad \left(h^{TT}\right)_i^i = 0 \quad \partial^j h^{TT}_{ij} = 0 \tag{1.16}$$

However, this set of equation admit solution only in vacuum. If $\Box\bar{h}_{\mu\nu} \neq 0$ (i.e. in the presence of matter or radiation), Lorenz gauge imposes

$$\partial^{\mu}\bar{h}'_{\mu\nu} = 0 \;\;\Rightarrow\;\; \Box\xi_{\mu} = 0 \tag{1.17}$$

$$\Rightarrow\;\; \Box\mathcal{D}_{\mu\nu\rho}\xi^{\rho} = 0 \tag{1.18}$$

So we cannot set to zero any further component of $\bar{h}_{\mu\nu}$ without falling into contradiction

$$\Box h_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \;\;\neq\;\; 0 = \Box\mathcal{D}_{\mu\nu\rho}\xi^{\rho} \tag{1.19}$$

$$h_{\mu\nu} \;\;\neq\;\; \Xi_{\mu\nu} \tag{1.20}$$

Being traceless, the trace-removal is redundant (i.e. $\overline{h}^{TT}_{\mu\nu} = h^{TT}_{\mu\nu}$), and the linearized vacuum equation reads

$$\Box h^{TT}_{ij} = 0 \tag{1.21}$$

and —being h^{TT}_{ij} also symmetric— the most general solution can be cast in the form of *tensor plane–waves* with wavevector $k^{\mu} = \left(\frac{2\pi f}{c}, \boldsymbol{k}\right)$

$$h_{ij}\left(x, k\right) = \sum_{A=+,\times} h_A(k)\exp(ik^{\mu}x_{\mu})(\boldsymbol{e}_A)_{ij} \tag{1.22}$$

with $\boldsymbol{e}_{A=1,2}$ corresponding to a basis for TT 3–tensors e_{ij} orthogonal to the propagation direction. This is again imposed by Lorenz gauge condition $k^i e_{ij}(k) = 0$. In the particular case of a 3–vector $\boldsymbol{k}$ in the z–direction, the two required tensors could be chosen as

$$\boldsymbol{e}_{+} \equiv \boldsymbol{e}_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \qquad \boldsymbol{e}_{\times} \equiv \boldsymbol{e}_2 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{1.23}$$

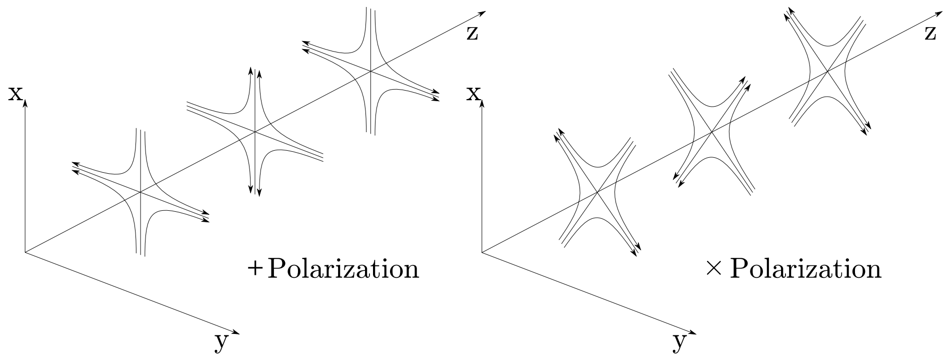

FIGURE 1.1: Polarization tensors *plus* $\boldsymbol{e}_{+}(\hat{z})$ (left) and *cross* $\boldsymbol{e}_{\times}(\hat{z})$ (right), depicted via the *linear maps* $\dot{\boldsymbol{r}} = \boldsymbol{e}_A \boldsymbol{r}$ along the $z-$axis.

Therefore, a GW propagating along a direction $\hat{\boldsymbol{n}}$ can be decomposed in a superposition of plane modes

$$h_{ij}^{TT}(t, \boldsymbol{x}) = \sum_{A=+,\times} e_{ij}^{A}(\hat{\boldsymbol{n}}) \int_{-\infty}^{+\infty} df \tilde{h}_A(f) \exp(-2\pi\imath f(t - \hat{\boldsymbol{n}} \cdot \boldsymbol{x}/c)) \tag{1.24}$$

An important distinction is necessary here. If a detector is most sensitive to GWs with wavelength much bigger than the typical detector size, the retarded time is uniform over it and the term $|\boldsymbol{x}|f/c$ is negligible. Such an approximation holds for ground-based detectors ($|\boldsymbol{x}| \ll 10^5$–$10^7$m), while it's not satisfied for space-based detectors across their whole sensitivity band ($|\boldsymbol{x}| \ll 10^9$–$10^{12}$m). Proposed satellites configuration are expected to be as large as 2.5×10^9m.

When the long-wavelength approximation holds, it is possible to introduce strain scalar timeseries $h_+, h_\times$, defined by the inverse Fourier transforms

$$h_A(t) = \int_{-\infty}^{+\infty} df \tilde{h}_A(f) \exp(-2\pi\imath f t) \tag{1.25}$$

through which (1.22) becomes

$$h_{ij}^{TT}(t) = \sum_{A=+,\times} e_{ij}^A(\hat{\boldsymbol{n}}) h_A(t) \tag{1.26}$$

1.2 Detectors and Observables

The tensor signal in (1.22) is coupled to each observatory through a given experimental setup, resulting in one (or more) observable. A single detector is therefore specified by a tensor $D^{ij}(t)$ whose contraction with $h_{ij}(t)$ gives a time-series $h(t)$:

$$h(t) = D^{ij}(t) h_{ij}(t) \tag{1.27}$$

The functional form and dependencies of the detector tensor are tied to the physical phenomena used to probe the gravitational field. In the following, Section 1.2.1 and 1.2.2, we will characterize the instrument response of ground-based and space-based detectors, respectively. The target in both setups is to construct observables compatible with GR (i.e. whose outcome is invariant upon change of reference system), involving light propagation, mirrors, and leading order perturbations of a background, slowly evolving metric over cosmological timescales and distances).

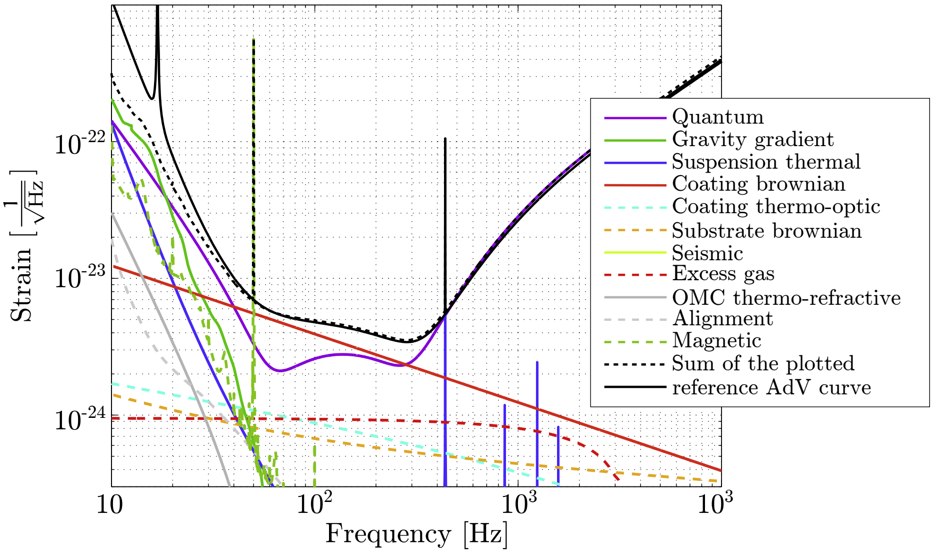

FIGURE 1.2: Advanced Virgo noise budget represented as strain amplitude spectral density. Seismic motions, Brownian noise, and quantum fluctuations of laser source are dominant in the low- ($\sim$ 10Hz), mid- ($\sim$ 100Hz), and high- ($\sim$ 0.5kHz) frequency, respectively. Additional subdominant contributions are described in [11], which this plot is adapted from.

1.2.1 Ground-based interferometers

We will focus here on interferometric measurements, where beams of coherent laser light are set to travel along multiple paths, bounce on test masses suspended from the environment and ultimately recombine through interference at specific locations, where a photodiode samples the readout data. In Figure 1.4 a simplified layout of a Michelson interferometer, sufficient to characterize the detector response, is shown: the laser emitted is split into two coherent beams along two orthogonal paths; then, it circulates along each of LIGO's 4km (3km for Virgo) arms, and is stored through multiple bounces in a Fabry-Perot cavity, providing an effective armlength of 1200km and thus accumulating the effect of GW on light travel times up to the required target sensitivity.

Passive suspensions and actuators (not represented in the layout) isolate the whole setup from external mechanical disturbances within the target frequency band ($\sim$ $10\text{Hz}-1\text{kHz}$), while additional optical elements and coatings are put in place to suppress thermal and opto-mechanical noise sources, as well as to mitigate laser amplitude and phase quantum fluctuations. In Figure 1.2 we illustrate a number of noise sources (as equivalent strain amplitude spectral densities [12]) for Advanced Virgo, as originally presented in [11]. Readily identifiable, three main components dominate current interferometers' noise budget: typically below 10Hz, seismic noise originates from micro-seismic activity effectively causing a displacement of the test-masses by μm; from 10Hz up to 100Hz thermal noise caused by thermo-kinetic excitations of suspensions and mirrors; above 100Hz shot noise arising from the inherent quantum nature of the laser beam photons. For a detailed description of LIGO and Virgo optical layouts, noise contributions and mitigation strategies, refer to [11, 13]. For the interested reader, we plot in Figure 1.3 strain sensitivity plots from the first three observing runs of LIGO and Virgo (the most recent ended in March 2020), alongside projected target sensitivities for the next two observing runs, O4 and O5 (expected to start by the end of 2022 and 2024, respectively) [14].

Going back to the detector coupling with a GW signal, the tensor structure of the metric perturbation is encapsulated in the two polarization tensors, therefore we can conveniently define two antenna pattern functions

$$F_A(\hat{\boldsymbol{n}}) = D^{ij} e^A_{ij}(\hat{\boldsymbol{n}}) \tag{1.28}$$

The detector orientation with respect to the wavevector $-\hat{\boldsymbol{n}}$ drives the functional form of the pattern functions. Solving the geodesic equation for rigid and equal–arm interferometers (oriented with respect to the GW polarization axis as described

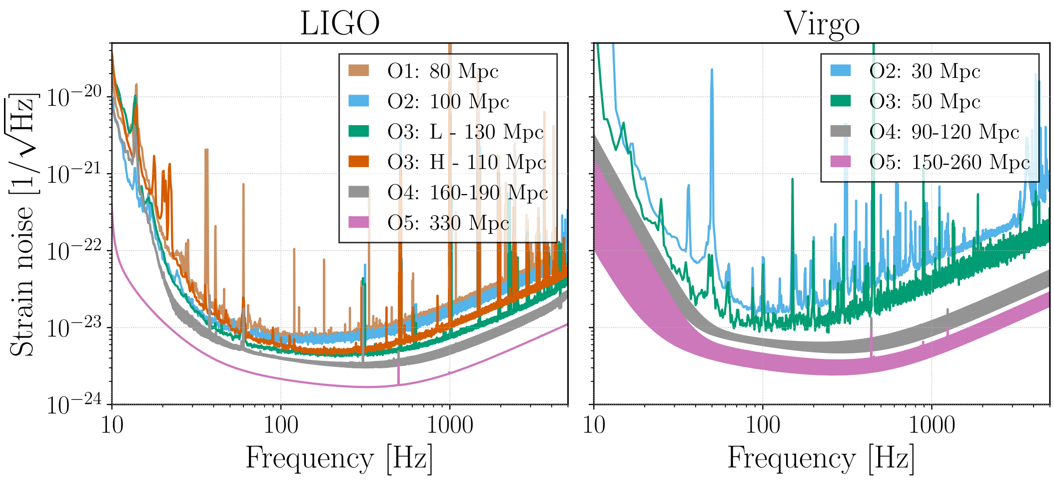

FIGURE 1.3: Past and upcoming observing runs strain noise spectral densities, for LIGO Hanford, Livingston(*left*), and Virgo (*right*) interferometers [14]. Among others, latest improvements before the start of O3 to mirror coatings, test-masses suspension silica fibers, higher input laser power, and the use of squeezed vacuum states brought both interferometers sensitivities down to best-to-date performances for LIGO Livingston (*green, left*), LIGO Hanford (*orange, left*) and Virgo (*green, right*)

in Figure 1.4), one obtains the instantaneous response [15]:

$$F_+(\hat{\boldsymbol{n}}, \psi) = \frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \cos 2\psi - \cos\theta \sin 2\phi \sin 2\psi \tag{1.29}$$

$$F_\times(\hat{\boldsymbol{n}}, \psi) = \frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \sin 2\psi + \cos\theta \sin 2\phi \cos 2\psi \tag{1.30}$$

where we have also included a generic rotation by an angle ψ of the detector axes with respect to the polarization tensors basis. The typical duration of a signal for ground based detectors is much shorter ($10\text{ms} \leqslant T \leqslant 1\text{min}$) than any appreciable time-evolution of the antenna pattern ($1\text{h} \leqslant T$). Therefore the instantaneous response is also a good approximation for the whole duration of an event observation.

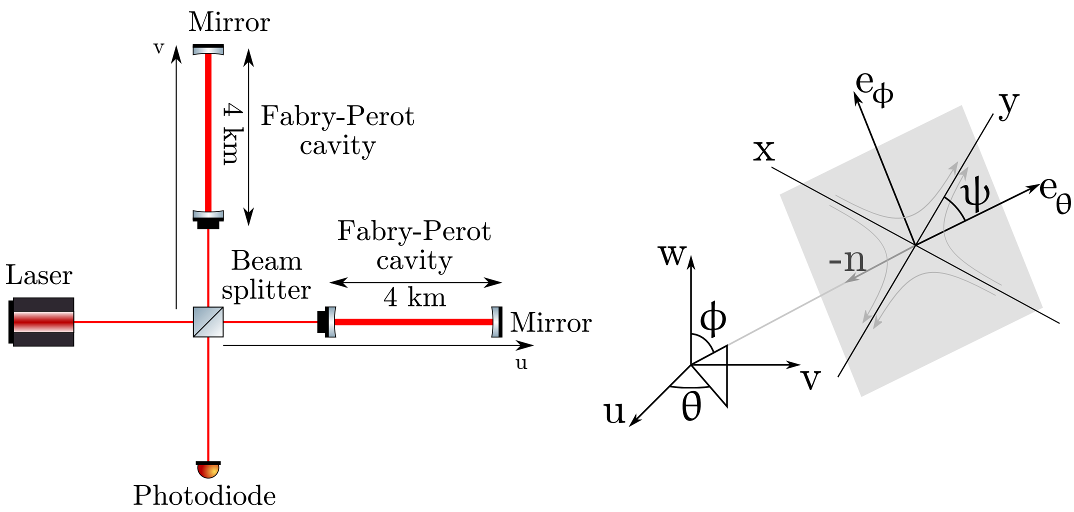

FIGURE 1.4: Simplified layout of ground-based gravitational wave detectors (*left panel*). The basic equal-arm Michelson interferometer is enhanced by the insertion of two Fabry-Perot cavities, one for each arm. The relative orientation of the interferometer arms (aligned with unit vectors u, v) with respect to the incoming waves (with polarization angle ψ) defines the detector istantaneous response through a suitable combination of angles (*right panel*), as in (1.28)

1.2.2 Space-based interferometers

We now turn our attention to the coupling of a GW signal to space-based detectors: current design of the LISA mission [16] (and similarly for other proposed missions, e.g. TianQin [17]) involves three identical spacecrafts, flying in equilateral triangular formation on Keplerian orbits around the Sun.

The constellation center of mass trails the Earth's orbit[1] while the satellites rotate at constant inclination with respect to the orbital plane (see Figure 1.5).

As a consequence of (i) the constellation orbital motion, (ii) the signals durations being comparable to the orbit timescale, and (iii) the satellites not being locked at fixed equal distances between each other further dependences would be introduced with respect to ground-based interferometers in the detector response function, resulting in amplitude, phase, and frequency modulations. Current LISA optical setup is not constituted by

[1] no trailing is involved for the TianQin mission, whose center of mass is centered on the Earth itself

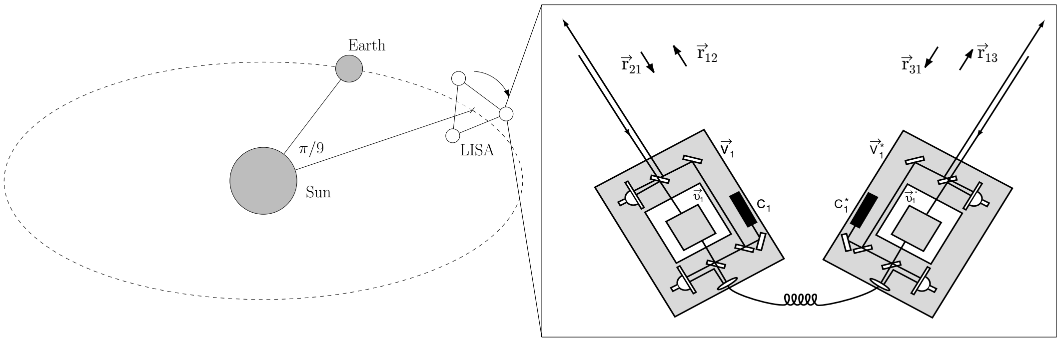

FIGURE 1.5: (*Left*) Satellite motion cartwheeling along the Earth orbit, with a constant offset angle of $\pi/9$ on the orbital plane. (*Right*) Simplified schematics (adapted from [18]) of the companion optical benches aboard each of the three spacecraft satellites. The two local lasers are used to construct intra-spacecraft and inter-spacecraft phase difference observables, the fundamental components to construct the virtual interferometric variables (i.e. TDIs).

phase–locked laser signals (alike LIGO and Virgo) across multiple spacecrafts. Instead, each satellite is comprised of two optical benches, following the motion of free falling test masses, thus providing insulation from environmental disturbances. In the simplified picture in Figure 1.5, each optical bench is equipped with a local laser used to (i) measure phase differences with respect to intra-spacecraft (i.e. between optical benches, through optical fibers connecting them) and inter-spacecraft (i.e. between satellites) incoming beams. Similarly, each local laser beam is sent out to perform analogous measurements on the other benches [18, 19].

The fundamental datum is therefore the phase difference between the signal $\Phi_{ij}(t_j)$ travelling from spacecraft i to j and the reference signal onboard of spacecraft j, evaluated at time t_j. This signal is our probe of changes in path–lengths (either due to the satellites motion $\Delta l_{ij}(t_i)$, or to the transit of a GW $\delta l_{ij}(t_i)$), however it is contaminated by laser phase noises $C(t)$, shot noises $n^s(t)$ and relative acceleration noises $\boldsymbol{n}^a(t)$

$$\begin{aligned}\Phi_{ij}(t_j) &= \omega_l[\delta l_{ij}(t_i) + \Delta l_{ij}(t_i)] \\ &\quad + C_i(t_i) - C_j(t_j) + n^s_{ij}(t_j) - \hat{\boldsymbol{r}}_{ij}(t_i)(\boldsymbol{n}^a_{ij}(t_j) - \boldsymbol{n}^a_{ji}(t_i))\end{aligned} \tag{1.31}$$

where $\hat{\boldsymbol{r}}_{ij}(t_i)$ is a unit vector connecting the space coordinates of emission and reception events on i-th and j-th detector, respectively[2]. We could in principle, following closely the observable definition in the context of ground-based detectors, define three Michelson–like variables by accumulating phases from the four one-way paths combined at each satellite

$$M_X(t) = \Phi_{12}(t_{21}) + \Phi_{21}(t) - \Phi_{13}(t_{31}) - \Phi_{31}(t) \tag{1.32}$$

$$M_Y(t) = \Phi_{23}(t_{32}) + \Phi_{32}(t) - \Phi_{21}(t_{12}) - \Phi_{12}(t) \tag{1.33}$$

$$M_Z(t) = \Phi_{31}(t_{13}) + \Phi_{13}(t) - \Phi_{32}(t_{23}) - \Phi_{23}(t) \tag{1.34}$$

with suitable delays

$$t_{ij} = t - l_{ij}(t_i) \tag{1.35}$$

and inter–satellites path lengths defined by the null–geodesic [20]:

$$l_{ij}(t_i) = \int_i^j \sqrt{g_{\mu\nu}dx^\mu dx^\nu} \tag{1.36}$$

[2] We reserve latin indices for constellation spacecrafts (they do not obey to any tensorial structure), and greek indices for space-time tensors, wherever needed. Alternatively, when convenient for readability, we use boldface fonts to denote space-time tensors.

For equal and constant armlengths, Equations (1.32),(1.33),(1.34) are sufficient to cancel the laser noise dominating the individual Φ_{ij}'s [21, 22].

For the sake of completeness, we report here an expression for the full LISA response in the *rigid adiabatic approximation (RAA)*: the natural timescale separation for this model is the round–trip light travel time between two spacecrafts, or equivalently $c/2\pi L \sim 19 \times 10^{-3}\,\mathrm{mHz}$. Using the Fourier decomposition in (1.24) one can write a generic time and frequency dependent detector tensor [19] defined implicitly by the time-dependent perturbation to the relative path-length between two spacecrafts

$$\frac{\delta l_{ij}(t)}{L} = \sum_n \mathrm{Tr}\left[\mathbf{T}\left(f_n, t, \hat{k}\right)\mathbf{h}_n\right] \tag{1.37}$$

where the sum is performed over the relevant frequencies satisfying the RAA approximation, while the trace Tr is over the space-time indices of the matrix product between the detector tensor $\mathbf{T}$ and the strain signal $\mathbf{h}_n$. Physically, this is equivalent to describe the phase measurements, in presence of a GW, while keeping the satellites at fixed positions and evolving them adiabatically through a rigid constellation orbit. Consistently, the sources' frequency evolution is required to be "slow enough" (i.e. $f/\dot{f} \ll 2\pi L/c$) to not evolve appreciably over the round–trip travel-time of a photon between the two spacecrafts. For all sources we will be focusing on in the following Chapters, i.e. quasi-monochromatic or slowly chirping detached binaries, we will assume this approximation [23]. Any additional modelling assumption will be explicitly stated (e.g. in Eq. (5.8) in Chapter 5).

From Equation (1.37) it is clear that to write explicitly a response, we need an expression for δ_{ij}. Every LISA one–arm phase observable $\Phi_{ij}(t_j)$ in (1.31) is affected by an incoming GW through the pathlength variation δl_{ij}

$$\delta l_{ij}(t) = \frac{1}{2}\mathrm{Tr}\left[\frac{\hat{\boldsymbol{r}}_{ij}(t)\otimes\hat{\boldsymbol{r}}_{ij}(t)}{1-\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{r}}_{ij}(t)}\int_{s_i}^{s_j}\boldsymbol{h}^{TT}(s)ds\right] \tag{1.38}$$

with: s denoting the retarded time introduced in Equation (1.22); s_i denoting the retarded time evaluated on the i-th spacecraft on the corresponding event (either the photon emission or its reception); t denoting the event of reception on spacecraft j; Tr and $\otimes$ denoting the trace and the tensor product over space-time indices. Equation (1.38) is the time domain single-arm response to a GW, the building block for more complex observable responses. By defining the strain Fourier transform

$$\tilde{h}(f) = \int_{-\infty}^{\infty} h(u)e^{-2\pi i f u}du \tag{1.39}$$

one can obtain the equivalent response in frequency domain

$$\frac{\delta l_{ij}(t)}{L_0} = \frac{1}{2}\hat{r}^a(t)\hat{r}^b(t)\int_{-\infty}^{\infty}\tilde{h}_{ab}(f)\mathcal{T}(f,t,\hat{\boldsymbol{k}})e^{2\pi i f(t-\Delta t)}df \tag{1.40}$$

with the transfer function $\mathcal{T}(f,t,\hat{\boldsymbol{k}})$

$$\mathcal{T}(f,t,\hat{\boldsymbol{k}}) = \mathrm{sinc}\left[\pi f L_0\left(1-\hat{\boldsymbol{k}}\cdot\hat{\boldsymbol{r}}(t)\right)\right] \tag{1.41}$$

where Δt is the light-travel time of a photon propagating from the Solar system baricenter to halfway along the LISA arm under consideration.

Going beyond the equal-arm treatment above, more realistic LISA satellites configurations must take into account unequal and time-varying armlengths. More complex combinations of phase measurements, a procedure commonly known as *Time–delay interferometry* (TDI), must be introduced to suppress phase noises whilst keeping the

GW signal unaffected. The construction of such variables (which are known themselves as TDIs) is a crucial modelling step towards realistic full LISA response to an arbitrary GW signal [24, 19, 25]. For constant unequal arm configurations, delays for different arms will be different. Following notation choices in literature [21], we slightly enhance the one introduced so far, defining individual inter-spacecraft distance measurements as y_{ijk} or $y_{ij'k}$: i and k refers to sender and receiver spacecrafts respectively, while j and j' emphasize the propagation arm of the laser beam. We will label the oriented arm from i-th to k-th detector as j if (ijk) is an even permutation of (123) or j' if it's an odd permutation. Accordingly, we define each delay operator with a comma subscript index

$$f_{,j}(t) = D_j f(t) = f\left(t - L_j\right) \tag{1.42}$$

With the above convention, we have sufficient flexibility to construct laser phase noise-cancelling Michelson variables X, Y, Z, as follows:

$$X = y_{231} + y_{13'2,3} + y_{32'1,3'3} + y_{123,2'3'3} - y_{32'1} - y_{123,2'} - y_{231,22'} - y_{13'2,322'} \tag{1.43}$$

$$Y = y_{312} + y_{21'3,1} + y_{13'2,1'1} + y_{231,3'1'1} - y_{13'2} - y_{231,3'} - y_{312,33'} - y_{21'3,133'} \tag{1.44}$$

$$Z = y_{123} + y_{32'1,2} + y_{21'3,2'2} + y_{312,1'2'2} - y_{21'3} - y_{312,1'} - y_{123,11'} - y_{32'1,211'} \tag{1.45}$$

which are known in literature as *first generation TDIs (1.0-g)*. Similar definitions can be made to take into account LISA constellation rotation during its motion around the Sun and armlenghts time-dependence. The interested reader is encouraged to consult exhaustive reviews for further details [26]. Finally, it is worth mentioning that in Chapter 5 we will make use of an additional set of variables, usually referred to as *noise-orthogonal* variables. They are obtained by linearly combining X,Y and Z as follows:

$$A = \frac{1}{\sqrt{2}} (Z - X) \tag{1.46}$$

$$E = \frac{1}{\sqrt{6}} (X - 2Y + Z) \tag{1.47}$$

$$T = \frac{1}{\sqrt{3}} (X + Y + Z) \tag{1.48}$$

in such a way that the acceleration and shot noise cross-correlations [21] cancel, therefore speeding up likelihood evaluations.

Chapter 2

Sources modelling and inference

Chapter 2 is a review Chapter, and no original work is presented.

2.1 Compact binary coalescences

In Chapter 1 we introduced the elementary observable (the "signal") in the context of ground– and space–based interferometry, focusing on the propagation of a gravitational-wave in vacuum and its coupling to interferometric detectors. Here we instead focus on the relevant sources for the studies presented in the following Chapters. Throughout the rest of the thesis we will be focusing on binary systems. It is a standard textbook result (originally derived in [27]) that, to the lowest order in the source velocity v/c and in the radiation zone, the emission is completely described by the quadrupole formula

$$\mathbf{h} = \frac{2G}{c^4 d_L}\ddot{\mathbf{Q}} \tag{2.1}$$

$$\mathbf{Q} = \int_{\mathcal{S}} \mathrm{d}^3\mathbf{r}\rho(\mathbf{r})\left(\mathbf{r}\otimes\mathbf{r} - \frac{1}{3}r^2\right) \tag{2.2}$$

where d_L is the luminosity distance from the source to the point where the strain tensor $\mathbf{h}$ is probed, $\rho(\mathbf{r})$ is the mass density (as a function of position $\mathbf{r}$, with $r = ||\mathbf{r}||$), and the integration is performed over its entire support domain $\mathcal{S}$. With a simple dimensional analysis argument, e.g. approximating $\mathbf{Q}$ in Eq. (2.2) with its typical domain lengthscale R, mass M and motion timescale $1/f$, we obtain for the metric perturbation in Eq. (2.1)

$$h \sim \frac{GMR^2 f^2}{c^4 d_L} \tag{2.3}$$

which for systems governed by Keplerian dynamics can be further simplified to

$$h \sim \frac{(GM)^{5/3} f^{2/3}}{c^4 d_L} \tag{2.4}$$

Targeting a strain amplitude $h \sim 10^{-21}$ at around either 1 mHz or 100 Hz (i.e. the frequencies of highest sensitivity for space- and ground- based detectors, respectively) is possible with close inspiralling astrophysical binary systems. We provide here three examples

$$h \sim 1 \times 10^{-21} \left(\frac{M}{100M_\odot}\right)^{5/3} \left(\frac{650\text{Mpc}}{d_L}\right) \left(\frac{f}{100\text{Hz}}\right)^{2/3} \tag{2.5}$$

$$\sim 1 \times 10^{-21} \left(\frac{M}{5M_\odot}\right)^{5/3} \left(\frac{4\text{Mpc}}{d_L}\right) \left(\frac{f}{100\text{Hz}}\right)^{2/3} \tag{2.6}$$

$$\sim 1 \times 10^{-21} \left(\frac{M}{1M_\odot}\right)^{5/3} \left(\frac{400\text{pc}}{d_L}\right) \left(\frac{f}{5\text{mHz}}\right)^{2/3} \tag{2.7}$$

These are indeed representative cases for tight orbiting galactic DWDs ($M \sim 0.5 - 1M_\odot$) in Chapter 3 and 4, BNSs $M \sim 1 - 3M_\odot$ in Chapter 6 , and stellar mass ($M \sim 1 - 100M_\odot$) black–holes in Chapter 3,5,6 .

2.2 Waveforms, parameter estimation

The GW emission from compacts sources is typically observed through noisy detectors. As a consequence, stringent and robust detection and parameter estimation criteria are required to extract astrophysical information. For a “modelled” search of GW signals —of interest in this thesis— a characterization of detector noise spectral shape and an accurate description of expected waveform based on general relativity are crucial. Therefore, we must go beyond the simplest quadrupolar description of a circular non-spinning binary system emission.

2.2.1 Signal morphology

Depending on the spectral sensitivity of a given detector, we might be able to observe the quasi-monochromatic signals of very early inspirals (e.g. DWD for LISA), the early inspiral of moderately chirping sources (e.g. SmBBH for LISA), or the final stages of the binary undergoing late inspiral, merger, and ringdown into a remnant compact object (*neutron star black hole* (NSBH) or BBH in LIGO and Virgo). Orbit circularization restrict the observability of eccentricity to the early inspiral of sources [29, 30, 31]. On the contrary, merger and ringdown excite higher order modes [32, 33, 34], and exhibit tangible, degeneracies-breaking, waveform signatures from unequal masses, aligned or precessing spins [35, 36]. This motivates the significant effort to construct highly detailed and fast to evaluate waveforms extending in frequency domain from a fraction of mHz (the lowest frequencies accessible to LISA) up to a few hundred Hz (the highest frequencies accessible to LIGO and Virgo).

The waveform models used currently by ground-based detectors data-analysis pipelines are focused mainly on two different approaches: (i) inspiral-merger-ringdown

[37] and effective one body formalism [38, 39]. Within the former, the GW emission is decomposed into three separate phases: the early inspiral, where perturbative post-newtonian (e.g. in powers of v/c) waveforms are employed to high degree of accuracy; the merger, which is characterized by strong non-linear gravity effects between the binary's compact objects and therefore involves an hybrid modelling with analytical techniques and numerical relativity simulations; the ringdown, which describes the excitation (and subsequent relaxation) of the remnant object into its equilibrium state. The three stages are then smoothly connected [40] into a single one for the whole source evolution. On the contrary, the latter formalism describes the two bodies as a single object evolving through an effective gravitational potential. Both approaches, initially developed for the quadrupolar emission of circular aligned–spin binaries, have subsequently been tailored to achieve higher accuracy in many specific scenarios, through the introduction of higher modes and precessing spins [41, 42, 43]. Such improvements come at a non-negligible computational cost, therefore a compromise is required between accuracy and speed, and we will specify our choice in each chapter based on the scientific question addressed therein. For the sake of exposition, a GW-emission "model", described here in frequency domain, is comprised of an amplitude and a phase, parameterized by the source parameters. To the zero-th post-Newtonian order it reads:

$$\tilde{h}_+(f) = Ae^{i\Psi_+(f)}\frac{c}{r}\left(\frac{G\mathcal{M}}{c^3}\right)^{5/6}\frac{1}{f^{7/6}}\left(\frac{1+\cos^2\iota}{2}\right) \tag{2.8}$$

$$\tilde{h}_\times(f) = Ae^{i\Psi_\times(f)}\frac{c}{r}\left(\frac{G\mathcal{M}}{c^3}\right)^{5/6}\frac{1}{f^{7/6}}\cos\iota \tag{2.9}$$

where $A = \frac{1}{\pi^{2/3}}\left(\frac{5}{24}\right)^{1/2}$, r is the source-detector distance, $\mathcal{M}$ is the binary chirp mass parameter

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}} \tag{2.10}$$

ι is the source angular momentum inclination with respect to the line-of-sight, and $\Psi_+(f)$ ($\Psi_\times(f)$) is the phasing associated to the plus (cross) polarization,

$$\Psi_+(f) = 2\pi f\,(t_c + r/c) - \Phi_0 - \frac{\pi}{4} + \frac{3}{4}\left(\frac{G\mathcal{M}}{c^3}8\pi f\right)^{-5/3} \tag{2.11}$$

$$\Psi_\times = \Psi_+ + (\pi/2) \tag{2.12}$$

This is the fundamental building block of any Bayesian inference on modelled signals. In Section 2.2.2 we outline the steps required to perform it, while in Section 2.3 we show how a collection of reconstructed event parameters can be used to infer the population distribution of different sources.

2.2.2 Single event inference

In all studies focusing on individual detections presented in this thesis a well-defined set of "events" is assumed, characterized by a certain definition of significance: a "catalogue". Depending on the problem at hand, we will provide different definitions (either frequentist or Bayesian) of the probability of a signal being astrophysical in origin versus being a noise instrumental artifact, e.g. the *signal-to-noise ratio*s (SNRs) used in Chapters 4 and 5. Any such classifier relies on our knowledge of noise statistical properties across the detector network (e.g. its amplitude spectral density as described in Chapter 1, stationarity, uncorrelatedness, gaussianity) and a parameterized model of GW signals that might be present in the datastream. In a Bayesian context, the inference is carried out by updating a *prior* assumption with information provided by the data, which then gets incorporated into a *posterior*. With the language of probabilities, we want to compute our posterior belief of a source (with parameters θ, e.g. component masses, spins, etc.) generating the observed, noise contaminated data

$\vec{d}$, given some prior belief on the source parameters $\pi(\theta)$. The latter can be modulated into the former using the Bayes theorem

$$p(A \mid B) = \frac{p(B \mid A)p(A)}{p(B)} \tag{2.13}$$

which, in our specific context reads

$$p(\theta \mid \vec{d}) = \frac{\mathcal{L}(\vec{d} \mid \theta)\pi(\theta)}{p(\vec{d})} = \frac{\mathcal{L}(\vec{d} \mid \theta)\pi(\theta)}{\int \mathrm{d}\theta \mathcal{L}(\vec{d} \mid \theta)\pi(\theta)} \tag{2.14}$$

The expression in (2.14) contains the likelihood $\mathcal{L}$ which (thanks to the Gaussianity and uncorrelatedness of the noise across different detectors) can be written as

$$\log \mathcal{L}(\vec{d} \mid \theta) = -\frac{1}{2} \sum_{\alpha=1}^{N_d} |d_\alpha - h_\alpha(\theta)|_\alpha^2 + \text{const.} \tag{2.15}$$

where the sum runs over each detector datastream d_α and signal s_α, which might differ due to the strain tensor coupling as described in (1.27), and the constant is independent on θ and therefore it does not affect the posterior. The norm $|\cdot|_\alpha$ is induced by the weighted scalar product by each detector noise one-sided power spectral density [12], a statistical measure of the average level of noise present in the data

$$\langle a \mid b \rangle_{(\alpha)} \equiv 4\Re \left\{ \int_0^\infty \mathrm{d}f \frac{\tilde{a}(f)\tilde{b}(f)}{S_\alpha(f)} \right\} \tag{2.16}$$

$$|a|_\alpha^2 \equiv \langle a \mid a \rangle_{(\alpha)} \tag{2.17}$$

There is an important distinction to be made here: in the context of ground-based detectors, GWs from bright binary coalescences are expected to be either individually present or absent in any given data chunk (i.e. very rarely two GWs will overlap).

Therefore, leveraging its approximate stationarity, the noise power spectral density is usually estimated from data segments neighbouring the ones containing signal. Through averaging over many such chunks the statistical uncertainty on the power spectrum is significantly reduced. This is not the case for proposed space-based detectors: multiple sources, frequently of multiple classes (e.g. SmBBHs and DWDs), will be overlapping in the datastream in time and frequency domain. Therefore estimating the noise will be more complex, taking into account the presence of other sources while estimating the parameters of a given one. This is the case, for example, of the population of unresolved galactic DWDs that will pile up in the mid-low frequency range ($\sim 0.2 \div 3$ mHz) of LISA sensitivity band, or the problem of joint inference on multiple sources presented in Chapter 3. Studies are ongoing to characterize the statistical properties of confused sources (see e.g. [44] and reference therein) and how they affect individual inference.

The methodology described above is an overview on how one can infer on sources' parameters in a Bayesian framework. This allows us to isolate chunks of data, either in frequency domain (e.g. the DWDs discussed in Chapters 3, 4) or time domain (e.g. BBHs detections from the first two observing runs [45, 46] used in Chapter 3, and BBHs and BNSs catalogues from the first half of the third observing run [47] used in Chapter 6), and identify parameters of resolvable astrophysical signals emerging from the noise. A collection of N such data chunks is denoted in the context of population inference as data $\left\{\vec{d}_i, i = 1, \ldots, N\right\}$.

2.3 Hierarchical inference

Once the data are identified, their probabilistic relation with the astrophysical population must be established. The posterior probability (or, equivalently, the likelihood) of a population given some data is constructed by combining into a single posterior [48, 49, 50, 51]

$$p\left(\vec{\lambda} \mid \left\{\vec{d}_i\right\}\right) = \frac{\pi\left(\vec{\lambda}\right)}{p\left(\left\{\vec{d}_i\right\}\right)} \prod_{i=1}^{N} \frac{\int \mathrm{d}\vec{\theta} p\left(\vec{d}_i \mid \vec{\theta}\right) p_{\mathrm{pop}}\left(\vec{\theta} \mid \vec{\lambda}\right)}{\int \mathrm{d}\vec{\theta} p_{\mathrm{det}}(\vec{\theta}) p_{\mathrm{pop}}\left(\vec{\theta} \mid \vec{\lambda}\right)} \tag{2.18}$$

the prior belief on the population parameter $\pi\left(\vec{\lambda}\right)$, the likelihood of the observed "detections" originating from a set of astrophysical sources, sampled from the population $p_{\mathrm{pop}}(\vec{\theta} \mid \vec{\lambda})$. The latter can be further broken down into the parameters describing the signal, each with likelihood $p\left(\vec{d}_i \mid \vec{\theta}\right)$ of having individual parameters $\left\{\vec{\theta}_i\right\}$. As each population is in principle not entirely detectable for a given detector network, a re-weighting factor $p_{\mathrm{det}}(\vec{\theta})$ accounting for selection effects is inserted to normalize each population according to its overall detectability. Finally an overall normalization factor is given by the evidence of the observed data. The equation presented here is the fundamental component of all population level inference in what follows, with the exception of Chapter 3, where we focus on the "observed" mass distribution of BBHs, therefore we set $p_{\mathrm{det}}(\vec{\theta}) = 1$ for every $\vec{\lambda}$.

2.4 Binary populations

This thesis deals with several observational scenarios of stellar-mass binaries of compact objects in the mass range $0.1 - 100 M_{\odot}$. They comprise a wide variety of binaries (DWDs, BNSs, BBHs) that depending on their evolutionary stage may or may not be observable both from the ground and from space. A wide variety of formation channels for the binary components have been proposed in literature. The astrophysics governing those is beyond the scope of the thesis, however we shortly, non-exhaustively, summarise here the main formation paths and the imprint they leave on the sources' parameters distribution (see, e.g., Sec. 1 in [52] for a more comprehensive list).

BBHs binaries may form in the galactic field and evolve as isolated dynamical systems, primarily governed by the interaction between the individual black-holes progenitors. Such evolutionary pathway requires a number of modelling assumptions to be introduced determining, among others, the specifics of common envelope evolution and the occurrence of mass transfer episodes [53, 54, 55].

Alternatively, young clusters, globular clusters, nuclear star clusters, or accretion disks in active galactic nuclei can host the formation and evolution of the binaries [56, 57, 58, 59, 60, 61, 62, 63, 64].

In all the above scenarios the black holes form from the collapse of stellar objects at the end of their lifetime. A somewhat specific sequence of interactions –e.g. including common envelope phase [65, 66, 67] or chemically homogeneous evolution [68, 69, 70]– may be involved in the formation of compact binaries endproducts. In addition, environmental and dynamical factors like metallicity and supernovae kicks play an essential role regulating the formation and stability of those systems [71, 67, 72, 66, 73]. It is also worth noting that binaries might originate from the dynamical evolution of triple systems [74, 75, 76], the coagulation from repeated hierarchical mergers (see [77] and reference therein for a comprehensive review) or the evolution of an initial population of primordial black-holes [78, 79].

The formation channels listed above leave a potential imprint on the distribution of sources populating the Universe, and subsequently on the events observable by current and future GW detectors [80]. The evidence for a gap in the BBH component mass distribution in the range $40 - 120 M_\odot$ –with current detector being most sensitive to its lowest end–, moderately small mass ratios, and evidence in favour of aligned spins or orbital precession [81] could all be confirmations of one or more of the proposed formation mechanisms [55, 82, 83, 84, 85, 86, 87, 88, 89].

Similarly, on the lowest end of the black holes mass spectrum (i.e. in the range between $\sim 2M_\odot$ and $5M_\odot$) observations of galactic neutron stars and black holes hint at the presence of another gap [90, 91, 92, 93]. However, the detection of gravitational-waves from two BNS mergers [94, 95, 96] and three compact binary mergers with the secondary mass in the same mass range [97, 98] has provided fresh insight on previous assumptions about the existence of such a gap [99, 100, 101, 102].

On the other end of the "low" mass gap, the observed binary neutron stars constitute a powerful testbed for a wide variety of questions: it's been suggested that BNS mergers observed through GW belong to a population similar to the one in our Galaxy [103]; similarly, there is growing evidence that recycled and non-recycled neutron stars belong to different sub-populations [104, 105]; finally, the isolated binary evolution channel, the accretion mechanisms and supernovae explosion models have been challenged together with other plausible formation scenarios [106, 107, 108, 109].

While the statistics for BNSs observed by LIGO and Virgo is low, the prospects for a detection of a population of double white dwarfs by LISA are outstanding. DWDs are by far the most numerous resolvable sources expected in the LISA sensitivity band [110, 111], ranging from thousands to a few tens of thousands.

With such a vast catalogue of sources, studies of the Milky Way morphology [112], its star formation history [113] and gravitational potential will be possible [114]. In addition, binary dynamics (e.g. common envelope and mass-transfer episodes) and internal WD structure [115, 116] will be observationally constrained, furthering synergies with electromagnetic observation [117], population synthesis models and cosmological simulation (see e.g. [118] and reference therein).

It is also worth noting that all the populations introduced here have been thought of as sets containing individually resolvable sources –for a given detector network– whose properties contributes to the population model evidence (see e.g. the mathematical structure of Equation 2.18). In principle, BBHs [119], BNSs [120], and DWDs populations are all characterized by a certain number of sources piling up in a stochastic signal [121]. Even if stochastic –without a deterministically predictable temporal structure– its statistic properties (isotropic, gaussian, unpolarized) are leveraged to detect it in presence of noise [122], thanks to detection algorithms tailored to the specific signal and detector network (see [123] for a comprehensive review of the formalism). Tentative searches for a SGWB have so far provided only upper limits [124, 125]. They constitute an important piece of data, since they affect modelling required in other fields of observational GW astrophysics [126, 127, 128]. This is for example the approach followed in the studies presented in Chapter 6.

Chapter 3

The Label Switching Problem

Contribution summary

This Chapter is a partially edited and reformatted version of [1]:

R. Buscicchio, E. Roebber, J.M. Goldstein, and C.J. Moore. Label switching problem in Bayesian analysis for gravitational wave astronomy, published in Physical Review D, 100(8):084041, (2019).

At the time of publication of [1], the number of detections available from the LVC collaboration where substantially less then today (10 BBHs and 1 BNS, as mentioned in 3.4.1). Since the analysis presented therein is performed using those detections only, I kept the exposition unaltered. For a comprehensive up-to-date list of detections, see [129].

I conceived the study, carried out the original analytical formulation, and developed the code supporting the two numerical simulations presented with the support of the co-authors. In particular, results presented in Sec. 3.4.2 rely on a long-term development of a codebase which all co-authors have contributed significantly to. I've produced all plots shown in this Chapter, and finalized the draft with the help of the co-authors.

3.1 Introduction

It sometimes occurs in Bayesian inference problems that the target distribution depends on several parameters whose ordering is arbitrary. Three examples are immediately apparent from the field of gravitational wave (GW) astronomy alone. Firstly, when describing a compact binary with component masses m_1 and m_2, the likelihood is symmetric under exchange of the labels 1 and 2 (provided all other relevant parameters are suitably adjusted simultaneously). Secondly, when analysing GW time series data containing two or more overlapping sources of the same type, the likelihood is invariant under exchanging all of the parameters of any pair of sources. And thirdly, when analysing the parameters of a population of observed GW events, mixture models can be used to model the population and/or to infer the presence of distinct astrophysical populations. In this case the hyper-likelihood for the population parameters may be invariant under exchanging the parameters of the population components.

Sometimes a simple reparametrisation and restricting the parameter range is enough to remove the degeneracy arising from the arbitrary ordering. In the first case of the binary with two component masses, it is possible to define, say, the total mass $M = m_1 + m_2$ and mass ratio $q = m_2/m_1$ and to sample these over the restricted ranges $M > 0$ and $q \leq 1$. This covers only the restricted portion of the parameter space $m_1 \geq m_2$, thereby removing the symmetry from the likelihood.

The second and third examples are more problematic as they are not restricted to just 2 degrees of freedom. In each case the target distribution has a high degree of symmetry and is invariant under permutations of some number of labels, K. A great deal of literature is devoted to this *label switching problem* in the context of mixture models [130, 131, 132, 133, 134, 135, 136]. The invariance of the target distribution under permutations means that if the posterior has a peak (or mode) at a particular point in parameter space it will necessarily have peaks at all $K!$ points related by symmetry. The extreme scaling of this multimodality poses a serious obstacle to any sampling algorithm in moderate or high dimensional problems.

The most natural way to solve the label switching problem is to impose an *artificial identifiability constraint*. Searching over the restricted region $m_1 \geq m_2$ of the binary component mass space is an example of such a constraint in 2 dimensions. In the K dimensional problem this can be generalised by demanding a certain ordering of the parameters; see, for example, [131, 133, 134, 135]. Restricting to this small region of parameter space avoids all symmetries and removes the excess multimodality. It is also obvious that if one can adequately explore the restricted parameter space satisfying the artificial identifiability constraint then, by symmetry, this is equivalent to exploring the full space.

It remains to implement a suitable artificial identifiability constraint in practical inference problems. This problem can be approached in several ways. For example, when using an Markov chain Monte Carlo (MCMC) to explore the target distribution the proposal can be augmented by composing with a sorting function; i.e. propose a point then reorder the parameters such that the constraint is satisfied [136]. Alternatively, the log-prior distribution can be crudely modified so that it returns $-\infty$ for any point not satisfying the constraint. Either of these will ensure the chain never leaves the desired region of parameter space. However, if not accompanied by tailored proposal distribution, the former method might introduce biases in regions of parameter space where ordering between variables changes (i.e. close to the hypersurfaces where two or more parameters take equal values). Similarly, the latter method would be significantly penalized, in terms of proposal acceptance rate, as the number of parameters to order increases.

While undoubtedly simple, neither of these approaches are completely satisfactory. The former approach requires the user to modify their MCMC proposal distribution and it is difficult to apply when using other stochastic sampling algorithms which don't have a user-accessible proposal distribution (such as nested sampling [137]). For this reasons such an approach is not compatible with the modern approach of treating the sampler, as far as possible, as a *black box* to which the user must only provide a likelihood and a prior. The latter approach is easy to implement for all samplers, but has the significant drawback of being extremely inefficient in high numbers of dimensions. This is because the sampler only proposes useful points satisfying the identifiability constraint a tiny fraction $1/K!$ of the time.

This paper presents a solution to the label switching problem. Our approach follows that of [138] (see, in particular, Eq. A13; however, this equation contains a typographical error as pointed out by [139]). This solution to the label switching problem has been implemented in [140] and has been widely used extensively in the astronomical and cosmological literature [141, 139, 142, 143, 144, 145, 146, 147, 148, 149, 150]. Here we describe the solution in detail, including proofs of certain important properties of the solution. This solution is mathematically elegant, efficient in high dimensions, and can be easily integrated with any sampling algorithm while treating it as a black box.

In Sec. 3.2 the label switching problem is described in detail and the idea behind our proposed solution is illustrated in 2 dimensions. Our solution, for an arbitrary number of dimensions, is presented in Sec. 3.3. In Sec. 3.4 the efficacy of our proposed solution is demonstrated by applying it to the second (Sec. 3.4.1) and third (Sec. 3.4.2) example problems described in the opening paragraph of this section. These example applications are drawn from the field of GW astronomy, but we stress that this method has been more generally applied to inference in astronomy already.

3.2 The Label Switching Problem

We wish to treat problems containing multiple indistinguishable components. Each of the K components is modeled by some parameters $\mathbf{\Lambda}_k \in U$, where the parameter space U is an open set of $\mathbb{R}^n$ and $k \in \{1, 2, \ldots, K\}$. We will choose to distinguish components based on the values of one of these parameters, $x_k \equiv \mathbf{\Lambda}_k^1 \in I$ where I is an set of $\mathbb{R}$. For simplicity, in this section we will consider $x_k \in (0, 1)$ and use a flat prior on each x_k, although these restrictions can be relaxed later.

In the case where there are two components, $K = 2$, our full parameter space is $U \times U$. However we will mainly be interested in the subspace spanned by $\vec{x} = (x_1, x_2)$, which covers the unit square $I \times I$ (in the general K-dimensional case this will be a hypercube which will be denoted $\mathcal{C}$). For the remainder of this section we suppress the other components of $\mathbf{\Lambda}_k$ from our notation for clarity.

Since the two components are indistinguishable, the points (x_1, x_2) and (x_2, x_1) are equivalent; both the likelihood, $\mathcal{L}(\vec{x})$, and prior distributions are symmetric under interchange of the labels 1 and 2 (provided we also remember to relabel all the other components of $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$ simultaneously). As a result, the parameter space is twice as large as it needs to be. Evaluating $\mathcal{L}(\vec{x})$ over the square will typically lead to a distribution with two global maxima (an exception occurs when the true maximum is on the boundary $x_1 = x_2$); secondary peaks, ridges and other structures in the likelihood are also duplicated. In higher dimensions this duplication and multimodality increases in proportion to $K!$ and becomes a serious obstacle to sampling the target distribution.

To avoid sampling multiple identical copies of the same likelihood modes we will enforce the identifiability constraint $x_2 \geq x_1$. This amounts to labelling the component with the smallest x as #1, the component with the next largest x as #2, and so on in higher dimensions. In two dimensions, this restricts the parameter space to the triangle

$x_2 \geq x_1$ (see the off principal diagonal panels in Fig. 3.1). In higher dimensions, the parameter space is restricted to the region $x_K \geq x_{K-1} \geq \cdots \geq x_1$, which is hereafter referred to as the hypertriangle and denoted $\mathcal{T}$.

Samplers naturally propose points in a hypercube. To avoid modifying the sampler itself, we wish to map points in the hypercube to points in the hypertriangle (following the strategy first introduced in the astronomy literature by [151]):

$$\phi : \mathcal{C} \to \mathcal{T} \,. \tag{3.1}$$

Naively, we might try to choose ϕ to be the sorting function. Unfortunately, although it does map into the hypertriangle, it doesn't solve the multimodality problem, since sorting is a many-to-one map. If the sampler proposes a point $\vec{x} = (x_1, x_2)$ in the hypercube and then the user applies the sorting function $\vec{x}' = \mathrm{sort}(\vec{x})$ before evaluating the likelihood $\mathcal{L}(\vec{x}')$, nothing restricts the sampler from searching over the full hypercube. In fact, this procedure is identical to sampling the original hypercube with no sorting.

This is to be distinguished from the procedure of sorting inside the proposal distribution, as referenced in the introduction, which does restrict sampling to the hypertriangle. This is because the newly sorted points are kept by the sampler and used for generating the next set of proposed points. However, this approach violates our desire to treat samplers as black boxes.

To solve the problem we seek a function, ϕ, which is one-to-one. One possibility, in 2 dimensions, is to leave the x_1 coordinate invariant and shift/rescale the x_2 coordinate such that it lies in the desired range:

$$
\begin{aligned}
x_1' &= x_1 \\
x_2' &= x_1' + (1 - x_1')x_2 \,.
\end{aligned}
\tag{3.2}
$$

To see that points are indeed mapped to $\mathcal{T}$ it is sufficient to note that the correct ordering is enforced by adding a positive quantity to x_1' to get a larger value for x_2'. The range $x_2' \in (0, 1)$ is in turn ensured by scaling x_2 with the factor $1 - x_1'$. This map is a indeed one-to-one map from the square to the triangle, thereby removing the problem of multiple modes. However, this map has the unfortunate property that it distorts the prior on the x_2 component, favoring larger values (see the red distribution in Fig. 3.1).

The map in Eq. 3.2 can be "fixed" by revising the x_1 coordinate downwards, before shifting/rescaling x_2:

$$
\begin{aligned}
x_1' &= 1 - \sqrt{1 - x_1} \\
x_2' &= x_1' + (1 - x_1')x_2 \,.
\end{aligned}
\tag{3.3}
$$

The new map in Eq. 3.3 solves the problem in 2 dimensions (see the blue distribution in Fig. 3.1). The sampler can propose points $\vec{x} = (x_1, x_2)$ uniformly in the square. These points are mapped to the triangle $\vec{x}' = \phi(\vec{x})$. Finally, the likelihood is evaluated at the mapped points, $\mathcal{L}(\vec{x}')$. This procedure correctly covers the parameter space just once with the desired flat prior. To prove that the proposed map does indeed maintain the desired flat prior on the individual components one can evaluate the Jacobian of the transformation $\vec{x}' = \phi(\vec{x})$ and show that is constant. This is done in the next section for the K-dimensional case. Because the Jacobian is constant, this transformation will correctly preserve the flat prior that is imposed on the original x_k. [1]

To state the problem formally: to solve the label-switching problem, we seek a bijection (a "one-to-one" and "onto" map) $\phi : \mathcal{C} \to \mathcal{T}$, for an arbitrary number of dimensions, with components $x'_\kappa = \phi_\kappa(x_k)$, such that the determinant $J = \det \mathbf{J}$ of the Jacobian matrix $\mathbf{J}_{\kappa k} \equiv \partial x'_\kappa / \partial x_k$ is a constant.

[1] For an extension of our solution to the wider class of separable priors, see Sec. 3.3.3.

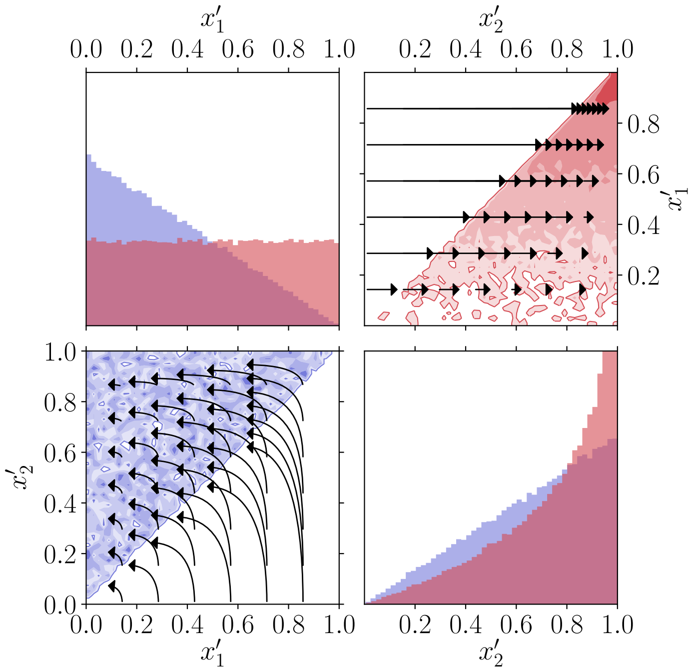

FIGURE 3.1: Two overlaid corner plots, one in the lower-left triangle (blue) and the other in the upper-right triangle (red). Points $\vec{x} = (x_1, x_2)$ were drawn uniformly in the unit square 10^5 times. Histograms of the points $\vec{x}' = \phi(\vec{x})$ are plotted for both the map in Eq. 3.2 (red) and the map in Eq. 3.3 (blue). Both maps correctly move points from the square to the triangle, but only Eq. 3.3 does so while preserving the correct uniform prior. The arrows illustrate how points in the square move under the action of the two maps.

3.3 The hypertriangle map in arbitrary dimensions

Our proposed generalization of the 2-dimensional map in Eq.3.3, $x' = \phi(x)$, is defined recursively as

$$x'_i = x'_{i-1} + (1 - x'_{i-1}) \left[1 - (1 - x_i)^{\frac{1}{K+1-i}}\right] , \tag{3.4}$$

where $i \in 1, \ldots, K$ and $x'_0 = 0$ by definition. This closely resembles Equation (A13) of [138], although here we have corrected a typographical error. Eq. 3.4 can be expressed non-recursively as:

$$x'_i = 1 - \prod_{j=1}^{i} (1 - x_j)^{\frac{1}{K+1-j}} . \tag{3.5}$$

If the inputs are in the correct range $x_j \in (0, 1)$, i.e. $x_j \in \mathcal{C}$, it can be shown that the output falls in $\mathcal{T}$ (the logic as outlined in Sec. 3.2 for Eq. 3.2 still applies). It can also be shown that this map is a bijection by inverting Eq. 3.5.

In the remainder of the section, we will first prove that Eq.3.5 is equivalent to Eq.3.4, and then that the Jacobian of Eq.3.5 is constant.

3.3.1 Equivalent Representations of ϕ

Starting with the recursive version of the map given in Eq.3.4, we rearrange it as follows:

$$\begin{aligned} x'_i &= x'_{i-1} + \left(1 - x'_{i-1}\right) \left[1 - (1 - x_i)^{\frac{1}{K+1-i}}\right] \\ &= \left[1 - (1 - x_i)^{\frac{1}{K+1-i}}\right] - x'_{i-1} \left[1 - (1 - x_i)^{\frac{1}{K+1-i}} - 1\right] \\ &= \left[1 - (1 - x_i)^{\frac{1}{K+1-i}}\right] + (1 - x_i)^{\frac{1}{K+1-i}} x'_{i-1} \\ &= 1 - (1 - x_i)^{\frac{1}{K+1-i}} \left(1 - x'_{i-1}\right) . \end{aligned} \tag{3.6}$$

This procedure can be repeated for the x'_{i-1} term inside the final set of parentheses, and then for x'_{i-2} and so on down to x'_1. This gives the equivalent representation to Eq. 3.5;

$$\begin{aligned} x'_i &= 1 - (1-x_i)^{\frac{1}{K+1-i}} \left[(1-x_{i-1})^{\frac{1}{K+1-(i-1)}} \left(1 - x'_{i-2}\right) \right] \\ &= \dots \\ &= 1 - \prod_{j=1}^{i} (1-x_j)^{\frac{1}{K+1-j}} \quad . \end{aligned} \tag{3.7}$$

3.3.2 The Jacobian of ϕ

As discussed in Sec. 3.2, to maintain the correct prior on the hypertriangle, it is necessary that the map ϕ has a constant Jacobian. To prove that our proposed hypertriangle map has this property, we start with the form of the map in Eq.3.5. The Jacobian matrix $\mathbf{J}_{ij}$ for this specific transformation is lower-triangular because the component x'_i depends only on x_j with $j \leq i$. Its determinant J is therefore equal to the product of the diagonal terms:

$$\begin{aligned} J &= \prod_{i=1}^{K} \mathbf{J}_{ii} \\ &= \prod_{i=1}^{K} \frac{\partial x'_i}{\partial x_i} \\ &= \prod_{i=1}^{K} \frac{1}{K+1-i} (1-x_i)^{\frac{1}{K+1-i}-1} \prod_{j=1}^{i-1} (1-x_j)^{\frac{1}{K+1-j}} \\ &= \frac{1}{K!} \prod_{i=1}^{K} \frac{1}{(1-x_i)} (1-x_i)^{\frac{1}{K+1-i}} \prod_{j=1}^{i-1} (1-x_j)^{\frac{1}{K+1-j}} \\ &= \frac{1}{K!} \prod_{i=1}^{K} \frac{1}{(1-x_i)} \prod_{j=1}^{i} (1-x_j)^{\frac{1}{K+1-j}} \quad , \end{aligned} \tag{3.8}$$

where in the final step a factor has been moved inside of the second product and the upper limit of the product has been changed accordingly. Writing out the products explicitly gives

$$\begin{aligned} J = & \frac{1}{K!} \\ & \times \frac{1}{(1-x_1)} \left[(1-x_1)^{\frac{1}{K+1-1}} \right] \\ & \times \frac{1}{(1-x_2)} \left[(1-x_1)^{\frac{1}{K+1-1}} (1-x_2)^{\frac{1}{K+1-2}} \right] \\ & \times \dots \\ & \times \frac{1}{(1-x_K)} \left[(1-x_1)^{\frac{1}{K+1-1}} \dots (1-x_K)^{\frac{1}{K+1-K}} \right] \end{aligned} \tag{3.9}$$

Careful counting of all the terms reveals that everything cancels and we are left with

$$J = \frac{1}{K!} \,. \tag{3.10}$$

The Jacobian is equal to one over the number of times the original parameter space was covered by the hypercube.

3.3.3 Extension to separable priors

The above derivation considered only flat priors on the x_k. Here we consider the applicability of our hypertriangulation map to separable priors of the form

$$\Pi(x_1, \dots, x_K) = \prod_{k=1}^{K} \pi(x_k) \tag{3.11}$$

In such cases it is first necessary to transform to new coordinates such that the prior is flat before proceeding to apply the hypertriangulation map as before.

In order to find the new coordinates with flat priors, first evaluate the cumulative distribution function

$$F(x) = \int_0^x \pi(s)\mathrm{d}s\,. \tag{3.12}$$

Then define new coordinates $y_k = F(x_k)$ which lie in the range $[0, 1]$. The prior on these new coordinates is now flat and the hypertriangulation map may now be applied to the y_k.

3.3.4 Implementation of ϕ

For concreteness, we provide here a pseudo-code implementation of Eq. 3.5. The input $\mathbf{x} = (x_1, \ldots, x_K)$ (in $\mathcal{C}$) and output $\mathbf{x}' = (x'_1, \ldots, x'_K)$ (in $\mathcal{T}$) are arrays where all values are in the prior range $(0, 1)$. The values of $\mathbf{x}$ may be in any order whilst the values of $\mathbf{x}'$ are, by construction, in ascending order. If a different prior range is needed then the input and output must be shifted and rescaled as appropriate. A full Python implementation (including the shifting and rescaling) is provided at the GitHub repository [152].

3.4 Example GW Applications

In this section we present two applications of our hypertriangle method to two rather different Bayesian inference problems drawn from the field of GW astronomy.

The first example in Sec. 3.4.1 is a Gaussian mixture model; models of this type have been studied extensively in the context of the label switching problem [135, 132, 134, 136, 133, 131, 130].

The second example in Sec. 3.4.2 involves the identification of multiple overlapping signals in time series data. The label switching problem has not often been explicitly

Algorithm Pseudo-code Implementation of Eq. 3.5

function $\phi(\mathbf{x})$:
$K \leftarrow$ **length**$(\mathbf{x})$
$i \leftarrow 1$
for $i \leq K$ **do**
  $j \leftarrow 1$
  $p \leftarrow 1$
  for $j \leq i$ **do**
    $p \leftarrow p(1 - x_j)^{\frac{1}{K+1-j}}$
    $j \leftarrow j + 1$
  end for
  $x'_i \leftarrow 1 - p$
  $i \leftarrow i + 1$
end for
return $x'_1, \ldots, x'_K$

considered in this context. However [153] discuss it when fitting multiple damped sinusoids to time series data.

3.4.1 The Observed Mass Function of LIGO/Virgo Binary Black Holes

LIGO and Virgo [13, 11] are ground-based GW detectors operating in the $(10^1 - 10^4)$ Hz frequency range. The network has been operating since September 2015 and has so far confidently detected 10 binary black hole (BBH) mergers and 1 binary neutron star merger [154]. The third observation run is ongoing and low latency pipelines [155, 156, 157, 158, 159] have produced a number of public alerts associated with event candidates [160, 161]. It is likely that by the end of the current run dozens more detections will be available [162] for further investigation [2].

[2]The statements in this introductory paragraph refer to the time of publication of the manuscript, which Chapter 3 is adapted from. We keep it unaltered, since it motivates the set of detections we considered in the study.

Detailed waveform models for BBH signal calibrated against numerical relativity are now available [163, 164, 165, 42, 166]. These are used in the LALInference Bayesian analysis software package [167] to construct posterior distributions on the parameters of each event. These include both intrinsic (component masses, spins, angular momentum, etc) and extrinsic (sky position, distance, inclination) parameters.

Of these parameters, the best measured and most astrophysically interesting are the individual black hole masses. Parameter reconstruction is crucial from an astrophysical perspective, because it allows both for in-depth studies of individual objects [154, 168, 169, 170, 171, 172] and of populations masses [46, 173, 174].

From a statistical point of view, Bayesian inference on a population of events with imperfect measurements has a well established formalism [50, 49]. A residual freedom remains in the choice of parameterization for the population. Previous studies have used astrophysically motivated functional dependencies [46, 174, 173, 175, 176, 177]. For example, one parameter in such models might be the location of a mass gap in the black hole population [178, 179]. Other studies have used a broader family of somewhat non-parametric models [180, 181, 182, 183, 184].

Within the latter formalism, greater flexibility can be achieved by fitting the observed data with an unknown number of sub-components. No *a priori* physical meaning is necessarily associated with these components, and they are usually sampled from a common hyper-parameter space. The lack of any hierarchy among these components naturally introduces a symmetry under permutations and leads to the label switching problem.

Here we apply our hypertriangle approach to inference on the population of observed BBH component masses, $m_1 \geq m_2$. We model the observed distribution of *source frame* [185] component black hole masses (in solar mass units) as a mixture $p_{\mathrm{pop}}(\log m_1, \log m_2)$

of K bivariate Gaussians;

$$\begin{bmatrix} \log m_1 \\ \log m_2 \end{bmatrix} \sim \sum_{k=1}^{K} w_k \mathcal{N}\left(\begin{bmatrix} \mu_k^{(\log m_1)} \\ \mu_k^{(\log m_2)} \end{bmatrix}, \mathbf{\Sigma}_k \right) . \tag{3.13}$$

Each component has a pair of means, $\mu_k^{(\log m_1)}$ and $\mu_k^{(\log m_2)}$, a symmetric 2×2 covariance matrix, $\mathbf{\Sigma}_k$, and a weight, w_k. The covariance matrix is described by its two eigenvalues, λ_k^1 and λ_k^2, and a rotation angle ϕ_k. Overall, each component is fully described by the parameter vector

$$\mathbf{\Lambda}_k = \left(\mu_k^{(\log m_1)}, \mu_k^{(\log m_2)}, \lambda_k^1, \lambda_k^2, \phi_k, w_k \right) . \tag{3.14}$$

We choose to enforce the artificial identifiability constraint $\mu_{k+1}^{(\log m_1)} \geq \mu_k^{(\log m_1)}$. This is done by applying our map ϕ from Eq. 3.5 to the vector of components $\mu_k^{(m_1)}$ with $k = 1, 2, \ldots, K$. We can sample on the modified parameter space covered by

$$\mathbf{\Lambda}_k = \left(\chi_k, \mu_k^{(\log m_2)}, \lambda_k^1, \lambda_k^2, \phi_k, w_k \right) . \tag{3.15}$$

where $\mu_k^{(m_1)} = \phi(\chi_k)$. In the language of Sec. 3.2, sampling on the parameter space in Eq. 3.14 covers $\mathcal{C}$ (with multimodality) while sampling on Eq. 3.15 covers $\mathcal{T}$.

The priors are taken to be flat on all of the components in Eqs. 3.14 and 3.15, except for the λ_k^1, λ_k^2 which we take log-uniformly distributed within their ranges.

The ranges for χ_k, $\mu_k^{(\log m_1)}$, $\mu_k^{(\log m_2)}$ are $(0, 2)$, with the additional constraint of $\mu_k^{(\log m_1)} > \mu_k^{(\log m_2)}$. The range on the angle ϕ_k is $(0, \pi/2)$ and the ranges on λ_k^1 and λ_k^2 are $(0.01, 4)$. Finally, the weights w_k were sampled in the range $(0, 1)$ and then normalized such that $\sum_k w_k = 1$. Prior choices on mixture parameters is summarized in Table 3.1.

$\mu_k^{(\log m_1)}$	$\mu_k^{(\log m_1)}$	λ_k^1	λ_k^2	ϕ_k [rad]	w_k
$(0, 2)$	$(0, 2)$	$(0.01, 4)$	$(0.01, 4)$	$(0, \pi/2)$	$(0, 1)$

TABLE 3.1: Prior ranges for the BBH observed population parameters.

We adopt a fully Bayesian hierarchical approach. At the lowest level there are the short segments of time series data $\{d\}$ surrounding each of the N_{obs} events. Each event is described by some parameters θ (e.g. masses, spins, etc). The likelihood that we wish to sample from is the probability of all the observed data given a certain value of the population parameters $\mathbf{\Lambda} = \{\mathbf{\Lambda}_k | k = 1, 2, \ldots, K\}$:

$$p\left(\{d\} \mid \mathbf{\Lambda}\right) = \prod_{i=1}^{N_{\text{obs}}} \frac{\int \mathrm{d}\theta \, p(d \mid \theta) \, p_{\text{pop}}(\theta \mid \mathbf{\Lambda})}{\int \mathrm{d}\theta \, p_{\text{pop}}(\theta \mid \mathbf{\Lambda})} \,. \tag{3.16}$$

Using Bayes theorem, the above likelihood can be turned into a posterior on the population parameters $\mathbf{\Lambda}$. This in turn can be expressed in terms of the N_i posterior samples on θ from each individual event [50]:

$$p\left(\mathbf{\Lambda} \mid \{d\}\right) = \varpi(\mathbf{\Lambda}) \prod_{i=1}^{N_{\text{obs}}} \frac{\frac{1}{N_i} \sum_{j=1}^{N_i} \frac{p_{\text{pop}}(\theta_i^j \mid \mathbf{\Lambda})}{\pi(\theta_i^j)}}{\int \mathrm{d}\theta \, p_{\text{pop}}(\theta \mid \mathbf{\Lambda})}, \tag{3.17}$$

where the posterior samples for each event are denoted θ_i^j (i labels the event and j labels the sample in the posterior chain), and $\varpi(\mathbf{\Lambda})$ and $\pi(\theta)$ respectively denote the priors on the population and individual event parameters. We will consider only the component masses as event parameters, $\theta = (m_1, m_2)$. Note that the normalization integral in the denominator of Eq. 3.17 is evaluated over the constrained prior range

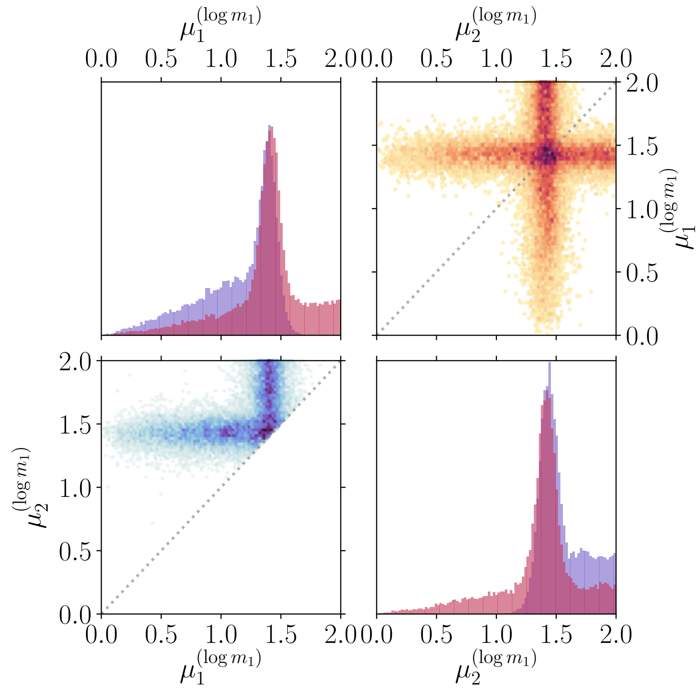

FIGURE 3.2: Two overlaid corner plots, one in the lower-left triangle (blue) and the other in the upper-right triangle (red). The red posterior is obtained by sampling in the parameter space of Eq. 3.14; this space covers the hypercube $\mathcal{C}$ and has a multimodal posterior. The blue posterior is obtained by sampling in the parameter space of Eq. 3.15 and then transforming to $\mu_k^{(m1)} = \phi(\chi_k)$; this only covers the hypertriangle $\mathcal{T}$ and has a single posterior mode. The grey dotted line marks equality between the two components.

$\log m_1 > \log m_2$. We use the publicly available posterior samples [186] for the 10 BBH events described in [154].

As our focus here is on the label switching problem, and its solution using the hypertriangle map, for simplicity we do not consider selection effects [189, 190]. Rather, we model the distribution of *observed* black hole masses. We defer a full treatment, including selection effects, to future work.

We model the observed mass distribution using $K = 1, \ldots, 4$ Gaussian components. We sample the distribution in Eq. 3.17 using the nested sampling algorithm [137] as implemented in CPNest [191]. The primary output of the algorithm is the model

K	$\log Z_{\mathcal{T}}$	$\log Z_{\mathcal{C}}$
1	-74.76 ± 0.09	
2	-78.37 ± 0.05	-78.30 ± 0.05
3	-81.82 ± 0.09	-81.66 ± 0.08
4	-84.58 ± 0.06	-84.2 ± 0.1

TABLE 3.2: Log-evidences for mixtures with different number of components K. The variables $Z_{\mathcal{T}}$ and $Z_{\mathcal{C}}$ denote the evidences obtained by sampling on the hypertriangle parameter space in Eq. 3.14 and the (multimodal) hypercube parameter space in Eq. 3.15 respectively. Mathematically we have already proved that these parameter spaces are equivalent and therefore the two evidences are equal; these two columns serve to demonstrate this numerically. For the $K = 1$ component case there is no distinction between the two parameter spaces (the map ϕ reduces to the identity in this case). The errors on the CPNest evidence integrals were estimated by a combination of the internal CPNest error estimate (as described in [137]) and examination of the spread of results from multiple runs. The $Z_{\mathcal{T}}$ and $Z_{\mathcal{C}}$ evidences are broadly consistent; however for large K there is some tension. We think this is due to CPNest systematically underestimating the $Z_{\mathcal{C}}$ evidence which comes from a high dimensional and highly multimodal posterior. Alternative nested samplers [151, 138, 187, 188] have been shown to reliably estimate evidences for problems of similar complexity.

evidence, which we use to determine which K is favored; we find that the data favors a description using a 1-component Gaussian mixture. Additionally, the algorithm produces samples from the posterior in Eq. 3.17. The log-evidences for different K are presented in Table 3.2, while the posterior samples for $K = 2$ ($K = 3$) on the $\vec{\mu}^{(\log m_1)}$ parameters are shown in Fig. 3.2 (Fig. 3.4). Median a posteriori values of $p(\log m_1, \log m_2)$ are shown in Fig. 3.5 for one and two mixture components. The full posterior chain on all of the parameters is provided at [192].

Because this is a relatively low-dimensional problem (we consider $K \leq 4$) the analysis can be performed both with and without the hypertriangle map. If the map is not used then the posterior has $K!$ degenerate modes. If the map is used then there is just a single mode and, importantly, no information is lost. The elimination of the excess multimodality is shown for two dimensions in Fig.3.2. More impressive

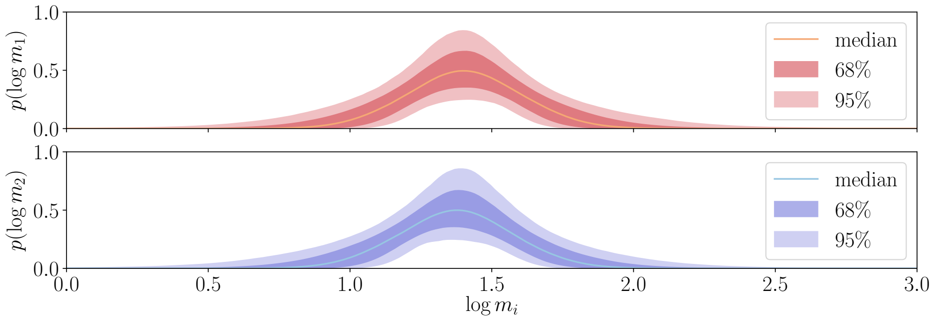

FIGURE 3.3: . The recovered marginal mass distributions on the *observed* component masses in source frame. The red (blue) curves show the marginal distribution on $\log m_1$ ($\log m_2$). All masses are measured in solar mass units. The central line in each case corresponds to the a posteriori median values of $p(\log m_i \mid \mathbf{\Lambda})$. The shaded regions denote the 1σ and 2σ confidence regions associated. These posteriors are obtained by marginalizing over K. The two dimensional mass distribution is shown in Fig. 3.5.

demonstrations of the elimination of the excess multimodality are possible in higher numbers of dimensions; a plot in 3 dimensions for $K = 3$ component mixture is shown in Fig. 3.4. The preservation of information is demonstrated by the fact that the evidence in unchanged. This fact can be shown analytically and is a consequence of the Jacobian for our transformation in Eq. 3.10; it is also demonstrated numerically for this specific problem in Table 3.2.

We can now use the posteriors on $\mathbf{\Lambda}$ to plot the observed black hole mass distribution. This can be done using the posterior on the $\mathbf{\Lambda}$ from either of Eqs. 3.14 or 3.15 with identical results. Although, the single mode posterior from Eq. 3.15 is naturally easier to sample from. The marginalised mass distributions on m_1 and m_2 are plotted in Fig. 3.3. As shown by the evidences in Table 3.2 a one component mass distribution is favoured. We stress again that we have not included selection effects; including these is expected to suppress the high mass tail (this is because high mass BBHs can be seen out to greater distances than lower mass systems) and therefore our results are not

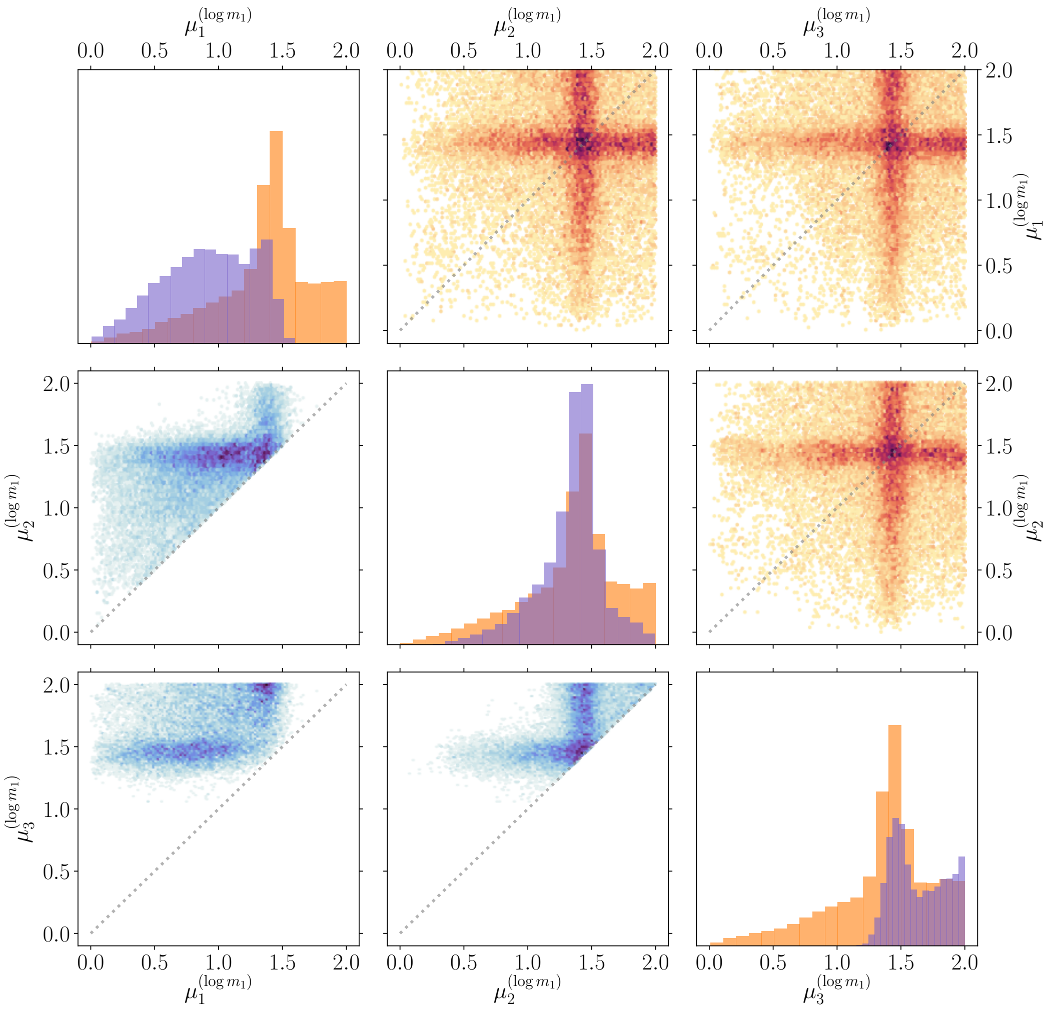

FIGURE 3.4: Posteriors on mixture parameters $\left(\mu_1^{(\log m_i)}, \mu_2^{(\log m_i)}, \mu_3^{(\log m_i)}\right)$, assuming $K = 3$. The bottom-left and top-right triangles show the corner plots for the analysis performed the with and without applying the hypertriangulation map, respectively hypertriangulation map. Along the diagonal the 1-dimensional samples histograms are overlaid with both configurations. Dashed gray lines denote equal mixture components primary mass means.

incompatible with the presence of a mass gap.

The hypertriangle map has demonstrated its utility. It eliminated the excess multimodality in the description of the observed BBH mass distribution. This renders the target posterior easier to sample. There is no loss in information incurred by

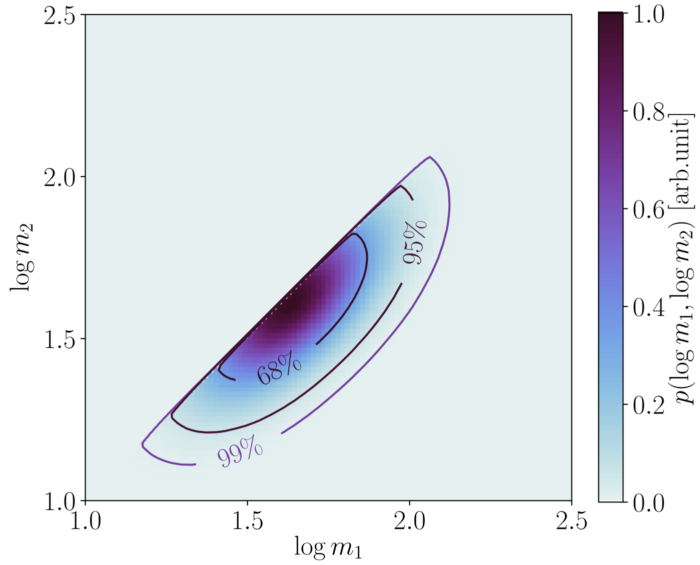

FIGURE 3.5: $p(\log m_1, \log m_2)$ a posteriori median values. Posterior samples from the one and two component mixtures are combined, according to their evidence in Table 3.2, into a single set of posterior samples. Lines denote the 1σ, 2σ and 3σ contour levels, respectively.

sampling this remapped parameter space compared to sampling the full original space.

3.4.2 Overlapping Galactic White-Dwarf Binaries in LISA

LISA [16] is a planned space-based mission which will observe GWs in the (0.1–100) mHz frequency range. The LISA band is source-rich, with many signals overlapping in both time and frequency. In particular, galactic white dwarf binaries (GBs) [111] are so numerous at low frequencies that they form a confusion noise foreground for LISA. Several GBs have already been identified electromagnetically and will serve as verification sources for LISA [193].

The label-switching problem arises in the analysis of multiple sources, since the parameters of any pair of sources are interchangeable. In this section we will show

how the application of the hypertriangle map allows for efficient Bayesian recovery of multiple GB signals without ambiguity arising from label switching.

The GWs emitted by a distant source are observed in the solar system as plane waves. There are two GW polarization components denoted $+$ and $\times$. Under the assumption that each source is monochromatic, these components are given by

$$\begin{aligned} h_+(t;\mathbf{\Lambda}) &= A\,(1+\cos^2\iota)\,\cos(2\pi f t-\Phi)\,, \\ h_\times(t;\mathbf{\Lambda}) &= -2A\,\cos\iota\,\sin(2\pi f t-\Phi)\,, \end{aligned} \tag{3.18}$$

where f is the GW frequency, ι is the inclination angle between the binary's orbital angular momentum and the line of sight, and Φ is a phase offset.

The LISA detector response additionally depends on the ecliptic longitude and latitude $\{\lambda,\beta\}$ of the source and a polarization angle ψ. The GW amplitude A can be further expressed in terms of physical quantities of the GB system (e.g. the component masses and the luminosity distance); however, these quantities are highly degenerate and are therefore not considered.

Each of the K sources is described by seven parameters:

$$\mathbf{\Lambda}_k = \{\log_{10}A_k, f_k, \lambda_k, \sin\beta_k, \cos\iota_k, \psi_k, \Phi_k\}\,. \tag{3.19}$$

We use flat priors on all parameters with ranges given in Table 3.4. The log-likelihood is given by

$$\log\mathcal{L}(\mathbf{\Lambda}_k) \propto -\frac{1}{2}\sum_{\alpha}\left|s_\alpha - \sum_{k=1}^{K}h_\alpha(\mathbf{\Lambda}_k)\right|^2_{(\alpha)}\,, \tag{3.20}$$

where k labels the various GBs, and where s_α denotes two approximately independent

$f - f_\star$	$\log_{10} A$	ι [rad]	λ [rad]	β [rad]	ψ [rad]	ϕ [rad]	ρ
0	-22.15	0.246	-0.096	0.218	1.640	1.795	10
2/yr	-22.13	0.403	0.091	0.294	1.066	4.249	10
4/yr	-22.13	0.376	-0.055	0.359	0.794	4.760	10
6/yr	-22.15	0.284	0.031	0.248	1.127	2.078	10
8/yr	-22.13	0.390	0.006	0.223	0.775	4.537	10
10/yr	-22.12	0.428	0.091	0.296	1.088	5.765	10

TABLE 3.3: The parameters of the six injected GBs, with $f_\star = 1\,\mathrm{mHz}$. The amplitudes were chosen such that the signal-to-noise ratio is 10 in each case.

LISA output channels, with $\alpha \in \{\mathrm{A}, \mathrm{E}\}$ (see, for example, [194]). The model h_α is the LISA response to sinusoidal signals of the form in Eq. 3.18. The line brackets indicate a norm with respect to the usual signal inner product

$$\langle a|b\rangle_{(\alpha)} = 4\Re\left\{\int_0^\infty \mathrm{d}f\, \frac{\tilde{a}(f)\tilde{b}(f)}{S_\alpha(f)}\right\}, \tag{3.21}$$

where $\tilde{a}(f)$ is the Fourier transform of $a(t)$. Each output channel is assumed to contain additive stationary Gaussian noise with a one-sided power spectral density $S_\alpha(f)$.

We simulate one year of mock LISA noise using LISA code [195]. For simplicity, we estimate the power spectral densities from these signal-free noise realizations using the Welch periodogram [196, 197].

We inject $K = 6$ sources, each with a signal-to-noise ratio $\rho_k = 10$ defined with respect to the inner product in Eq. 3.21. The six sources were chosen to have regularly spaced frequencies; other source parameters were chosen randomly and are given in Table 3.3. For simplicity, we perform a noise-free analysis.

The simulated data has a cadence of 5 s and a total duration of 1 yr, resulting in arrays of length 6.3×10^6. This data was heterodyned, filtered, and downsampled to isolate a narrow range of frequencies $f \in (f_\star - 1/\mathrm{yr}, f_\star + 11/\mathrm{yr})$, where $f_\star = 1$ mHz. For

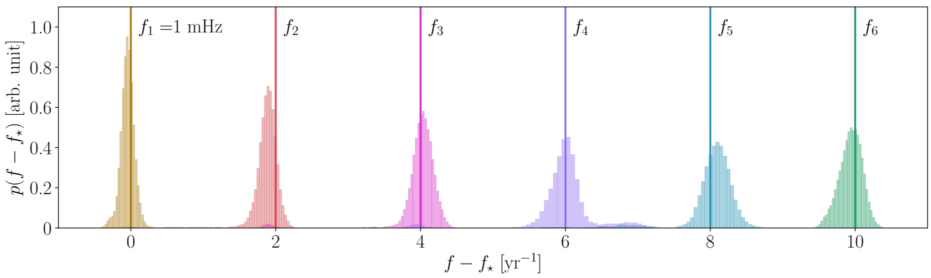

FIGURE 3.6: The 1-D marginalized posterior distributions on the physical frequencies f_k of the six GBs. Vertical lines mark the injected frequencies. As we used a zero noise injection, we expect the posteriors to be peaked at the injected values. We observe that some neighboring sources (notably 4 and 5) show some correlation . This effect is not an artifact of the hypertriangle map or the sampling. Rather, it is a genuine feature of the posterior caused by the non-zero overlap between sources closely spaced in frequency. A systematic study of such cross-contamination, and its dependence upon the sources' parameters is subject of ongoing investigation.

a one-year observation period, the expected frequency resolution of LISA is $\sim 1/1\,\mathrm{yr}$, so this frequency range covers 12 bins.

We assume the number of sources K is already known by other means; we do not address the problem of searching for an unknown number of sources (see, for example, [198, 199, 200]).

This is a $6 \times 7 = 42$-dimensional Bayesian inference problem. The likelihood in Eq. (3.20) is invariant under permutations of the index k (i.e. relabelling the GBs numbered 1 to 6). Naively sampling this distribution in the specified prior ranges will return a posterior distribution with (at least) $6! = 720$ peaks. To remove this problem we enforce the artificial identifiability constraint $f_{k+1} \geq f_k$ by sampling on the parameters

$$\boldsymbol{\Lambda}_k = \{\log_{10} A_k, \chi_k, \cos \iota_k, \lambda_k, \sin \beta_k, \psi_k, \phi_k\} \,. \tag{3.22}$$

Here $f_k = \phi(\chi_k)$ (see Eq. 3.5), and the prior on χ_k is the same as the prior on f_k. In

$(f - f_\star)[\text{yr}^{-1}]$	$\log_{10} A$	λ [rad]	$\sin\beta$	$\cos\iota$	ψ [rad]	Φ [rad]
$(-1, 11)$	$(-23.0, -21.8)$	$(0, 1)$	$(-0.75, 0.75)$	$(0, 1)$	$(0, \pi)$	$(0, 2\pi)$

TABLE 3.4: Prior ranges on the GB parameters. The frequency prior spans twelve bins around $f_\star = 1\,\text{mHz}$. Priors are taken to be uniform over the respective ranges.

Table 3.4 we list the prior boundaries for each of the seven parameters, which are taken to be uniformly distributed over the respective ranges.

The sampler explores the space of $\chi_k \in \mathcal{C}$ which is mapped to the physical frequencies $f_k \in \mathcal{T}$. The resultant distribution has a single global maximum and is therefore relatively easy to sample from (albeit in 42 dimensions).

We use CPNest [191] to sample the distribution and correctly recover all sources. We note that without applying our hypertriangle map, it would be excessively difficult to sample from this 720-fold degenerate distribution.

In Fig. 3.6 we focus on the the 1D marginalized posteriors on the physical frequencies f_k.

The full posterior parameters are publicly available [192]. A selection of these parameters are plotted in Fig. 3.7, for the subset of parameters corresponding to the third source tabulated in Table 3.3.

3.5 Discussion

We discussed a general solution to the label switching problem which allows the sampler to be treated as a black box, and is therefore widely applicable. To enforce the identifiability constraint, we map the sampled points from a hypercube with the desired prior to a hypertriangle, taking care to preserve the prior. We have successfully used this for two real-world problems from gravitational wave astrophysics. The hypertriangle transformation has the potential to greatly simplify a wide class of highly-degenerate Bayesian inference problems, with no loss of information.

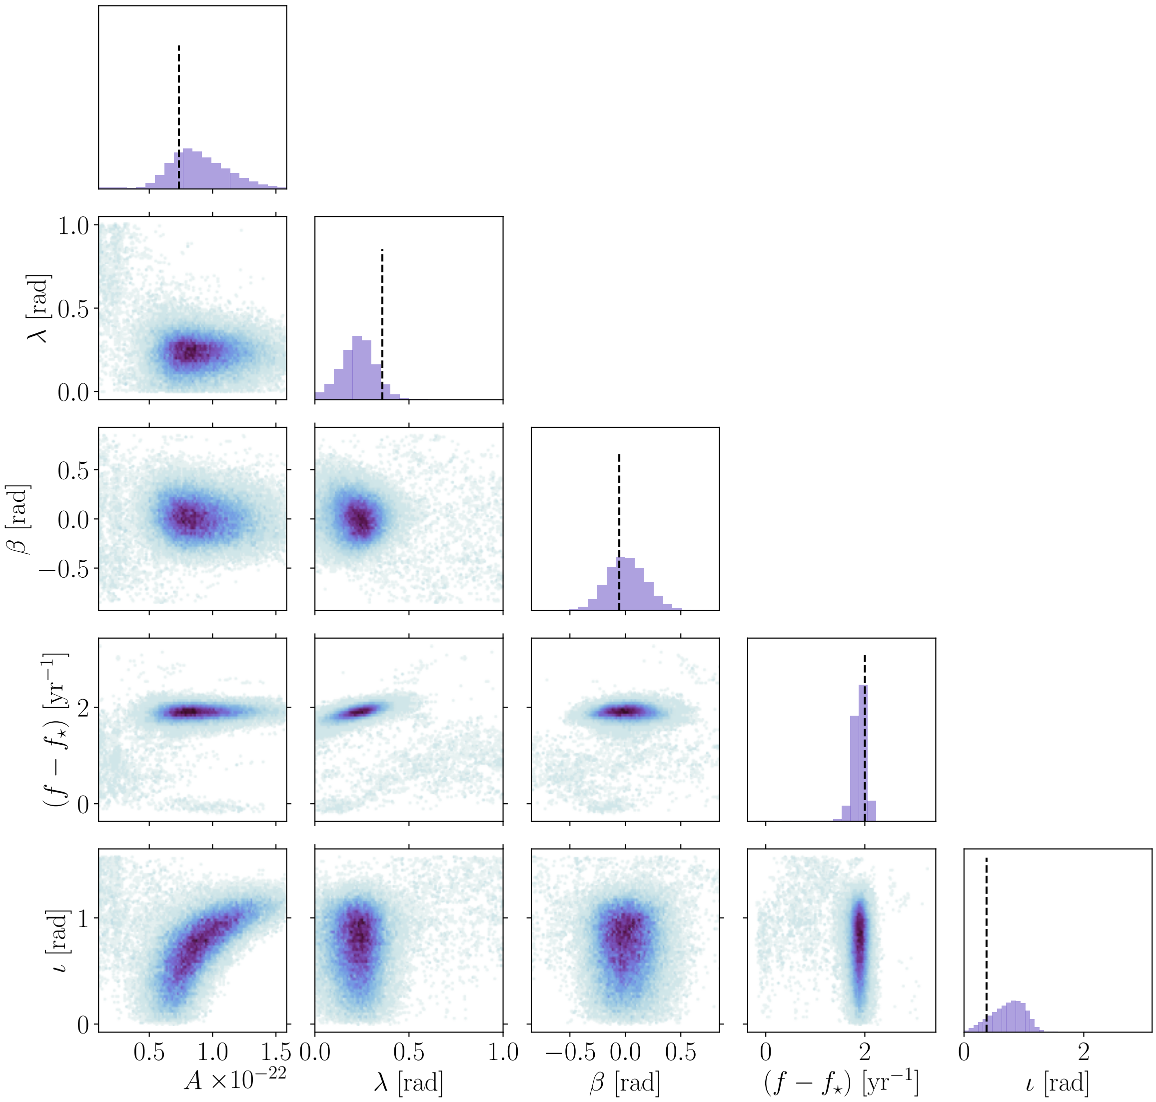

FIGURE 3.7: Posteriors on selected parameters from the third galactic binary. Vertical lines show the true injected values.

Chapter 4

Binary White Dwarfs in Milky Way Satellites

Contribution summary

This Chapter is a partially edited and reformatted version of [2]:

E. Roebber, R. Buscicchio, A. Vecchio, C.J. Moore, A. Klein, V. Korol, S. Toonen, D. Gerosa, J.M. Goldstein, S.M. Gaebel, T.E. Woods - Milky Way satellites shining bright in gravitational waves - published in Astrophysical Journal Letters, Volume 894 Issue 2:L15, (2020)

I contributed to conceive the study, to design and perform the parameter estimation campaign, and to produce the content reported in all sections except Section 4.2, which was led by the co-author V. Korol. Results presented rely on a long-term development of a codebase which all co-authors have contributed to. The code used for performing parameter estimation is not publicly available, yet. Injected source parameters and posterior samples are released in [201].. I've produced all plots shown in this Chapter, drafted and finalized the draft in collaboration with the co-authors.

4.1 Introduction

The identification and characterization of Milky Way (MW) satellite galaxies lie at the intersection of several outstanding problems in cosmology, astrophysics and fundamental physics [202]. These include the nature of dark matter, the formation and evolution of the faintest galaxies, and their reionization history. Faint satellites also offer us the opportunity to study star formation in low-metallicity environments and systems chemically different compared to the MW, which may be relevant for the origin of r-process and heavy elements.

Following the serendipitous discovery of the Sculptor dwarf galaxy [203], only a dozen other MW satellites were known up until approximately 2010. The Sloan Sky Digital Survey (SSDS) (and subsequently DECam, DES and Pan-STARRS, with the recent addition of Gaia [204]) has transformed the field raising the number to around 60; however, at least twice as many satellites are thought to exist, and the number could be nearly an order of magnitude higher [205]. The observational effort to complete the census of the MW satellites is made particularly arduous by the need to detect galaxies with luminosities down to $\sim 10^5\,L_\odot$. The next leap is expected with the Large Synoptic Survey Telescope [206]. By the end of the decade, LSST should provide a complete sample for distances up to $\sim 1\,\mathrm{Mpc}$ and luminosities down to $\sim 2 \times 10^3\,L_\odot$, and could detect any novae and supernovae in faint dwarf galaxies out to much greater volumes [207]. The spectroscopic characterization of these satellites will remain a major challenge, probably requiring 30 m class telescopes, and no survey will be able to observe within $\sim \pm 10°$ of the galactic plane [205].

The Laser Interferometer Space Antenna [208] is a millihertz gravitational-wave (GW) observatory planned for launch in 2034. LISA will survey the entire sky with a depth of a few hundred kpc for double white dwarfs (DWDs) and other solar-mass binary compact objects with orbital periods $\lesssim 10$ min [209].

In this Letter we show that LISA could provide new and complementary information about MW satellites using populations of short-period DWDs as tracers of these dwarf galaxies, and as markers of the astrophysical processes and conditions within their unusual (compared to the MW) environments. We will also show that LISA will observe tens of DWDs within the Large and Small Magellanic Clouds (LMC and SMC) and will unambiguously place them within specific regions of the Clouds. A handful of short-period DWDs should also be observable in other satellites. If located above $\sim 30°$ of the galactic plane, they can be easily associated to their host, since galactic DWD foreground sources are rare. At frequencies above a few mHz, LISA can also probe the zone of avoidance around the galactic plane.

4.2 Expected DWD population

To date no undisputed DWD is known in MW satellites. An X-ray source, RX J0439.8-6809, has been tentatively identified as a compact accreting WD system with a He WD donor in the LMC [210, 211], although later spectral modeling suggests this object may also be consistent with an unusually hot WD in the MW halo [212]. This lack of observational evidence is due to the faintness of these systems. They are undetectable by optical telescopes at the distance at which satellites are typically found—the median distance of known satellites is ~ 85 kpc, see [205].

4.2.1 Astrophysical modeling

A companion paper by [215] investigates the population of DWDs radiating in the LISA sensitivity band in MW satellite galaxies (see also [118]). In that paper, a suite of

	LMC	SMC	Sagittarius	Fornax	Sculptor
Stellar Mass ($M_\odot$)	1.5×10^9	4.6×10^8	2.1×10^7	2.0×10^7	2.3×10^6
Distance (kpc)	50.0	60.6	26.7	139	86
Ecliptic latitude β	-85.4°	-64.6°	-7.6°	-46.9°	-36.5°
Galactic latitude b	-32.9°	-44.3°	-14.2°	-65.7°	-83.2°
Galaxy area (deg^2)	77	13	37	0.17	0.076
Foreground sources	1	0.2	20	10^{-3}	3×10^{-4}
Expected sources (optimistic)	>100	>25	10	0.2	0.07
Expected sources (pessimistic)	70	15	3	0.1	<0.04
Sky localization (deg^2)	2.1	3.1	2.3	–	9.3

TABLE 4.1: Promising satellites for GW detection. Mass, distance, and sky location are taken from [213, 214, 205]. The expected number of LISA sources is estimated using the models of [215]. The sky localization is the 90% area recovered for the fiducial DWD described in Section 4.4 for each host satellite. We assume a 4 year mission duration and the SciRD noise spectral density [216].

models that span metallicity, star formation history (SFH) and unstable mass transfer phase are constructed using the population synthesis code `SeBa` and calibrated against state-of-the-art observations of DWDs [217, 110, 218, 219] . Here we summarize the main assumptions and results, and we refer the reader to the companion paper for details.

Despite the many uncertainties surrounding the composition and formation history of these satellites, the parameters crucial for determining the number of sources detectable by LISA are: (i) the total stellar mass $M_\star$, which sets the fuel supply used to generate stars and (ii) the star formation history (SFH), which controls the mass and frequency distribution of DWDs within the LISA sensitivity band at the present time.

Star formation histories in MW dwarf satellites vary greatly, ranging from purely old populations (formed over 12 Gyr ago) to constantly star forming [220, 221, 222]. To cover the range of possible SFHs we consider a constant star formation rate of $1\,\mathrm{M}_\odot\,\mathrm{yr}^{-1}$ and an exponentially decaying one with characteristic timescale $\tau_{\mathrm{SF}} = 5\,\mathrm{Gyr}$ [221], as optimistic and pessimistic star formation models, respectively.

By setting the metallicity to $Z = 0.01$, the binary fraction to 50% and the initial mass function to [223], the optimistic (pessimistic) SFH model predicts 0.2 (0.1) detectable sources for a satellite with $M_\star = 10^7 \mathrm{M}_\odot$ at the distance of 100 kpc. Results scales linearly with the mass of the satellite. Other unconstrained parameters, such as metallicity, binary fraction and unstable mass transfer have very minor impacts on the detectable DWD rate and, together, affect predictions by only a factor of a few.

4.2.2 Known satellites

Table 4.1 summarizes the properties of selected known MW satellites and the expected number of DWDs that can be observed by LISA according to the population synthesis models. We assume a mission duration $T_{\rm obs} = 4\,\mathrm{yr}$ and a noise spectral density corresponding to the LISA Science Requirements Document [216]. The choice of noise spectral density has a significant effect on the number of sources expected—the SciRD sensitivity curve is a factor 1.15, 1.4, and 1.5 *worse* than the original LISA noise curve (see Figure 1 in [208]) at 3, 5, and $\gtrsim 10$ mHz, respectively. Using the more optimistic noise curve of [208], as in many previous studies [224, 225, 226], would roughly double the expected number of sources.

The number of DWDs that we expect to see in a particular satellite depends strongly on the mass of the satellite, on its SFH, and on its distance. It depends somewhat less strongly on the ecliptic latitude of the satellite via the weakly directional "pointing" of the LISA instrument. The Magellanic Clouds and the Sagittarius, Fornax, and Sculptor dwarf spheroidal galaxies are promising systems to host detectable LISA sources [215].

The LMC and SMC are by far the most massive known satellites of the MW. They are expected to contain 10^2–10^3 detectable DWDs [215]. Sagittarius is expected to host several detectable sources, even for a pessimistic SFH model. The rates for Fornax, Sculptor, and smaller galaxies are lower, but these predictions depend on the specific details of the SFH.

Other satellites can be reached by LISA, but may already have exhausted their reservoir of observable DWDs. LISA is thus in a position to study details of the LMC/SMC, detect a handful of DWDs in some of the more massive satellites, and identify systems in other satellites if they have undergone recent star formation. Furthermore, LISA has the unique opportunity to discover new MW satellites.

(a) Masses

	$m_1\,(M_\odot)$	$m_2\,(M_\odot)$	$\mathcal{M}\,(M_\odot)$
He WDs	0.4	0.35	0.33
Typical WDs	0.6	0.55	0.5
Heavier WDs	0.7	0.65	0.59
Extremely low-mass	0.7	0.2	0.31
Type Ia SN progenitors	0.9	0.85	0.79

(b) Frequencies

f_0 (mHz)
2
3
4
5
10

(c) Inclinations

	ι (rad)
face-on	0
intermediate	$\pi/3$
edge-on	$\pi/2$

TABLE 4.2: Parameters used in our 4200 runs. We grid over these parameters as well as our sample of 56 dwarf galaxies.

4.3 LISA signal recovery

Having established that LISA can and will observe DWDs hosted by MW satellite galaxies, we need to consider whether it will be possible to associate these DWDs with the actual host satellite. The challenge is further exacerbated by the fact that LISA will observe ten to fifty thousand galactic DWDs, in addition to the unresolved stochastic foreground produced by $\sim 10^6$ DWDs. Making these associations will depend on LISA's ability to measure source sky locations and distances for these sources. We investigate this by performing full parameter estimation analyses on mock LISA data which we generate for a range of plausible sources within the satellites.

We consider a number of DWD systems that we expect to populate these satellites, spanning mass, frequency and binary inclination (see Table 4.2). For each choice of these parameters (75 combinations in total) we place a binary randomly within each of the 54 satellites in [205], together with the LMC and SMC. The distance and angular size of these satellites are taken from [213] and [214].

We generate mock LISA data sets lasting $T_{\rm obs} = 4\,{\rm yr}$ and containing the individual DWDs with zero noise. We recover the sources using the conservative LISA SciRD noise power spectral density [216] generated with LISACode [227], with an estimation of the galactic confusion noise taken from [228]. Gravitational radiation from the DWDs is treated as a quasi-monochromatic signal with linear drifts in frequency:

$$f_{\rm GW}(t) = f_0 + \dot{f}_0 (t - t_0). \tag{4.1}$$

We model the effect of these GW signals on the three noise-orthogonal channels A, E and T [229, 194], and process the resulting data using a coherent Bayesian analysis.

Each of the signals is described by 8 unknown parameters:

$$\{\mathcal{A}, f_0, \dot{f}_0, \lambda, \beta, \iota, \psi, \phi_0\}\,, \tag{4.2}$$

where $\mathcal{A}$ is the GW amplitude, (λ, β) are the ecliptic longitude and latitude, respectively, ι is the inclination angle, ψ is the polarization angle, and ϕ_0 is an arbitrary initial phase. The GW amplitude is given by

$$\mathcal{A} = \frac{2(G\mathcal{M})^{5/3}}{c^4 D}(\pi f_0)^{2/3}. \tag{4.3}$$

This is set by the source's distance D and chirp mass

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \tag{4.4}$$

for component masses m_1 and m_2.

For each signal injection, the GW amplitude, frequency, sky position, and inclination are chosen from our grid defined in Table 4.2. The polarization and initial phase are chosen randomly with a flat distribution. Finally, $\dot{f}_0$ is chosen according to the gravitational radiation reaction:

$$\dot{f}_0 = \frac{96}{5}\frac{(G\mathcal{M})^{5/3}}{\pi c^5}(\pi f_0)^{11/3}. \tag{4.5}$$

During parameter estimation, we treat $\dot{f}_0$ as an unknown parameter which can take either positive or negative values to account for the possibility of accretion affecting the period evolution of the system [230]. Population synthesis studies predict that $< 10\%$ of DWD systems in this frequency range will be in mass-transfer states [231]. Priors are chosen to be flat in $\log \mathcal{A}$, $\sin \beta$, $\cos \iota$, and flat in all other parameters in Eq. (4.2).

Our grid over parameters and satellites covers a range of sources from the very quiet to the very loud. We consider a source to be detected if the coherent signal-to-noise ratio exceeds 7. This is a conservative threshold with respect to previous ones [232] and [233], where detection thresholds of 5 and 5.7 were chosen, respectively, for monochromatic sources. A more accurate choice of detection threshold will be informed by the distribution of false alarms originating from noise in the real data, once available. Of our 4200 injected sources, 1954 are detected. For all satellites, at least one combination of the parameters produces a detectable DWD. For nearby satellites, a large range of parameters produce detectable systems.

To help summarize our results, hereafter we will focus on the following five satellites: the LMC, SMC, Sagittarius, Sculptor, and Fornax (see Table 4.1). These satellites span a broad range in distance, ecliptic latitude, and angular scale and are the most likely to host detectable DWDs [215]. We will also focus our discussion on a fiducial source with a chirp mass of $\mathcal{M} = 0.5\, M_{\odot}$, radiating at $f_0 = 5$ mHz, and with an intermediate inclination of $\iota = \pi/3$. The detectability of such a source for the complete set of satellites is shown in Figure 4.1. If any such system is present within $\sim$120 kpc, it will be detectable by LISA. This represents roughly half the known MW satellites, including all our highlighted satellites except Fornax, which is at a distance of 139 kpc.

4.4 Host satellite identification

We have shown that LISA will be sensitive to DWDs radiating at a few mHz in the MW satellites. However, it is not immediately obvious that these sources can be robustly associated with their host satellites. In this section, we will consider three pieces of information to solve this problem: the source sky localization, the anisotropic

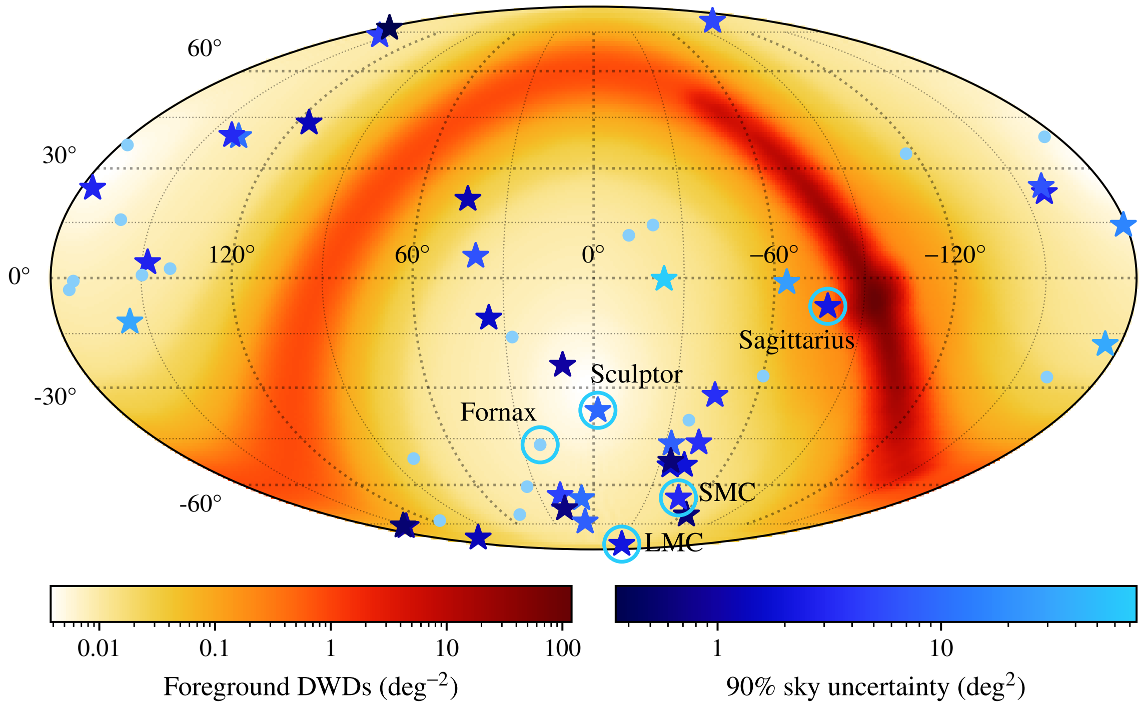

FIGURE 4.1: LISA sensitivity to a fiducial source with $\mathcal{M} = 0.5\,M_\odot$, $f = 5$ mHz, and $\iota = \pi/3$ in each satellite. Light blue dots are 'undetected' sources ($\rho < 7$). Stars are 'detected' sources ($\rho > 7$), and are color-coded according to the quality of their sky localizations. For a fixed frequency, this is largely governed by distance and ecliptic latitude. The largest sky uncertainty, for a source with $\rho = 7.6$, $D = 118$ kpc, and $\beta = -0.3°$, is 75 deg^2. The smallest, for a source with $\rho = 36$, $D = 22$ kpc, and $\beta = 77°$, is 0.3 deg^2. Satellites of interest are highlighted with blue circles. In this case, LISA is sensitive to systems at distances of $\lesssim 120$ kpc, which excludes Fornax.

distribution of foreground MW DWDs, and measurements of the source distance.

At mHz frequencies, LISA's angular resolution is good. Most injections of our fiducial source can be located to within $\lesssim$10 deg^2 (see Figure 4.1); the exceptions are low-SNR sources near the ecliptic. This means that sources inside the LMC, SMC, and Sagittarius (all of which are larger than 10 deg^2), can potentially be localized to specific regions of the satellites. The sky uncertainty depends strongly on the SNR and GW frequency ($\propto \rho^{-2} f_0^{-2}$), and also on the ecliptic latitude (a source on the ecliptic has an order of magnitude more uncertainty than a source at the poles).

Equally important to the sky localization is the foreground of MW DWDs for each satellite. At frequencies $\gtrsim$ 3 mHz, MW DWDs become resolvable [228], so the stochastic MW foreground is not a significant concern here. We model the MW sources following [234], including a stellar halo generated with a single burst SFH, a power law density distribution according to [235], and a total mass of $1.4 \times 10^9\,\mathrm{M}_\odot$ [236]. The resulting foreground is strongly anisotropic, closely following the galactic plane (see Figure 4.1). Most known satellites are well away from the galactic plane, in regions with a foreground density of $\sim$ 0.01/deg^2. For a sky localization of $\sim$ 1–10 deg^2, this corresponds to $\sim$ 0.01–0.1 contaminating foreground sources . For known satellites (as well as unknown ones with similar sky positions and DWD source densities) this in turns corresponds to a typical false alarm probability between $\sim 5 \times 10^{-5}$ and $\sim 5 \times 10^{-3}$. At lower (higher) frequencies, the sky localization is worse (better) and the false alarm rate rises (falls).

In addition to associations based on the sky localization, the frequency evolution for sources above 3–4 mHz will be measurable. Our 90% fractional errors on $\dot{f}$ are distributed according to:

$$\Sigma_{\dot{f}} \approx 0.07 \left(\frac{\rho}{10}\right)^{-1} \left(\frac{f}{5\ \mathrm{mHz}}\right)^{-11/3} \left(\frac{\mathcal{M}}{0.5\, M_\odot}\right)^{-5/3}. \tag{4.6}$$

Assuming the inspiral is driven by radiation reaction (Equation 4.5), measurements of $\dot{f}$ and $\mathcal{A}$ permit the measurement of the distance to the satellite with a precision of $\sim 30\%$. Stellar interactions within DWDs will reduce, not increase, the total $\dot{f}$ [230]. This implies that a lower limit can safely be set on the distance (see Figure 4.2), thereby further reducing the chance of a false positive.

Let us examine some cases in detail. The LMC and SMC are large satellites with many expected sources (~ 100 and ~ 20, respectively). Both galaxies have large angular extents ($77\,\mathrm{deg}^2$ and $13\,\mathrm{deg}^2$,respectively) and high ecliptic latitudes, so sub-galaxy localizations of sources are likely, with typical source localization of a few degrees squared. The SMC is in a region of the sky with very few foreground sources, so statistical associations can be readily made , with a false alarm probability of 4% within the source sky localization uncertainty. This is also true for the LMC, but its situation is complicated by a partial degeneracy in the LISA response at extreme ecliptic latitudes. This may result in a larger foreground than stated in Table 4.1, due to the presence of MW DWD sources at the other ecliptic pole. Distances for sources in the Magellanic clouds should be well-measured with 40% to 50% fractional error, and will help make associations robust.

Sagittarius is a relatively massive and nearby dwarf spheroidal galaxy, so several detectable sources are expected. Its unusual SFH means that the 'optimistic' case of 10 sources from Table 4.1 is quite plausible [215]. Unfortunately, its location near the galactic bulge and its large angular scale lead to a large number of foreground sources (although note that the foreground varies from 0.1 to 1 per squared degree across the satellite, with false alarm probabilities that can be as low as 0.2.). Frequency measurements can be used to partially remove the foreground. If we consider only sources with $f > 3$ mHz, the foreground drops from 20 to 5. Distances for this satellite

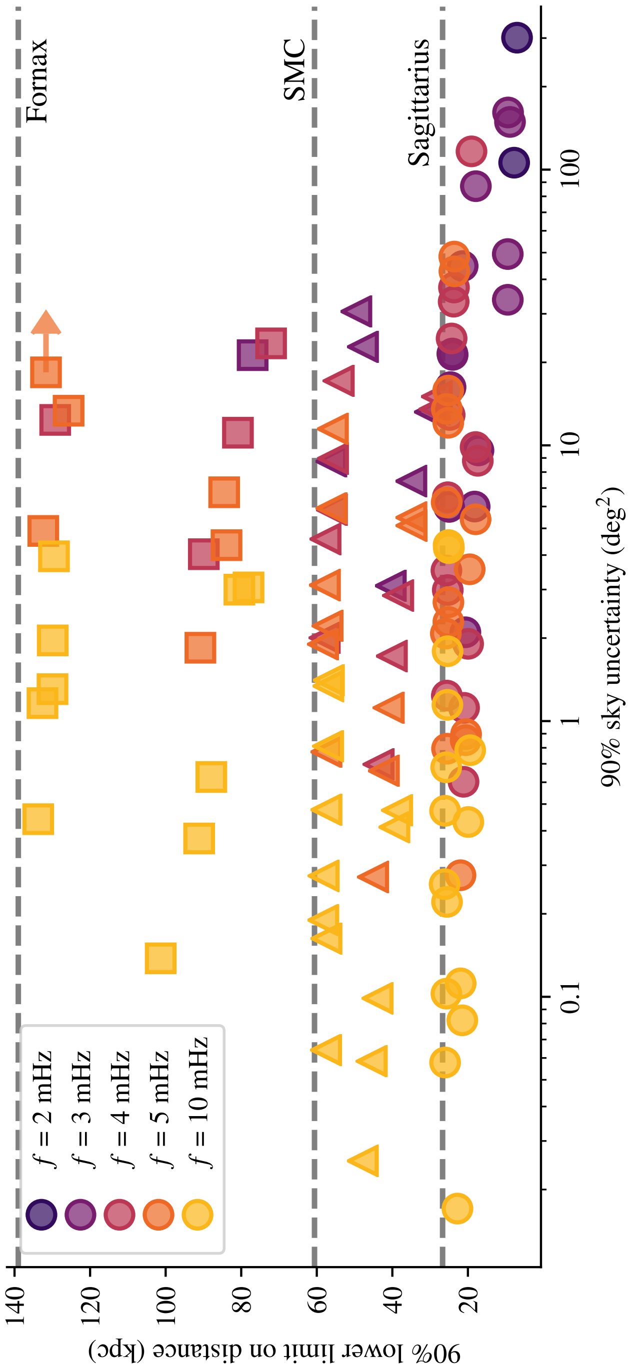

FIGURE 4.2: Distance lower limits and sky localizations for all "detected" runs in Sagittarius (circles), the SMC (triangles), and Fornax (squares). Dashed lines mark the true distance of each satellite. For a given frequency, the sky localization is primarily affected by the source's mass, and the lower limit on the distance is primarily affected by the source's inclination. Lower limits on the sky localization are given for one source which does not have a well-defined 90% sky area, but does have a well-defined area at lower confidence.

are likely to be well-measured, but as Sagittarius is well within the MW halo, their additional constraining power will be somewhat reduced. Robust associations with Sagittarius will be non-trivial, but careful modelling of the MW population should make it possible.

Fornax, Sculptor, and the other dwarf satellites are too small or too distant to be likely hosts of LISA sources. However, it is possible that uncertainties in the SFH, perhaps combined with a more optimistic LISA noise curve will produce detectable DWDs. In this case, the sources will be readily identifiable as the foregrounds are small, with false alarm probabilities of 1% for pessimistic sky localizations of $10\,\mathrm{deg}^2$. Moreover distance lower limits (particularly in the case of Fornax) would provide strong evidence for the satellite association.

4.5 Discovering hidden satellites of the Milky Way

Unlike light, GWs are not impeded by dust and gas. Moreover, above a few mHz, DWDs become individually resolvable and the MW no longer acts as a GW confusion-noise foreground. This gives LISA an advantage over EM telescopes in that it can peer through the galactic plane and possibly make discoveries on the far side. Currently, the best example of a satellite near the galactic plane is the recently discovered Antlia 2 which has a galactic latitude of $\sim 11^\circ$ [237]. However, at lower latitudes dust extinction increases dramatically; therefore, even objects as large as the LMC could have remained undetected.

If such an object exists, LISA could potentially detect high-frequency DWDs from it. The question is then whether these detections are sufficient to infer the presence of the hidden satellite. This task is complicated by the high density of resolvable foreground

DWD sources in the galactic plane.

To illustrate the discovery potential of LISA, consider a hypothetical satellite, similar to the LMC, at a distance of 50 kpc behind the disk of the MW. We assume that it has an angular diameter of $10°$, a mass of $1.5 \times 10^9\,M_\odot$, a fixed metallicity of $Z = 0.005$, a constant star formation rate, and an age of 13.5 Gyr [215]. This object could be completely covered by the galactic disk.

The foreground density of DWD sources in the disk is $\sim 100/\mathrm{deg}^2$ (see Figure 4.1). If galactic sources are distributed uniformly throughout the disk, which has a total area of $\sim 3000\ \mathrm{deg}^2$, then a simple Poisson counting argument suggests that an excess of ~ 100 sources in an 80 deg^2 patch of the sky would be a significant overdensity at the 90% level.

Based on previous studies (see e.g. Fig.4 in [215]) we expect ~ 100 detectable sources in our hypothetical satellite, so an LMC-like satellite at $\lesssim 50$ kpc should appear as a statistically significant overdensity. At greater distances, it would have too few sources to overcome the foreground. This calculation assumes a similar stellar density to the LMC; a denser (sparser) satellite would be detectable at a greater (lower) maximum distance. Furthermore, we assume a uniform, Poissonian distribution of DWDs in the galactic disk—a more realistic non-uniform distribution will require a larger overdensity to be significant.

However, we have not yet considered distance measurements. Section 4.4 suggests that the majority of detectable extragalactic sources will be chirping, meaning that lower bounds can be placed on their distances. This will allow us to distinguish them from the foreground and detect satellites out to greater distances. We assume that chirping sources allow us to place a lower limit on the distance of $\sim 60\%$ of the true value (although many sources do considerably better—see Figure 4.2). A satellite at

50 kpc with multiple detected sources can be confidently placed at $\gtrsim 30$ kpc, which is greater than the distance to any DWD in the galactic disk. At 150 kpc (200 kpc) we expect to detect ~ 10 (~ 3) sources from our hypothetical satellite which can likewise be distinguished from disk foreground sources (although a small number of halo sources remain as contaminants).

The galactic plane obscures $\sim 10\%$ of the sky. For the first time, LISA will be able to survey this region for major MW satellites out to astrophysically interesting distances of $\lesssim 200$ kpc.

4.6 Conclusions

We have shown that if a population of DWDs emitting GWs at $\gtrsim 3$ mHz exists in the MW satellites, LISA will be able to detect them. Although the exact rate depends on the star formation history of each satellite, it is probable that many such DWDs will be detected in several different satellites. Moreover, in this frequency band, LISA will provide sky localizations of ~ 10 deg^2 and distance measurements with errors of $\sim 30\%$. This means that LISA should be able to associate these DWDs to their host satellites. Finally, at frequencies above a few mHz, the galactic confusion noise clears, and LISA can see through the galactic disk and bulge. This fact, combined with the arguments above, suggests that LISA might be capable of discovering hidden satellites of the MW, provided they are sufficiently massive.

Observations of short-period extragalactic DWDs will naturally occur as part of the LISA survey of the galactic DWD population. These observations will complement those of large optical surveys, since the selection effects are very different. The possibility of detecting short-period DWDs in MW satellites highlights the discovery space opened up by a GW observatory and its potential impact on a wide range of open questions in astrophysics and cosmology, from low-metallicity star formation history and heavy element nucleosynthesis to small-scale cosmology in the nearby Universe.

Chapter 5

Stellar-mass Binary Black Holes with LISA

Contribution summary

This Chapter is a partially edited and reformatted version of [3]:

R. Buscicchio, A. Klein, E. Roebber, C.J. Moore, D. Gerosa, E. Finch, and A. Vecchio - Bayesian parameter estimation of stellar-mass black-hole binaries with LISA - published in Physical Review D, Volume 104(4):044065, (2021).

I contributed to conceive the study, performed the parameter estimation campaign, and produced the results reported in all sections. Results presented rely on a long-term development of a codebase which all co-authors have contributed to. A. Klein is the major contributor to the code required for the results presented in this Chapter. I've produced all plots shown in this Chapter, drafted and finalized the text incorporating comments and suggestions from A. Klein and the other co-authors.

5.1 Introduction

The Laser Interferometer Space Antenna (LISA) [16] is a gravitational-wave (GW) observatory targeted at the discovery and precise study of compact binary systems ranging from white dwarfs of masses ~ 0.1–$1\,M_\odot$ to black holes with masses up to $\sim 10^7\,M_\odot$. Cosmological phenomena with characteristic timescale between ~ 1 hr and ~ 10 sec might also be detectable.

One of the sources of great interest are stellar-mass binary black holes (hereafter SmBBHs, also referred to as stellar-origin binary black holes, SOBBHs[1]) in the mass range ~ 10–$100\,M_\odot$ which populate LISA's sensitivity window at frequencies $f \gtrsim 10$ mHz. These systems are now routinely observed merging at ~ 100 Hz by the ground-based laser interferometers LIGO and Virgo [238, 239]. LISA is expected to observe ~ 1–10 SmBBHs during the whole mission [240, 241, 242, 243], an estimate that crucially depends on the upper-mass cutoff of SmBBHs, the detection strategy, as well as the LISA performance at high frequencies. Each of these systems will contain valuable information in terms of both astrophysical formation channels [31, 244, 245, 246] and fundamental physics constraints [247, 248, 249, 250]. The subset of sources that merge on a timescale of $O(10)$ yr will be even more unique, allowing for its multiband characterization using GWs from both space and the ground [251], as well as advanced planning of electromagnetic follow-up campaigns [240, 252].

Because of the long duration and complex morphology of their signals, detecting and characterizing the physics of SmBBHs with LISA is a highly nontrivial challenge. Previous work has been carried to study parameter estimation for circular precessing systems [253]. In this context, we tackle for the first time the effects of orbital eccentricity

[1] We prefer to use SmBBHs instead of SOBBHs to acknowledge that the problem of detecting and characterizing these sources is independent of the (astro)physics that determines their formation.

coupled with spin precession.

The importance of SmBBHs for the LISA science case is well recognized. To this end, a set of LISA data challenges (LDCs) are in progress under the auspices of the LISA consortium as part of the core preparation activities for the mission adoption (see lisa-ldc.lal.in2p3.fr). The first set of these challenges (LDC–1) contains mock datasets populated with SmBBHs. The analyzed systems from LDC–1 are illustrated in Fig. 5.1. Most of the sources appear as quasimonochromatic. However, a few of them merge within the observing time (which in LDC–1 was set to 2.5 yr), thus allowing a finer characterization of their parameters through the chirping morphology.

Here we report on the results of the analysis of all SmBBHs in LDC–1 using the generic Bayesian codebase we are developing, hereafter referred to as Balrog. While the LDC–1 sources were injected assuming quasicircular binaries with aligned spins, we also present preliminary results on the more general problem of analyzing systems with orbital eccentricity and spin precession. Overall, this paper quantifies how well SmBBHs can be characterized with LISA once they have been detected.

This paper is organized as follows. In Sec. 5.2, we describe our data analysis strategy and outline its technical implementation. In Sec. 5.3, we present the challenge dataset we analyzed and our inference results. In Sec. 5.4, we provide our conclusions and pointers to future work. Throughout the paper we use $G = c = 1$.

5.2 Analysis approach

5.2.1 Inspiralling stellar-mass black-hole binaries

SmBBHs in the early inspiral region probed by LISA are long-lived sources radiating for most or all of the mission duration, depending on their masses and orbital period at

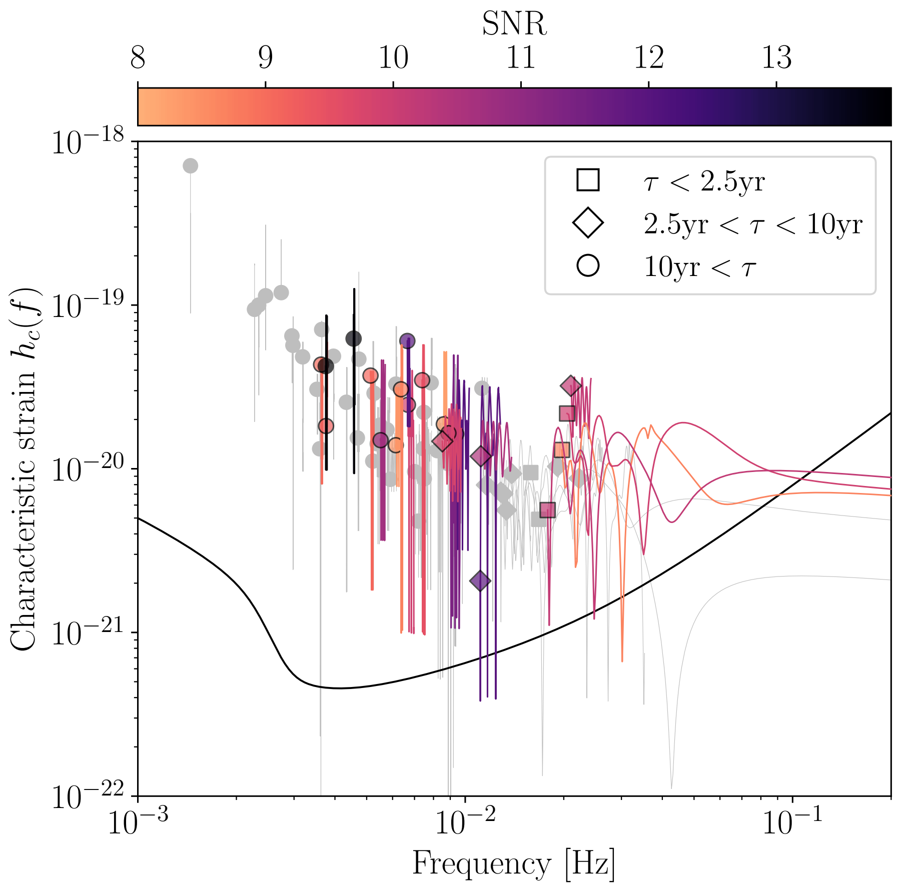

FIGURE 5.1: Characteristic strain of the injected challenge dataset sources. For low-frequency, quasimonochromatic sources, the characteristic strain is modulated by the LISA orbit throughout the mission. Markers denote the sources' initial frequencies, continuous lines denote their spectral characteristic strain amplitude. Binaries merging within the dataset duration of 2.5 yr are marked by squares, while diamonds and circles indicate binaries coalescing in 2.5-10 and over 10 years, respectively. More details on the source properties are outlined in Sec. 5.2. Lines and markers are colored according to the signal-to-noise ratios (SNRs); unresolved sources with SNR < 8 are grayed. Note that characteristic strain amplitudes were constructed using the low-frequency approximation of the LISA response [254]. The solid black line denotes the LISA characteristic noise spectral amplitude [255, 12].

the start of the mission. In fact, for a binary with component masses m_1 and m_2 at redshift z, the leading Newtonian order time to coalescence is [29, 256]

$$\tau \approx 4.1 \left(\frac{\nu}{1/4}\right)^{-1} \left(\frac{f}{20\,\mathrm{mHz}}\right)^{-8/3} \left(\frac{M_z}{50\,M_\odot}\right)^{-5/3} \mathrm{yr}\,, \tag{5.1}$$

where $M_z = (1+z)(m_1+m_2)$ is the redshifted total mass, $\nu = m_1 m_2/(m_1+m_2)^2$ is the symmetric mass ratio, and f is the GW frequency. During this period the (leading Newtonian order) number of wave cycles to merger is

$$\mathcal{N} \approx 4.1 \times 10^6 \left(\frac{\nu}{1/4}\right)^{-1} \left(\frac{f}{20\,\mathrm{mHz}}\right)^{-5/3} \left(\frac{M_z}{50\,M_\odot}\right)^{-5/3}. \tag{5.2}$$

Consequently, if the source merges in a few years, i.e. unless $\dot{f} T_{\mathrm{obs}} \ll f$, most of the wave cycles are accumulated in the LISA band. These cycles need to be matched by the analysis, in sharp contrast with the current LIGO–Virgo observations, for which only a few or tens of cycles are in band.

The signal will also have complex features induced by spin-precession and the effects of eccentricity. This adds complexity to the waveform and to the structure of the likelihood function, and increases the dimensionality of the parameter space.

If the black-hole spins are misaligned with the orbital angular momentum, this will induce a precession of the orbital plane around the axis of the total angular momentum characterized by a number of spin-precession cycles before merger [257, 258]

$$\begin{aligned} \mathcal{N}_{\mathrm{spin}} \approx 1.9 \times 10^3 \left(1 - \frac{\delta\mu^2}{7}\right) \left(\frac{\nu}{1/4}\right)^{-1} \\ \times \left(\frac{f}{20\,\mathrm{mHz}}\right)^{-1} \left(\frac{M_z}{50\,M_\odot}\right)^{-1}, \end{aligned} \tag{5.3}$$

where $\delta\mu = (m_1 - m_2)/(m_1+m_2)$ is the dimensionless mass difference.

Eccentricity may also be non-negligible in the LISA band, and an important parameter to measure as it is a tracer of the environment in which these binaries reside and the formation channel(s) of these systems. The number of periastron precession cycles before merger is [29, 30]

$$\mathcal{N}_{\rm ecc} \approx 6.4\times 10^3 \left(\frac{\nu}{1/4}\right)^{-1}\left(\frac{f}{20\,{\rm mHz}}\right)^{-1}\left(\frac{M_z}{50\,M_\odot}\right)^{-1}. \tag{5.4}$$

Note that these estimates are valid in the low-eccentricity limit.

It is therefore clear that to accurately reconstruct the physics of SmBBHs, one needs to deal with the full complexity of the 17 dimensions parameter space that describes GWs radiated by a binary system in general relativity. The morphology of these SmBBH signals is very different from both those currently observed by LIGO and Virgo as well as the supermassive BBH merger signals expected in LISA. In fact, these signals have more in common with the extreme-mass-ratio inspiral (EMRI) signals also expected in LISA which also contain $10^5 - 10^6$ wave cycles in band and can exhibit strong relativistic precession effects. The data analysis challenges presented by this source type are well-known [259, 260, 261]. In addition, EMRI present a severe modelling challenge, see e.g. [262, 263, 264].

5.2.2 Statistical inference

In this paper we are not concerned with the (significant) challenge of actually searching for SmBBHs [243], but we restrict ourselves to the problem of measuring the source parameters once candidates have been initially identified through a first search stage. We will therefore assume that a preceding pipeline provides an initial, possibly poor guess of the source parameters on which we can deploy our Bayesian parameter-estimation approach.

Our analysis is performed using the three noise-orthogonal time-delay-interferometry (TDI) outputs that are generated by combining the readouts of the LISA phase-meters [265]. This stage suppresses by a factor $\approx 10^8$ the laser noise leaving the data stream only affected by the secondary noise sources and GWs. The details of this complex procedure are under active investigation and development, see e.g. Refs. [266, 267, 268, 269].

We employ a coherent analysis of the full LISA TDI outputs, $d = \{d_k; k = A, E, T\}$, by means of Bayesian inference. The likelihood, $\mathcal{L}(d|\boldsymbol{\theta})$, of the data d given the parameters $\boldsymbol{\theta}$ of the source is [270]

$$\ln \mathcal{L}(d|\boldsymbol{\theta}) = -\sum_k \frac{\langle d_k - h_k(\boldsymbol{\theta})|d_k - h_k(\boldsymbol{\theta})\rangle_k}{2} + \text{const}, \tag{5.5}$$

where h_k is the TDI output k produced by the GW $h(t;\boldsymbol{\theta})$, or, equivalently, in the Fourier domain, $\tilde{h}(f;\boldsymbol{\theta})$. The inner-product is defined as

$$\langle a|b\rangle_k = 2\int_0^{+\infty} \mathrm{d}f \, \frac{\tilde{a}(f)\tilde{b}^*(f) + \tilde{a}^*(f)\tilde{b}(f)}{S_k(f)} \,, \tag{5.6}$$

where $\tilde{a}(f)$ is the Fourier transform of the time series $a(t)$, $S_k(f)$ is the noise power spectral density of the kth data stream, and the extrema $[f_{\rm low}, f_{\rm high}]$ corresponds to the frequency range spanned by a GW with parameters $\boldsymbol{\theta}$ over the duration of the observation.

Once a prior $p(\boldsymbol{\theta})$ is specified, we compute the joint posterior probability density function (PDF) on the parameters of the source

$$p(\boldsymbol{\theta}|d) \propto \mathcal{L}\left(d|\boldsymbol{\theta}\right) p(\boldsymbol{\theta}) \tag{5.7}$$

through stochastic sampling. BALROG is designed to work with different sampler flavors and implementations. For the analysis presented here we use a nested sampling algorithm based on CPNEST [271].

We model the gravitational waveforms $h(t; \boldsymbol{\theta})$ in their adiabatic inspiral regime through a post-Newtonian (PN) expansion, using two different waveform models. One of them is a new implementation under active development [258] which includes both spin precession and orbital eccentricity. Improving upon previous work [272], the new formulation is substantially more efficient in terms of computational requirements, making the analysis presented here possible. We also use a 3.5PN TAYLORF2 waveform (see e.g. [273, 274]) restricted to aligned spins and quasicircular orbits when analyzing the LDC–1 dataset, in agreement with the signals injected in it. The full set of waveforms we use are computed at sufficiently high PN order to ensure that no systematic effect is introduced in the analysis [275]. We describe the TDI outputs from such a model as in Ref. [276], which allows us to fully reproduce the waveforms used in LDC–1. Each source is described by 17 (11) parameters in the precessing and eccentric (spin-aligned and circular) case.

5.2.3 Implementation

The noise-orthogonal TDI outputs d on which the LISA GW analysis is based need to be computed from intermediate TDI data products, e.g. the TDI Michelson observables X, Y, and Z [277, 18]. Note that the noise-orthogonal data channels A, E, and T first introduced in the literature [265] were constructed from the Sagnac variables α, β, and γ and are therefore slightly different from the ones we are using here. Here for concreteness and to consistently interface with the data currently generated within the LDCs, we start from X, Y, and Z. This step will need to be revised in the future as the interplay between the raw phase-meter data and the actual GW analysis becomes clearer.

In order to improve our computational efficiency, we use a rigid adiabatic approximation (RAA) of the TDI variables [278], that is approximately related to the 1.5-generation variables injected into the datasets as

$$\tilde{X}_{1.5\text{-g}}(f) \approx \left(1 - e^{-4\pi i f L}\right) \tilde{X}_{\mathrm{RAA}}(f), \tag{5.8}$$

where $L = 2.5 \times 10^9$ m is the mean LISA armlength, and similarly for the other two TDI variables Y and Z. We note that SmBBH sources accumulate most of their SNR at the high frequency end of the LISA bandwidth ($f \gtrsim 5\,\mathrm{mHz}$, where $fL \gtrsim 1$; see Fig. 5.1). Therefore, a long-wavelength approximation to the detector response is not appropriate. We note that the RAA is not faithful to the full TDI response at very high frequencies. Since we recover source parameters from full TDI signals with RAA signals, this study also serves as a test of the RAA for SmBBH signals.

In order to compute inner products [cf. Eq. (5.6)] involving a discrete time series, we use a hybrid method based on Clenshaw-Curtis quadrature [279]. In its simplest form, this numerical integration technique consists of approximating an integrand (e.g. 5.6) by its Chebishev polynomial, and obtaining the integral value from those of a Fourier series of periodic functions, which are fast to compute and accurate to an extent under control by the Chebishev polynomial truncation.

In this work, we hybridize the method above as follows: first, the time series representing the data (having a cadence of 5 s in the LDC–1 case) is related through a discrete Fourier transform to a frequency series from $f_{\mathrm{min}} = 0$ to $f_{\mathrm{max}} = 0.1$ Hz, with a resolution of $\Delta f = 1/T_{\mathrm{obs}} \approx 1.27 \times 10^{-8}$ Hz. This defines a finite set of data points in the Fourier domain f_i^{DFT} with $f_{\mathrm{min}} \leq f_i^{\mathrm{DFT}} \leq f_{\mathrm{max}}$. We transform the frequency interval into a log-frequency interval, and split the latter into ten subintervals of equal length. In each of them, we compute a 21-point Clenshaw-Curtis quadrature rule, resulting overall in a set of $N = 201$ distinct frequencies f_i^{CC} with corresponding weights w_i^{CC}. For each of these points, we then find the closest frequency in the discrete set f_i^{DFT} to form the set f_i^{H}. In order to construct the associated weights w_i^{H}, we first note that each frequency f_i^{CC} satisfies either $f_i^{\mathrm{CC}} < f_0^{\mathrm{H}}$; $f_i^{\mathrm{CC}} > f_{N-1}^{\mathrm{H}}$; or $f_k^{\mathrm{H}} \leq f_i^{\mathrm{CC}} \leq f_{k+1}^{\mathrm{H}}$. In the first case, we associate the weight w_i^{CC} with w_0^{H}; in the second case we associate the weight w_i^{CC} with w_{N-1}^{H}; and in the third case we distribute the weight w_i^{CC} linearly (in log) between w_k^{H} and w_{k+1}^{H} according to the respective distance to f_i^{CC} of their corresponding frequencies. Finally, some frequencies in the set f_i^{H} might be duplicates, in which case we combine them and their weights for minor gains in computational efficiency. This results in a set of $N_H \leq N$ hybrid frequencies and weights allowing us to approximate the integral as

$$\int_a^b \mathrm{d}f\; g(f) \approx \sum_{k=0}^{N_H-1} w_k^{\mathrm{H}} g(f_k^{\mathrm{H}}). \tag{5.9}$$

We verified that the loss of accuracy in the integral evaluation due to the modification of the quadrature rule does not impact the result significantly. With this method, we drastically reduce the number of waveform evaluations necessary to evaluate each log-likelihood by selecting a few relevant datapoints in the discrete Fourier transform of the data. Note that this algorithm needs only to be used once at the beginning of the run, and that the computational efficiency of the resulting run is independent of the length of the time series, and weakly dependent on its cadence. We also stress that the choice of 10 subintervals with a 21-point Clenshaw-Curtis quadrature rule applied to them is somewhat arbitrary. We found that it yielded fast parameter estimation with good reliability in the case at hand, but it can certainly be optimized depending on the particular source analyzed. Moreover, further optimization of subintervals and quadrature's number of points would be required in the presence of noise in the data, e.g. tuning both choices to narrow noise spectral features in the data.

5.2.4 Sampling parameters

An appropriate choice of the sampling parameters is crucial to complete the inference. In order to remove the influence of uncertain cosmological effects from the analysis, in what follows we express all mass parameters in the detector frame (i.e. we use redshifted masses), unless stated otherwise. We choose to use the following set of 11 parameters to describe the circular, spin-aligned SmBBHs:

- For the two mass parameters, we use the chirp mass $\mathcal{M}_c$ and the dimensionless mass difference $\delta\mu = (m_1 - m_2)/(m_1 + m_2)$.

- The two amplitude parameters A_L and A_R are related to the luminosity distance D_L and the inclination ι by $A_L = (1 + \cos\iota)/\sqrt{2D_L}$ and $A_R = (1 - \cos\iota)/\sqrt{2D_L}$. They are the square roots of the amplitudes of the left- (right-)handed components of a GW.

- The two phase parameters ψ_L and ψ_R are related to the initial orbital phase ϕ_0 and the polarization phase ψ by $\psi_L = \phi_0 + \psi$ and $\psi_R = \phi_0 - \psi$. These are the initial phases of the left- (right-)handed components of a GW.

- The spins are described by the parameters $\chi_{1,\ell}$ and $\chi_{2,\ell}$, corresponding to the dimensionless spin magnitudes of the two binary components. In the general case of arbitrarily oriented spins(e.g.), these two parameters would correspond to the dimensionless spin projections onto the orbital angular momentum (hence the subscript ℓ) axis of the two binary components.

- The initial orbital frequency of the source $f_0^{\rm orb}$, related to the initial GW frequency by $f_0 = 2f_0^{\rm orb}$.

- The sine of the ecliptic latitude $\sin b$ and the ecliptic longitude l, as sky location parameters.

For the case of eccentric sources with precessing spins, one needs to modify and extend the sampled set of parameters. In particular, we choose the following:

- We parametrize eccentric orbits with the square eccentricity at $f = 10$ mHz e_{10}^2 and the initial argument of periastron ϕ_e.

- Precessing sources require six spin parameters; we choose the dimensionless spin magnitudes $\chi_{1,2}$ and the spin orientations in an ecliptic frame, which we describe by their sine latitudes $\sin b_{1,2}^{\chi}$ and longitudes $l_{1,2}^{\chi}$.

- For these runs, we also use the approximate time to merger t_M defined in Eq. (5.16) instead of the initial orbital frequency $f_0^{\rm orb}$.

We use flat priors for all the above parameters. Assuming that some information on the source will be provided by the preceding search stage, we restrict the prior range for (at least some of) the parameters around the injected values, which is essential to keep the computational burden at a manageable level (cf. Secs. 5.3.2 and 5.3.4).

This specific choice of parameters greatly simplifies the likelihood structure, thus facilitating the sampling process. We use the chirp mass $\mathcal{M}_c$ because this is the mass parameter entering the frequency evolution at lowest PN order and is thus better constrained than any other mass parameter. For the second mass parameter, our choice of $\delta\mu$ has a key advantage over the more traditional alternatives of the symmetric mass ratio ν or the mass ratio q: the Jacobian of the transformation into the (m_1, m_2) space is symmetric and regular in the $m_1 = m_2$ limit, avoiding potential issues related to this reparametrization. As shown in Fig. 5.2, the amplitude parameters A_L and A_R are weakly correlated, in contrast with the more common choices of luminosity distance D_L and inclination $\cos\iota$. Furthermore, a flat distribution in A_L and A_R corresponds to a flat distribution in $\cos\iota$, which we expect from an isotropic distribution of source locations and orbital angular momenta. Similarly, in the interest of avoiding strongly correlated quantities, we opted for the phase parameters ψ_L and ψ_R instead of ϕ_0 and ψ.

Figure 5.2 shows a comparison of the two-dimensional posteriors for an illustrative LDC–1 run (source 15), which has signal-to-noise ratio (SNR) 12. We contrast the parameter spaces described by (A_L, A_R) and $(D_L, \cos\iota)$ as well as that described by (ψ_L, ψ_R) and (ϕ_0, ψ). The posterior distributions of parameters related to the circular polarization of the GW are significantly less correlated compared to those involving linear polarization. Moreover, we only sample half of the (ϕ_0, ψ) plane, thus removing the multimodality arising from the following symmetries of gravitational radiation: $(\phi_0 \to \phi_0 + n\pi)$, $(\psi \to \psi + n\pi)$, $(\phi_0 \to \phi_0 + \pi/2, \psi \to \psi + \pi/2)$, $n \in \mathbb{Z}$. Note that the source we chose for illustrative purposes offers an unbiased measurement of the phases $\psi_{L,R}$, while we observed significant biases in the recovery of those two parameters for most sources. However, we did not observe such biases when analyzing data containing a single source, and we thus argue that this effect arises from the confusion between overlapping sources and is independent of the chosen parametrization. This illustrative source is thus representative of the single source injection results, and most relevant to this discussion.

5.3 Results

As an initial test of the analysis approach described in the previous section, we have applied it to the data sets released for LDC–1. The data sets are briefly described in Sec. 5.3.1. Details of prior choices and parameter estimation results are presented in Sec. 5.3.2. As LDC–1 had limited scope and contained only BHs on circular orbits with aligned spins, we also present in Sec. 5.3.4 a proof-of-concept analysis on a generic precessing and eccentric system.

5.3.1 LISA data challenge

The first round of the LISA Data Challenges (LDC–1, lisa-ldc.lal.in2p3.fr) consisted of several datasets, each of which was dedicated to a specific source class: massive black hole binaries, extreme-mass ratio inspirals, galactic binaries, SmBBHs, and stochastic backgrounds.

The LDC–1 SmBBH data consists of two sets, each containing the same 66 SmBBH injections, one noise-free and the other including a realization of the expected LISA Gaussian stationary noise. In this work, we focused on the noiseless dataset, as our initial goal is to test the performance of the Bayesian analysis scheme to accurately recover the source parameters. We are currently developing functionalities to jointly estimate the unknown level of noise that affect the source measurements, which we will report about in the future. Each LDC–1 dataset consists of the 1.5-generation TDI observables X, Y, and Z [277, 18] with a cadence of 5 seconds and a duration of 2.5 years. By linearly combining X, Y, and Z, we construct the data, d, for the GW analysis, consisting of the three noise-orthogonal TDI observables A, E and T [265].

The parameters of the injected sources are released as part of the dataset. Figure 5.1 provides a summary of the main features of the signals that were injected. These sources all have GW frequencies of $\sim$ 1–10 mHz at the beginning of the LISA mission. This set of sources covers the chirp mass range $7 - 61\ M_\odot$, see Figs. 5.3, 5.4. The source chirp masses and initial frequencies determine the merger time [see Eq. (5.1)]. Five sources inspiral and merge within $T_{\rm obs} = 2.5$ yr, with chirp masses in the range 30–61$M_\odot$ and initial frequencies between 16 and 20 mHz. Five more, with chirp masses in the range 20–47$M_\odot$ and frequencies between 12 and 22 mHz merge within 5 years. Six other, with chirp masses in the range 13–55$M_\odot$ and frequencies between 8 and 21 mHz merge within 10 years. The longest lived ones, with chirp masses in the range 7–55$M_\odot$ and frequencies between 1 and 11 mHz merge within 3000 years. This set of sources covers a range of SNRs which is governed primarily by the distance to the source, the inclination angle and the source sky position, the latter of which is shown in Fig. 5.5. In total, 22 sources yield an optimal (and coherent across the 3 TDI observables) SNR > 8.

5.3.2 Parameter estimation and results

Preliminarily, we analyzed the same noiseless dataset in distinct runs, where we tuned the priors to a corresponding target SmBBH. We chose priors

- Flat in the dimensionless mass difference $\delta\mu$ in $[0, 0.9]$, corresponding to a mass ratio between $1:1$ and $1:19$.
- Flat in A_L and A_R (i.e. the amplitudes parameters introduced in Sec. 5.2.4) in $[0, A_{\rm max}]$, where $A_{\rm max} = 2\sqrt{2/D_L}$ (twice the overall amplitude of an optimally oriented source at the injected distance).

- Flat in ψ_L and ψ_R (i.e. the initial phases of the left- and right- handed components of the GW signal, respectively) in $[0, \pi]$.

- Flat in the dimensionless spin components of the two binary components, $\chi_{1,\ell}$ and $\chi_{2,\ell}$, in $[-1, 1]$.

- For $\mathcal{M}_c$, f_0, $\sin b$, and l (the chirp mass, the initial orbital frequency, the sine of the ecliptic latitude, and longitude, respectively), we emulated the output of a prior GW search by performing searches on single source simulated data in steps. At each step, we adjust the priors using the posteriors resulting from the previous step to $m \pm 4\sigma$, where m is the median of the posterior, and σ its standard deviation. In order to improve the convergence of the method, when computing m and σ we neglected posterior samples with log-likelihood smaller than one obtained with $A_L = A_R = 0$ (i.e. with $\log \mathcal{L} < -\mathrm{SNR}^2/2$). This method required at most three steps for each target before convergence. Note that this was not possible for all systems, particularly those with low SNR, which we flagged as not detected.

The method we used to determine the search priors produced a set of single-injection runs, where the same waveforms were used for injection and recovery. We found it a useful set of analyses to compare to our main results.

This set of sources also covers a range of SNRs which is governed primarily by the distance to the source, the viewing inclination angle and the source sky position in addition to the chirp mass. With our parameter-estimation pipeline, we were able to obtain good quality posterior distributions, and hence measurements of the source parameters for the 22 sources with SNR > 8. Eight sources with SNRs in the range

5.7–7.9 offered good quality posteriors as well, but we choose to use a fixed SNR threshold and exclude them from the analysis.

For each of the 22 sources that we selected, we computed the 3-dimensional volume within which LISA is able to localize it. Because these sources are generally long-lived, and are at the high-frequency end of the LISA bandwidth, relatively good (by the standards of GW astronomy) sky position measurements with uncertainty regions spanning $\Delta\Omega = 1$ to 100 square degrees are obtained. However, because these sources have relatively low SNRs in the range 8–14 there is a comparatively large fractional uncertainty in the distances spanning 30%–150%. These results are summarized in Figs. 5.5 and 5.6.

Of the intrinsic source parameters, by far the best measured is the chirp mass; for the loudest (quietest) of the recovered sources with SNR 14 (8) we find that we are able to measure the chirp mass to a fractional accuracy better than 0.5% (2%). Our parameter-estimation pipeline sampled directly in the chirp mass $\mathcal{M}_c$ and the dimensionless mass difference $\delta\mu$ as explained in Sec. 5.2.4. The resulting posteriors are shown in Fig. 5.7. The more astrophysically interesting component masses m_1 and m_2 for the individual BHs can be obtained from $\mathcal{M}_c$ and $\delta\mu$; see Fig. 5.3. Notably, we find fractional uncertainties on chirp masses —measured in the frame at rest with the Hubble flow— to be comparable or smaller ($\Delta\mathcal{M}_c^{\rm H}/\mathcal{M}_c^{\rm H} \lesssim 2\times10^{-2}$) to the uncertainties arising from source proper motion redshifts ($v_{\rm pec}/c \lesssim 10^{-2}$). Similarly, the choice of cosmology yields uncertainties in redshift up to 10^{-2} for the most distant source recovered at 500 Mpc, over a broad range of cosmological parameters [281, 282].

Of the other intrinsic parameters, the most interesting are arguably the component spins. While $\chi_{1,\ell}$ and $\chi_{2,\ell}$ cannot be individually measured, it is helpful to identify intrinsic parameters entering the PN frequency evolution series at different orders [256, 283, 270]. As mentioned above, the parameter entering the series at leading order is $\mathcal{M}_c$, the parameter entering at 1PN is $\delta\mu$, while spins first enter at 1.5PN via the combination

$$\beta = \sum_{i=1}^{2} \left(\mu_i + \frac{75\mu_j}{113} \right) \mu_i \chi_{i,\ell}, \tag{5.10}$$

where $i \neq j$, and $\mu_i = m_i/(m_1 + m_2)$ are the dimensionless individual masses. We normalized this parameter so that $|\beta| \leq 1$ for arbitrary mass ratios, implying $|\beta| \leq 94/113$ for equal-mass systems.

The marginal posterior distributions of these three parameters are shown in Fig. 5.4, together with those of the individual masses m_1 and m_2. While the chirp mass is measured extremely well for all sources, $\delta\mu$ and β can be measured with some confidence only for SmBBHs that are merging within the observation window. This is because those sources are the only ones with a sufficient frequency evolution such that the subdominant terms in the PN expansion become observable.

Overall, comparing runs performed on the LDC–1 data with ones performed on single-source injections, we find that parameters are recovered with similar precision. Biases in the LDC–1 runs are comparable to those expected from random noise fluctuations. The exception to that are the two phase parameters ψ_L and ψ_R. These were recovered without any significant bias in the single-source runs, but with large biases comparable to the prior range in the LDC–1 runs in almost all cases. These biases did propagate to both parameters when converted to the (ϕ_0, ψ) plane.

We summarize in Table 5.1 of the appendix the injected parameters of the 22 sources we analyzed, and in Table 5.2 their recovered values.

5.3.3 Challenging systems

Let us now discuss those few systems which showed posteriors that were found to be particularly challenging to analyze. All the SmBBH injections and recoveries done in LDC–1 were performed using noiseless injections. Therefore, in the absence of noise fluctuations, we might expect the likelihood (posterior) to be peaked at (near) the true (i.e. injected) source parameters. However, this is not guaranteed to be the case because (i) we are using different waveforms for recovery than the ones that were injected, and (ii) some sources are overlapping in the LDC–1 data and could therefore be confused.

In particular, we highlight four systems.

- For source number 5 (SNR = 11.36), we obtain a frequency posterior that is peaked significantly away from the injected values.

- Sources number 20 (SNR = 8.68), and 36 (SNR = 9.93) resulted in a 2-dimensional posterior on the chirp mass and mass difference parameters (or equivalently on the component masses) that only include the injected values on the boundary of their $\sim 99.8\%$ confidence interval.

- For source number 16 (SNR = 10.14) the marginalized, 1-dimensional posterior on $\mathcal{M}_c$ includes the injected value only in its 99.4% confidence interval.

We note that for the first two bullet points listed above, the issues described are not present in the single-injection run results used for comparison. These differences could possibly be due to the difference in the employed waveforms, the signal overlap in LDC–1, sampling issues, or a combination of these. Work toward analyzing jointly the overlapping sources [1] and characterizing performances of different samplers is ongoing.

On the other hand, the bias in $\mathcal{M}_c$ observed for source 16 was also present in the single-injection result. In the following we argue that it's a genuine effect of the signal parametrization. Figure 5.8 shows the marginalized posteriors in the $(\mathcal{M}_c, \delta\mu)$ plane for source number 16 (SNR $= 10.14$, $\tau = 7.9$ yr), together with a reparametrization of it. While the true parameters still lie in the main confidence region of the two-dimensional posterior, the chirp mass posterior only includes the injected value in the tail of the distribution. Notably, its injected mass ratio of $q \approx 1/11.3$ ($\delta\mu \approx 0.84$) is the most asymmetric among all detected sources. The flat posterior in $\delta\mu$ that we observe in Fig. 5.8 suggests that this parameter is not measurable. The flatness of this posterior together with the shape of the confidence region implies a bias in the marginalized posterior for $\mathcal{M}_c$ for highly asymmetric mass ratios. The shape of the two-dimensional posterior can be explained by an examination of the PN GW phase series [256]:

$$\begin{aligned}\Phi = \Phi_c - \frac{(\pi\mathcal{M}_c f)^{-5/3}}{16}\Bigg[1 \\ + \frac{(\pi\mathcal{M}_c f)^{2/3}}{2^{1/5}\left(1-\delta\mu^2\right)^{2/5}}\left(\frac{2435}{252} - \frac{55\delta\mu^2}{24}\right)\Bigg] + \mathcal{O}\left(f^{-2/3}\right).\end{aligned} \tag{5.11}$$

As $\delta\mu$ increases, the resulting change in the number of accumulated cycles can be compensated by an increase in $\mathcal{M}_c$. This behavior is all the more pronounced that the system is observed closer to merger, as the strength of the 1PN term gets more comparable to the 0PN one. To reduce the correlation in the $(\mathcal{M}_c, \delta\mu)$ plane, we can define a new parameter $\mathcal{M}_\phi(\mathcal{M}_c, \delta\mu, f_0, T_{\rm obs})$, such that the number of accumulated cycles during the observation is independent of $\delta\mu$ up to some given PN order. Note that the likelihood is not exclusively determined by the number of accumulated cycles of phase, hence one should not expect $\mathcal{M}_\phi$ and $\delta\mu$ to be completely uncorrelated. At

1PN order, we get

$$\mathcal{M}_\phi = \mathcal{M}_c \left\{ 1 - \frac{(5\mathcal{M}_c)^{1/4}\left[974(1-A) - 231\delta\mu^2\right]}{168 \times 2^{1/5} A} \right. \tag{5.12}$$

$$\left. \times \frac{\left(\tau_0^{3/8} - \tau_{\rm f}^{3/8}\right)^2}{5\tau_0 - 8\tau_0^{5/8}\tau_{\rm f}^{3/8} + 3\tau_{\rm f}} \right\},$$

$$A = \left(1 - \delta\mu^2\right)^{2/5}, \tag{5.13}$$

$$\tau_0 = \frac{5(\pi\mathcal{M}_c f_0)^{-5/3}}{256\pi f_0}, \tag{5.14}$$

$$\tau_{\rm f} = \max\left[\tau_0 - T_{\rm obs}, \frac{5(\pi\mathcal{M}_c f_{\rm max})^{-5/3}}{256\pi f_{\rm max}}\right], \tag{5.15}$$

where $f_0 = 2f_0^{\rm orb}$ is the initial GW frequency, and $f_{\rm max}$ is the higher limit of the observation frequency band. As shown in Fig. 5.8, the 2-dimensional posterior in the masses plane yields a milder correlation, and hence a smaller bias in the marginalized, 1-dimensional posterior for $\mathcal{M}_\phi$.

5.3.4 Eccentric precessing system

We also ran as a proof of concept a Bayesian parameter estimation run on a fully general eccentric precessing system. We chose a 95-55$M_\odot$ binary system, with spin magnitudes $\chi_1 = 0.7$ and $\chi_2 = 0.73$ respectively, initial spin misalignment angles $\theta_1 = 179°$ and $\theta_2 = 135°$ respectively, eccentricity at 10 mHz of $e_{10} = 3.1 \times 10^{-3}$, and SNR 15. These values were inspired by the most massive event detected by LIGO/Virgo to date, GW190521 [284]. This particular source accumulated $\mathcal{N} \approx 1.89 \times 10^6$ cycles of orbital phase, $\mathcal{N}_{\rm spin} \approx 892$ cycles of spin precession, and $\mathcal{N}_{\rm ecc} \approx 4060$ cycles of periastron precession.

For this run, we used the same sampling parameters as for the LDC–1 runs with a few modifications. We used different spin parameters, we added eccentricity parameters, and we replaced the initial orbital frequency with the approximate merger time parameter [29]

$$t_M = t_0 + \frac{5\mathcal{M}_c(\pi\mathcal{M}_c f_0)^{-8/3}}{32\sqrt{1-e_{10}^2}\,(8+7e_{10}^2)}, \tag{5.16}$$

where t_0 is the time at the start of data gathering. Note that, for simplicity, this relation is obtained from the leading PN order frequency evolution equation, assuming a constant eccentricity. It is more accurate for circular systems, and becomes gradually less so as the initial eccentricity increases.

We note that since spin-induced precession causes $\cos\iota$ and ψ to evolve with time, we use their initial values to define the parameters A_L, A_R, ψ_L, and ψ_R. Additionally, we set the priors on ψ_L and ψ_R as $[0, 2\pi]$ since the addition of eccentricity breaks the waveform symmetry from $(\phi_0 \to \phi_0 + n\pi)$ to $(\phi_0 \to \phi_0 + 2n\pi)$, $n \in \mathbb{Z}$.

Of note, we report on the measurability of a few chosen parameters, shown in Fig. 5.9. The merger time t_M could be recovered with accuracy ~ 2 hours, a figure comparable to the corresponding merging circular sources. The chirp mass could be recovered with accuracy of $\sim 0.004 M_\odot$. The eccentricity at 10 mHz could be recovered in the range $2.7\times10^{-3} < e_{10} < 4.9\times10^{-3}$ at 90% confidence, and was clearly distinguishable from zero. The injected dimensionless mass difference was recovered within the 24% confidence interval. The (initial) spin parameter β could be recovered with 90% confidence in the range $-0.41 < \beta < -0.14$, including the injected value of $\beta_{\rm inj} \approx -0.37$, (with $\beta = 0$ excluded at more than 99.9% confidence). The source is located on the sky at 90% confidence level within 8 deg^2. The recovered values were mostly consistent with the injected ones. More work is ongoing to assess the robustness of the parameter estimation pipeline across the full parameter space.

One interesting additional information to gather from these results is to determine whether the effects of spin-precession are measurable for such systems. In order to do this, we looked at the average precession parameter χ_p [285, 286], and found that the posterior did not differ significantly from the prior, suggesting that precession effects might not be measurable for SmBBHs with LISA, thus strengthening the case for multiband GW astronomy.

These Bayesian results obtained for a fully precessing eccentric binary show promise for an extension of the present work, investigating the full 17-dimensional parameter space of SmBBHs more extensively.

5.3.5 Computational performances

Parameter estimation runs were carried out on the high performance computing infrastructure provided by the Birmingham BlueBEAR cluster, with each run using 8 Intel Xeon (2.50GHz) sibling cores on a single computing node. The total CPU time for each run on the LDC–1 dataset was distributed with a median of 5 hours for the 22 sources with SNR > 8.

The three runs with sources coalescing within the dataset duration where the most computationally demanding with CPU times of 36, 20, and 45 hours for sources number 20, 36, and 47 respectively. All runs had small memory footprint throughout, with usage peaks below 1.6 Gigabytes.

The run with eccentricity and spin precession was substantially more expensive (~ 3000 CPUh), approximately 100 times more than its merging spin-aligned circular counterparts. This is due to a combination of factors including the increased dimensionality of the parameter space, the additional structure of the likelihood, and the increased complexity of the waveforms.

5.4 Conclusions

In this paper we presented a fully Bayesian parameter-estimation routine for the observation of SmBBHs with LISA. As part of the LISA data challenge LDC–1, we employed our codebase BALROG for the accurate estimation of 66 circular, spin-aligned SmBBHs' parameters. We confidently recovered all 22 sources with SNR > 8. Our results show that LISA will be able to localize SmBBHs over the sky within a few tens of squared degrees, and constrain their detector-frame chirp mass down to $\pm 0.01\ M_{\odot}$. Additionally, for sources merging within the mission lifetime, the chirping morphology of the signals allows us to measure parameters entering at higher order in the post-Netwonian expansion, namely the dimensionless mass difference $\delta\mu$, and the spin combination β.

On the technical side, we presented a novel choice of the sampling parameters that substantially reduce the correlations in the high-dimensional likelihood, thus vastly increasing the resulting computational efficiency. In particular, this relies on decomposing the signal into circular polarizations. We also presented an algorithm that drastically reduces the number of waveform evaluations needed to estimate likelihoods, by adapting a nonuniform quadrature rule to work with uniformly sampled data. This allowed us to successfully perform full Bayesian parameter estimation studies for individual spin-aligned, circular SmBBH sources undergoing $\mathcal{O}(10^6)$ wave cycles that required just a few CPU hours to complete. Focusing on a selected number of sources that exhibit mild biases in the recovered parameters, we characterized the effect of binaries cross-contamination, waveform differences, and inherent likelihood structures which make SmBBHs parameters challenging to sample.

Finally, we presented a proof-of-concept analysis where we tackle the full SmBBH parameter-estimation problem, which includes eccentricity and spin precession. We recover the parameters of a specific, but generic, source in the resulting 17-dimensional parameter space. We report a measurable eccentricity at 10 mHz of a few 10^{-3} together with a merger time determination within a time window of $\lesssim 1$ hour. We also report the immeasurability of spin-induced precession effects, suggesting that individual component spins cannot be recovered. This suggests that joint space and ground based detector GW observations might be crucial to fully characterize SmBBHs. More work is necessary to fully explore potential challenges for this type of sources. This analysis brings us closer to LISA's goal of efficiently and accurately reconstructing the parameters of SmBBHs, which constitute an unmatched tool to discriminate their formation history and evolution.

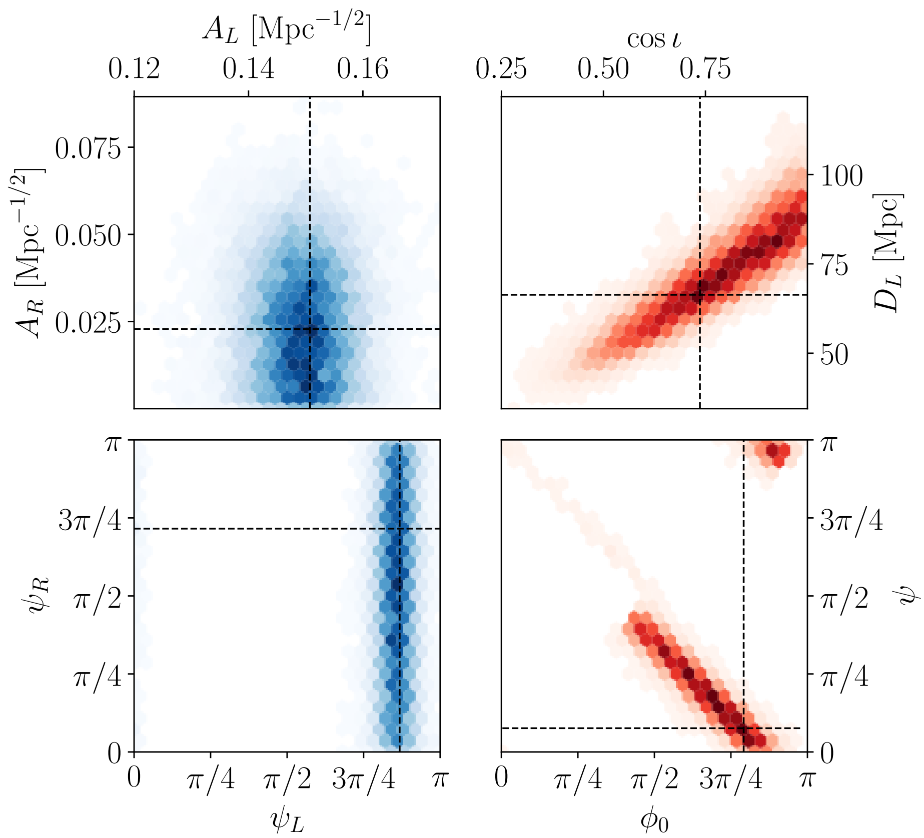

FIGURE 5.2: Two-dimensional posteriors for LDC–1 source 15 (SNR 12) using different parameters. Quantities related to the circular (linear) polarizations are indicated in blue (red). The amplitudes A_L and A_R (top left) are related to D_L and $\cos\iota$ (top right). The phases ψ_L and ψ_R (bottom left) are related to ϕ_0 and ψ (bottom right). The posteriors have been weighted in each plot so that the parameters shown in it are flat distributed. Quantities on the left (blue) are significantly less correlated than those on the right (red).

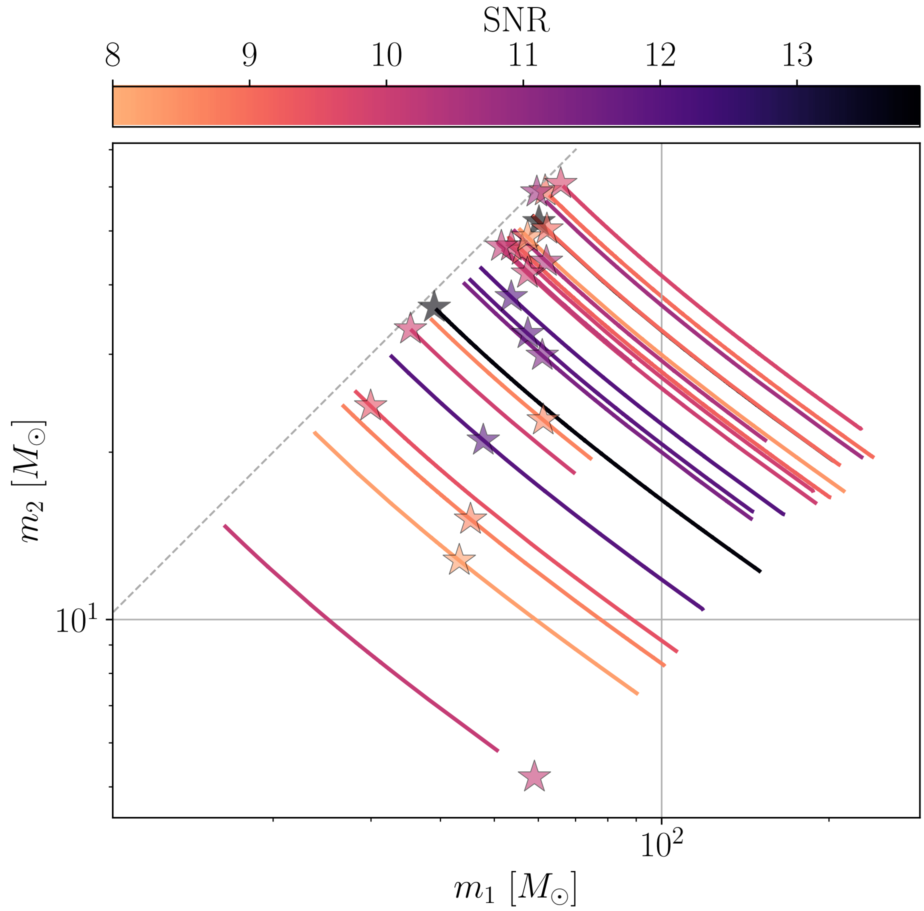

FIGURE 5.3: Posterior samples of detector frame component masses for the 22 recovered sources. Solid lines extend in the 90% confidence interval of the symmetric mass ratio posterior. The thickness of the curve is comparable to or greater than the posterior distribution widths, indicating the very high accuracy of the chirp-mass measurements. Injected values are marked by stars. Lines and markers are colored according to the sources' SNRs. All injected values lie within their posterior's 90% contour levels, except for source 16 (SNR = 10) whose true, high dimensionless mass difference is within the 98% confidence interval (cf. Sec. 5.3.3).

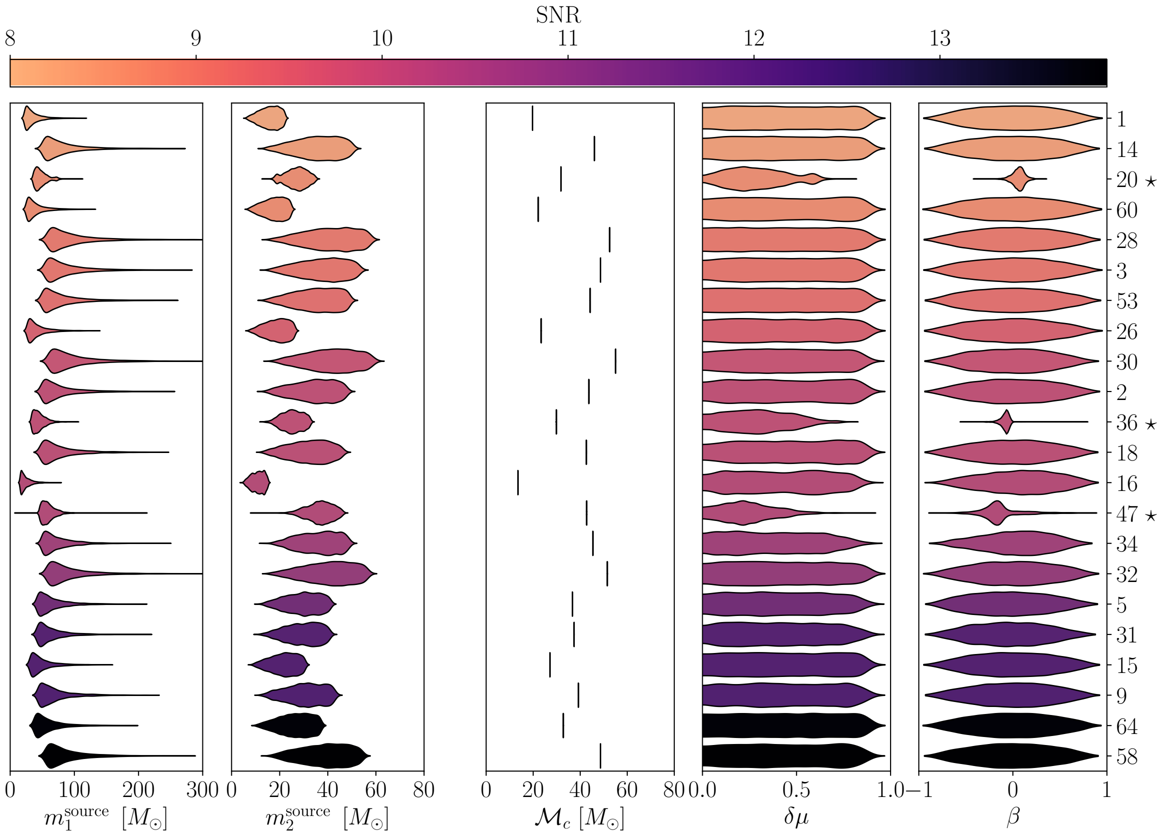

FIGURE 5.4: Marginal posteriors, represented through kernel density estimators for five selected parameters. From left to right, we show the primary component mass, the secondary component mass (both measured in the source frame), the redshifted chirp mass, the dimensionless mass difference, and the 1.5 PN spin parameter reported in Eq. (5.10). Source posteriors are sorted and colored by their SNRs; their index in the LDC-1 injections catalogue is reported to the right. Stars mark sources that merge within the LDC–1 dataset duration (2.5 yr). Posteriors are reweighted to an effective prior uniform in the column's parameter, except for the spin parameter β. The redshifted chirp mass, appearing in the leading-order PN term of the frequency evolution, is much better constrained than any of the other parameters. Parameters entering at higher PN order like $\delta\mu$ and β can only be constrained for systems that merge within the mission lifetime.

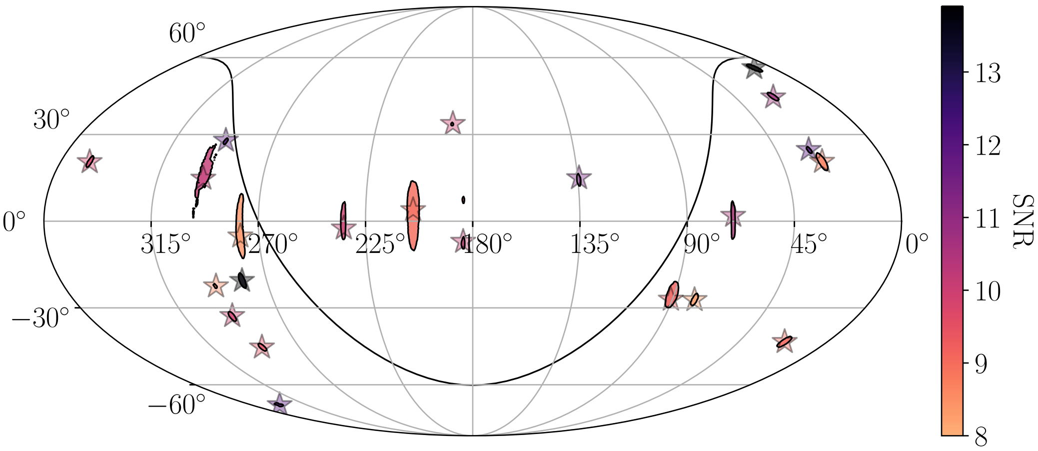

Figure 5.5: Posteriors on the sky position of the recovered sources in ecliptic coordinates using a Mollweide projection. Stars denote the true locations of the injected sources. All true locations are enclosed within the 90% confidence intervals of their posterior. The solid black line shows the galactic plane. Note that sources close to the ecliptic have an approximate symmetry involving the ecliptic latitude $b \to -b$ (see e.g. [280]), resulting in elongated posteriors in that region, and even a bimodal posterior as seen in the source close to 180° longitude.

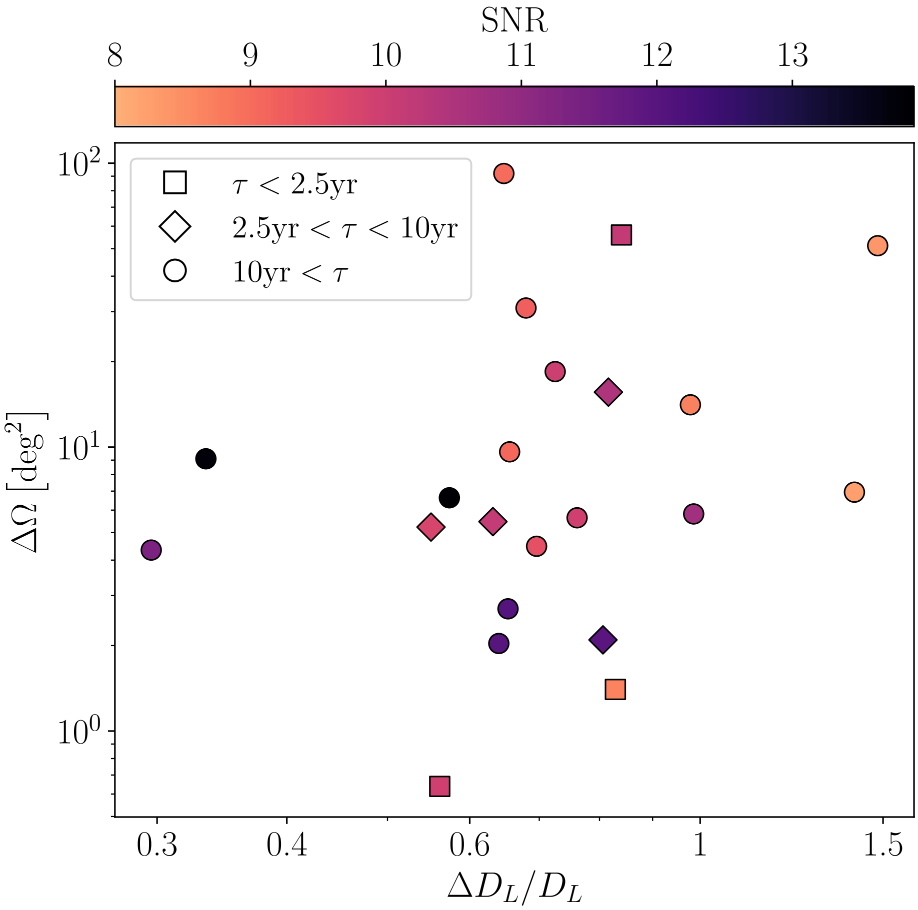

FIGURE 5.6: Marginalized posterior uncertainties on the distance and sky location of the 22 recovered sources. We show the 90% confidence intervals of the two–dimensional sky location posteriors and the uncertainties on the distance relative to its true injected value. All sources above SNR > 8 are resolved with angular resolution better than 100 deg^2. At SNR above 11 the localization improves by an order of magnitude and the distance is measured with 70% accuracy or better. Markers label the source time to coalescence as shown in legend and described previously in Fig. 5.1.

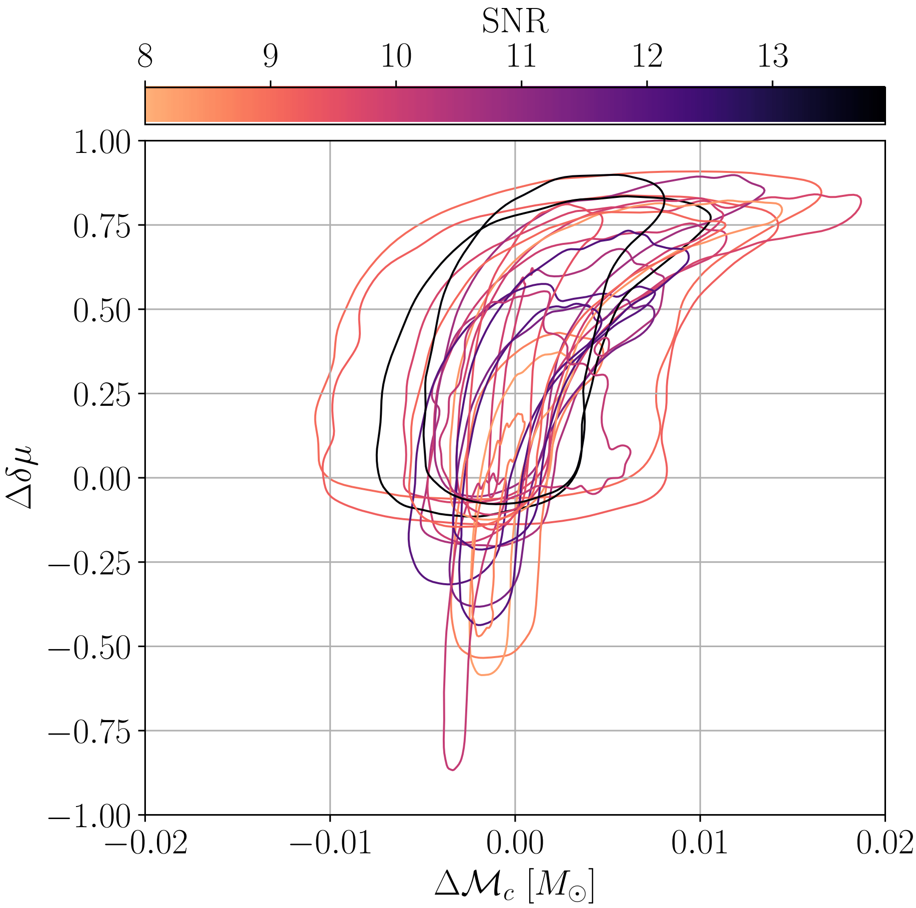

FIGURE 5.7: 90% marginalized posterior contours in the chirp mass-dimensionless mass difference plane for all 22 observed sources. The colors of the contours are set by the SNR of the corresponding source. All contours are offset so that the injected value of the corresponding source lies at the origin.

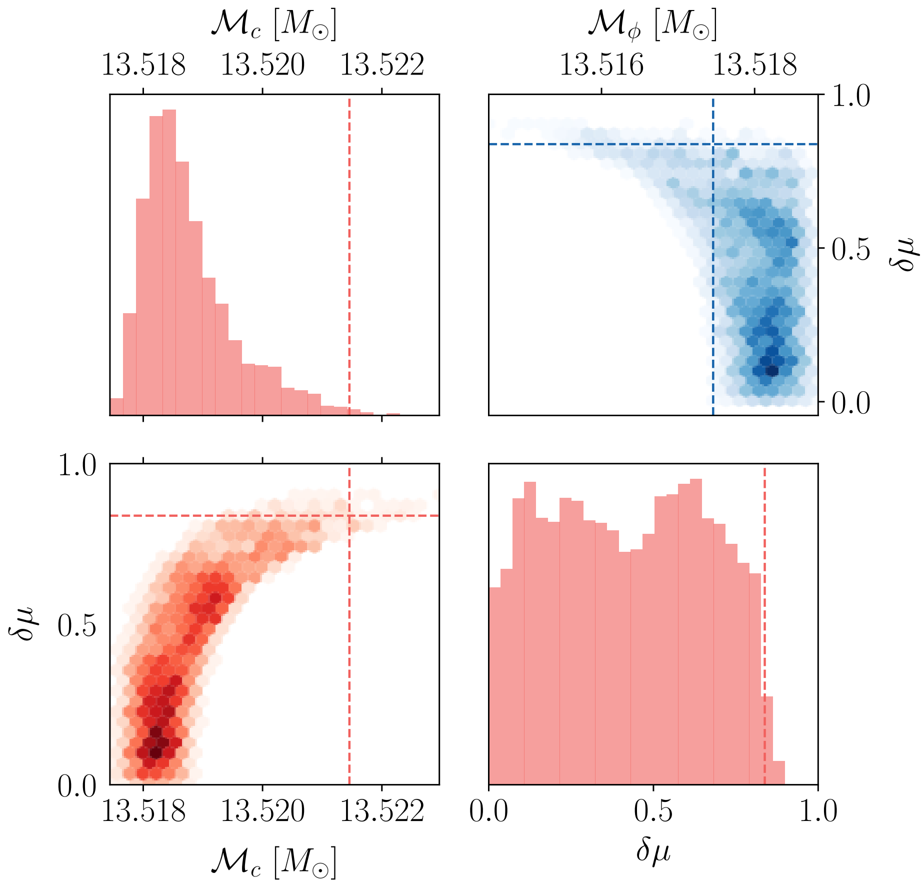

FIGURE 5.8: Marginalized posterior distributions in the masses plane for source number 16. The two-dimensional posterior for $(\mathcal{M}_c, \delta\mu)$ (red-shaded, lower left plot) shows support in the region containing the true injected value (dashed red line). However, the correlation structure in these parameters together with the fact that $\delta\mu$ is not measurable generate a strong bias in $\mathcal{M}_c$ when marginalized over $\delta\mu$, as shown in the upper left and lower right histograms. The corresponding plot in the single-source injection run showed a similar pattern. By re-parameterizing the masses plane with $(\mathcal{M}_\phi, \delta\mu)$ (blue shaded, upper right plot), we observe a milder correlation, hence a smaller bias on the posterior marginal distribution for $\mathcal{M}_\phi$ compared to $\mathcal{M}_c$.

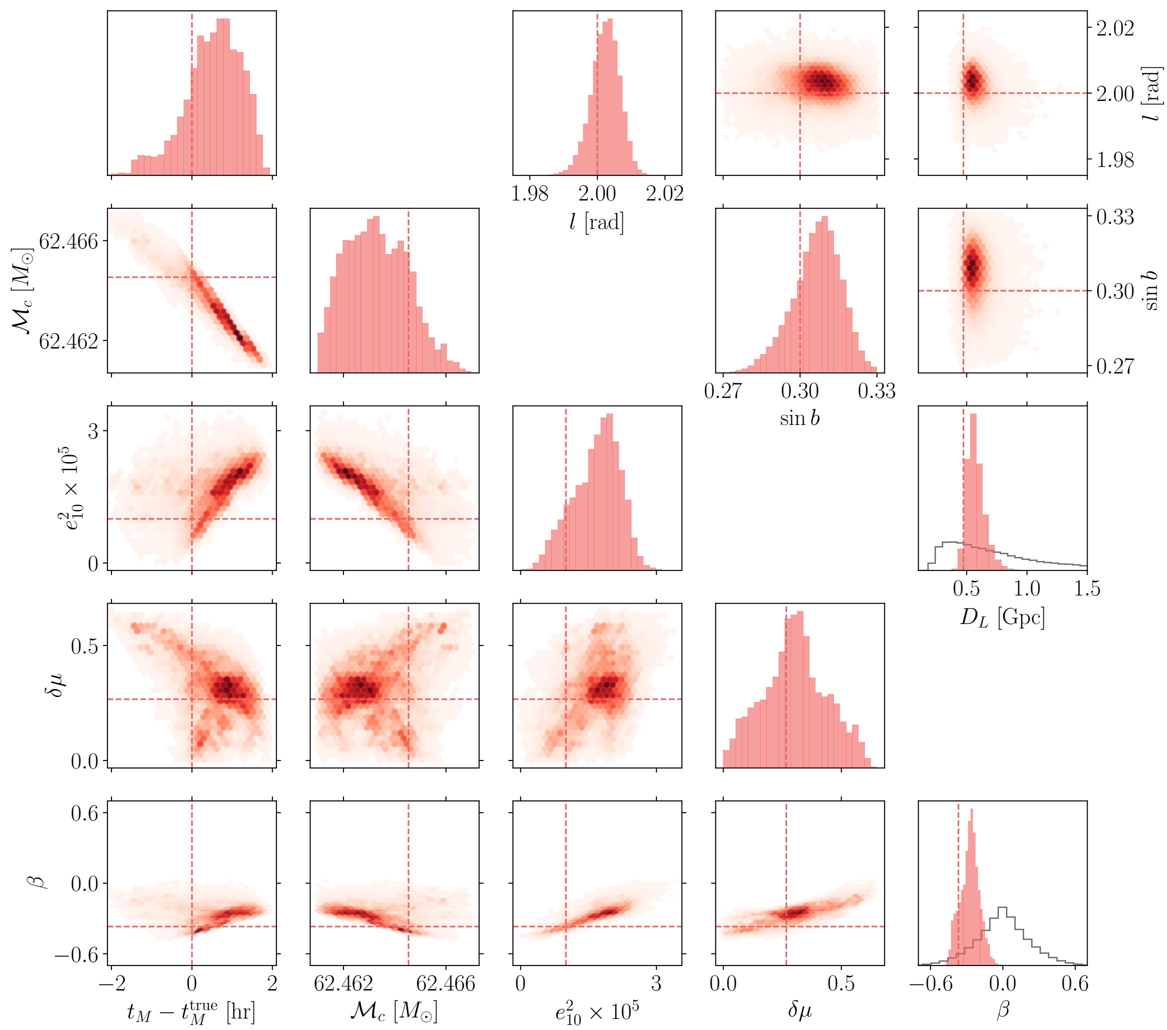

FIGURE 5.9: Posterior distributions on a selection of parameters of the eccentric precessing system. Priors are uniform for (upper right panel) the ecliptic longitude l, the sine of ecliptic latitude $\sin b$, the amplitude parameters A_L and A_R. Consequently, the luminosity distance prior is $\propto 1/D_L^2$, and shown as solid black line. Equally, uniform priors are used for (lower left panel) the merger time t_M, the chirp mass $\mathcal{M}_c$, the square eccentricity e_{10}^2, the dimensionless mass difference $\delta\mu$, the dimensionless spin magnitudes, and the spin unit vectors on the sphere. The resulting prior is shown for the spin parameter β as a solid black line. The merger time can be measured within approximately an hour, and the eccentricity and spin parameter can be well distinguished from zero. The source is localized in the sky within 8 deg^2.

5.5 Appendix: LDC–1 injected and recovered parameters

In this Appendix, we provide the parameters of our LDC–1 analysis in tabular format. In particular, in Table 5.1 we list the parameters of all the injected SmBBHs, while in Table 5.2 we present the results of our parameter-estimation recovery.

SNR	τ [yr]	f_0 [mHz]	l [rad]	$\sin b$	D_L [Mpc]	$\cos\iota$	m_1 [$M_\odot$]	m_2 [$M_\odot$]	$\mathcal{M}_c$ [$M_\odot$]	$\delta\mu$	β	ID
8.26	45.7	8.6020	1.40	-0.46	55.1	0.39	43.2	12.8	19.756	0.54	0.30	1
8.40	27.5	6.1377	4.85	-0.09	147.8	-0.28	57.4	48.7	46.016	0.08	-0.18	14
8.69	2.3	19.7430	5.12	-0.38	193.7	0.52	61.2	22.8	31.808	0.46	0.00	20⋆
8.70	84.3	6.3652	0.48	0.35	68.0	0.62	45.3	15.1	22.158	0.50	-0.07	60
9.04	89.5	3.6289	3.58	0.07	237.2	0.99	61.7	59.0	52.539	0.02	0.13	28
9.07	40.2	5.1370	0.38	-0.68	238.7	0.84	62.2	50.4	48.659	0.10	0.18	3
9.21	108.5	3.7577	1.58	-0.45	180.8	-0.98	56.1	46.1	44.244	0.10	0.21	53
9.48	51.7	7.3862	5.05	-0.70	87.8	0.88	30.0	24.2	23.418	0.11	0.13	26
9.81	8.5	8.5140	6.06	0.35	493.4	-0.96	65.8	60.8	55.072	0.04	0.00	30
9.92	23.9	6.6834	5.10	-0.54	183.1	0.69	53.6	47.0	43.669	0.07	-0.14	2
9.93	2.3	20.4730	3.30	0.56	263.1	-0.98	35.3	33.3	29.853	0.03	-0.03	36⋆
10.00	11.6	8.8876	4.09	-0.04	247.3	-0.73	57.4	42.0	42.599	0.16	0.13	18
10.14	7.9	21.0170	3.21	-0.12	77.4	-0.97	59.0	5.2	13.521	0.84	-0.04	16
10.15	1.8	17.8620	5.16	0.25	390.5	-0.91	51.5	46.9	42.763	0.05	-0.19	47⋆
10.48	5.7	11.1590	1.23	0.03	285.0	-0.66	62.0	44.1	45.386	0.17	0.01	34
10.71	30.3	5.5147	0.44	0.69	151.3	-0.47	59.6	58.7	51.511	0.01	-0.08	32
11.36	14.4	9.0172	2.35	0.25	79.9	0.05	61.0	29.8	36.680	0.34	-0.07	5
11.96	8.0	11.1150	5.09	0.47	176.4	0.65	57.4	32.7	37.418	0.27	0.40	31
12.00	53.4	6.6569	5.88	-0.94	66.3	-0.74	47.8	21.0	27.128	0.39	0.16	15
12.07	11.6	9.3873	0.53	0.42	191.7	0.76	53.6	38.0	39.212	0.17	-0.15	9
13.77	106.6	4.5605	4.90	-0.35	33.9	-0.16	39.0	36.4	32.804	0.03	-0.44	64
13.90	93.3	3.7494	0.19	0.82	106.6	-0.84	60.1	51.9	48.620	0.07	0.00	58

TABLE 5.1: Properties of the LDC–1 injected sources. Rows are ordered by increasing source SNR and labelled by the injection ID in the LDC dataset. Sources merging within the mission lifetime (here set to 2.5 yr) are marked with stars. For a description of the parameters see Sec. 5.2.4.

CPUh	$\tau - \tau^{\rm inj}$ [day]	$f_0 - f_0^{\rm inj}$ [nHz]	$\Delta\Omega$ [deg^2]	l [rad]	$\sin b$	D_L [Mpc]	$\cos\iota$	m_1 [$M_\odot$]	m_2 [$M_\odot$]	$\mathcal{M}_c$ [$M_\odot$]	$\delta\mu$	β	ID
3.7	$1.28^{+1.56}_{-4.11}$	$-1.0^{+6.2}_{-6.0}$	6.9	$1.4^{+0.02}_{-0.02}$	$-0.455^{+0.026}_{-0.024}$	$77.6^{+49.1}_{-28.5}$	$0.57^{+0.38}_{-0.3}$	$37.2^{+53.0}_{-13.4}$	$14.5^{+7.2}_{-7.1}$	$19.755^{+0.003}_{-0.001}$	$0.44^{+0.41}_{-0.39}$	$0.0^{+0.7}_{-0.7}$	1
5.1	$-0.17^{+1.03}_{-3.45}$	$-1.7^{+6.2}_{-6.1}$	51.2	$4.85^{+0.02}_{-0.02}$	$-0.067^{+0.19}_{-0.114}$	$182.3^{+166.5}_{-52.6}$	$-0.36^{+0.21}_{-0.54}$	$87.2^{+125.2}_{-31.6}$	$33.5^{+16.8}_{-16.4}$	$46.017^{+0.009}_{-0.003}$	$0.45^{+0.41}_{-0.39}$	$0.0^{+0.6}_{-0.7}$	14
35.9	$0.07^{+0.02}_{-0.06}$	$-6.1^{+17.0}_{-21.4}$	1.4	$5.12^{+0.01}_{-0.01}$	$-0.378^{+0.012}_{-0.013}$	$166.3^{+105.5}_{-55.0}$	$0.52^{+0.42}_{-0.28}$	$48.7^{+25.8}_{-10.1}$	$27.9^{+6.8}_{-8.4}$	$31.806^{+0.002}_{-0.001}$	$0.27^{+0.31}_{-0.22}$	$0.1^{+0.1}_{-0.1}$	20⋆
4.0	$1.01^{+5.38}_{-7.41}$	$-0.2^{+6.3}_{-6.4}$	14.1	$0.48^{+0.02}_{-0.02}$	$0.349^{+0.036}_{-0.04}$	$79.4^{+39.4}_{-27.2}$	$0.65^{+0.3}_{-0.29}$	$41.9^{+59.0}_{-15.1}$	$16.2^{+8.0}_{-7.9}$	$22.158^{+0.003}_{-0.002}$	$0.44^{+0.41}_{-0.39}$	$0.0^{+0.7}_{-0.7}$	60
2.5	$-0.4^{+9.1}_{-10.69}$	$-0.3^{+5.4}_{-5.4}$	92.0	$3.58^{+0.03}_{-0.03}$	$0.056^{+0.136}_{-0.179}$	$189.0^{+90.9}_{-62.7}$	$0.66^{+0.29}_{-0.29}$	$99.4^{+140.1}_{-36.3}$	$38.3^{+19.4}_{-18.6}$	$52.54^{+0.01}_{-0.009}$	$0.44^{+0.4}_{-0.4}$	$0.0^{+0.7}_{-0.7}$	28
3.3	$0.44^{+1.72}_{-4.14}$	$-0.1^{+4.8}_{-4.9}$	9.6	$0.38^{+0.03}_{-0.03}$	$-0.676^{+0.02}_{-0.019}$	$203.8^{+91.0}_{-65.5}$	$0.68^{+0.28}_{-0.28}$	$91.7^{+116.7}_{-32.8}$	$35.6^{+17.5}_{-16.6}$	$48.658^{+0.008}_{-0.003}$	$0.44^{+0.39}_{-0.39}$	$0.0^{+0.7}_{-0.7}$	3
2.7	$-0.76^{+12.77}_{-13.73}$	$-0.5^{+5.2}_{-5.1}$	30.9	$1.59^{+0.03}_{-0.03}$	$-0.43^{+0.057}_{-0.053}$	$121.7^{+81.0}_{-41.9}$	$-0.52^{+0.27}_{-0.42}$	$83.6^{+116.7}_{-30.4}$	$32.3^{+16.3}_{-15.7}$	$44.245^{+0.009}_{-0.009}$	$0.44^{+0.4}_{-0.4}$	$0.0^{+0.7}_{-0.7}$	53
3.2	$0.47^{+2.22}_{-4.34}$	$-0.0^{+6.7}_{-6.9}$	4.5	$5.05^{+0.02}_{-0.02}$	$-0.702^{+0.014}_{-0.013}$	$78.3^{+35.3}_{-25.9}$	$0.68^{+0.28}_{-0.29}$	$44.2^{+61.9}_{-16.0}$	$17.1^{+8.6}_{-8.3}$	$23.418^{+0.003}_{-0.002}$	$0.44^{+0.4}_{-0.4}$	$0.0^{+0.7}_{-0.7}$	26
7.0	$0.01^{+0.42}_{-1.27}$	$-1.1^{+5.1}_{-6.0}$	5.2	$6.06^{+0.01}_{-0.01}$	$0.349^{+0.026}_{-0.027}$	$325.4^{+159.2}_{-112.5}$	$-0.63^{+0.3}_{-0.33}$	$103.5^{+124.5}_{-36.7}$	$40.3^{+19.6}_{-18.3}$	$55.072^{+0.014}_{-0.004}$	$0.44^{+0.38}_{-0.39}$	$0.0^{+0.6}_{-0.6}$	30
5.0	$-0.19^{+0.89}_{-2.67}$	$-3.4^{+6.4}_{-6.6}$	5.6	$5.11^{+0.02}_{-0.02}$	$-0.546^{+0.02}_{-0.019}$	$188.9^{+79.9}_{-59.4}$	$0.69^{+0.27}_{-0.28}$	$82.0^{+105.1}_{-29.1}$	$32.0^{+15.6}_{-15.0}$	$43.669^{+0.008}_{-0.003}$	$0.44^{+0.4}_{-0.39}$	$0.0^{+0.7}_{-0.6}$	2
19.8	$0.04^{+0.02}_{-0.06}$	$-1.0^{+6.7}_{-7.9}$	0.6	$3.31^{+0.01}_{-0.01}$	$0.555^{+0.006}_{-0.006}$	$187.3^{+86.9}_{-60.8}$	$-0.65^{+0.29}_{-0.3}$	$46.4^{+23.0}_{-10.9}$	$25.8^{+7.3}_{-7.4}$	$29.852^{+0.001}_{-0.001}$	$0.29^{+0.29}_{-0.25}$	$-0.1^{+0.1}_{-0.1}$	36⋆
5.2	$0.09^{+0.49}_{-1.55}$	$0.4^{+4.7}_{-4.9}$	18.4	$4.09^{+0.01}_{-0.01}$	$-0.001^{+0.092}_{-0.088}$	$242.0^{+100.0}_{-79.4}$	$-0.69^{+0.3}_{-0.27}$	$82.7^{+106.2}_{-30.8}$	$30.4^{+15.7}_{-14.1}$	$42.598^{+0.009}_{-0.003}$	$0.46^{+0.38}_{-0.4}$	$0.0^{+0.6}_{-0.7}$	18
14.2	$1.03^{+0.26}_{-0.64}$	$6.8^{+6.0}_{-7.7}$	5.5	$3.21^{+0.01}_{-0.01}$	$-0.122^{+0.259}_{-0.033}$	$62.5^{+28.2}_{-20.7}$	$-0.67^{+0.29}_{-0.29}$	$24.8^{+25.7}_{-8.4}$	$10.1^{+4.6}_{-4.3}$	$13.519^{+0.002}_{-0.001}$	$0.42^{+0.37}_{-0.37}$	$0.1^{+0.6}_{-0.7}$	16
45.2	$0.02^{+0.05}_{-0.13}$	$-8.0^{+66.5}_{-96.4}$	55.9	$5.17^{+0.02}_{-0.04}$	$0.286^{+0.116}_{-0.206}$	$305.5^{+215.2}_{-112.8}$	$-0.55^{+0.36}_{-0.4}$	$62.2^{+25.5}_{-11.4}$	$39.2^{+8.3}_{-10.0}$	$42.762^{+0.006}_{-0.002}$	$0.23^{+0.27}_{-0.19}$	$-0.2^{+0.5}_{-0.3}$	47⋆
8.7	$0.0^{+0.29}_{-0.48}$	$-2.5^{+8.6}_{-8.7}$	15.6	$1.23^{+0.01}_{-0.01}$	$0.009^{+0.089}_{-0.099}$	$287.0^{+140.1}_{-92.5}$	$-0.64^{+0.28}_{-0.33}$	$78.6^{+75.0}_{-24.2}$	$35.7^{+14.3}_{-14.6}$	$45.385^{+0.006}_{-0.004}$	$0.38^{+0.38}_{-0.33}$	$0.1^{+0.5}_{-0.6}$	34
3.0	$-0.01^{+1.07}_{-3.59}$	$-0.7^{+4.6}_{-4.5}$	5.8	$0.44^{+0.02}_{-0.02}$	$0.695^{+0.015}_{-0.016}$	$181.1^{+91.1}_{-58.1}$	$-0.61^{+0.27}_{-0.34}$	$95.6^{+133.2}_{-33.4}$	$38.2^{+18.2}_{-18.6}$	$51.511^{+0.01}_{-0.003}$	$0.43^{+0.41}_{-0.38}$	$0.0^{+0.6}_{-0.7}$	32
10.1	$0.15^{+0.49}_{-1.46}$	$-12.5^{+4.9}_{-3.6}$	4.3	$2.35^{+0.01}_{-0.01}$	$0.244^{+0.027}_{-0.029}$	$78.3^{+13.9}_{-9.8}$	$0.06^{+0.08}_{-0.08}$	$66.6^{+78.1}_{-22.4}$	$27.7^{+12.5}_{-12.5}$	$36.68^{+0.006}_{-0.002}$	$0.41^{+0.4}_{-0.37}$	$0.0^{+0.7}_{-0.6}$	5
5.9	$0.26^{+0.32}_{-0.89}$	$1.6^{+5.1}_{-5.7}$	2.1	$5.09^{+0.01}_{-0.01}$	$0.462^{+0.013}_{-0.014}$	$192.4^{+81.6}_{-60.7}$	$0.67^{+0.29}_{-0.29}$	$67.9^{+77.6}_{-22.6}$	$28.2^{+12.6}_{-12.6}$	$37.416^{+0.007}_{-0.003}$	$0.41^{+0.39}_{-0.36}$	$0.0^{+0.6}_{-0.6}$	31
2.2	$0.97^{+1.89}_{-4.62}$	$0.0^{+4.1}_{-4.2}$	2.0	$5.88^{+0.03}_{-0.03}$	$-0.936^{+0.004}_{-0.004}$	$66.4^{+23.4}_{-19.0}$	$-0.72^{+0.26}_{-0.24}$	$50.9^{+67.5}_{-18.2}$	$19.9^{+9.8}_{-9.5}$	$27.128^{+0.004}_{-0.002}$	$0.44^{+0.4}_{-0.39}$	$0.0^{+0.7}_{-0.7}$	15
5.3	$-0.05^{+0.46}_{-1.24}$	$0.2^{+4.0}_{-4.1}$	2.7	$0.53^{+0.01}_{-0.01}$	$0.414^{+0.015}_{-0.015}$	$182.9^{+68.1}_{-57.1}$	$0.7^{+0.26}_{-0.28}$	$73.1^{+92.3}_{-25.7}$	$28.9^{+13.9}_{-13.5}$	$39.212^{+0.007}_{-0.003}$	$0.43^{+0.4}_{-0.38}$	$0.0^{+0.6}_{-0.7}$	9
4.7	$-0.82^{+8.24}_{-10.05}$	$-0.1^{+4.2}_{-4.3}$	9.1	$4.9^{+0.02}_{-0.02}$	$-0.349^{+0.034}_{-0.032}$	$36.9^{+6.6}_{-4.7}$	$-0.16^{+0.08}_{-0.09}$	$62.6^{+87.3}_{-23.2}$	$23.7^{+12.3}_{-11.5}$	$32.804^{+0.005}_{-0.004}$	$0.45^{+0.4}_{-0.4}$	$0.0^{+0.7}_{-0.7}$	64
2.4	$0.49^{+6.31}_{-8.35}$	$0.3^{+3.6}_{-3.6}$	6.6	$0.2^{+0.03}_{-0.03}$	$0.824^{+0.012}_{-0.012}$	$99.0^{+33.5}_{-27.7}$	$-0.73^{+0.26}_{-0.24}$	$91.0^{+112.2}_{-32.3}$	$35.7^{+17.3}_{-16.4}$	$48.619^{+0.007}_{-0.005}$	$0.44^{+0.39}_{-0.39}$	$0.0^{+0.7}_{-0.7}$	58

TABLE 5.2: Recovered parameters for the LDC sources. Rows are ordered by increasing source SNR and labelled by the injection ID in the LDC dataset. Sources merging within the mission lifetime (here set to 2.5 yr) are marked with stars. For each parameter (cf. Sec. 5.2.4), we quote median and 90% confidence intervals. In addition, we quote the area enclosed by the 90% contour level of the sky localization posterior $\Delta\Omega$ and the number of CPU hours required to perform each run (cf. Sec. 5.3.5).

Chapter 6

Gravitational lensing and stochastic background

Contribution summary

This Chapter is an edited and reformatted version of [4] and [5]:

R. Buscicchio, C.J. Moore, G. Pratten, P. Schmidt, M. Bianconi and A. Vecchio - Constraining the Lensing of Binary Black Holes from Their Stochastic Background - published in Physical Review Letters, Volume 125(14):141102, (2020)

and

R. Buscicchio, C.J. Moore, G. Pratten, P. Schmidt, and A. Vecchio - Constraining the lensing of binary neutron stars from their stochastic background - published in Physical Review D, Volume 102(8):081501, (2020)

I conceived the study with the help of prof. A. Vecchio and Dr. C. Moore, designed the analytical formalism and wrote the code required to produce the results presented here. I've produced all plots shown in this Chapter, drafted and finalized the text incorporating comments and suggestions from the other co-authors.

6.1 Constraining the lensing of BBHs from the SGWB

6.1.1 Introduction

Several binary black hole (BBH) mergers have been detected so far [238, 287] and a number of additional candidates reported [97, 288, 289, 290]. Forthcoming gravitational-wave (GW) detector upgrades will provide increased sensitivity, which will allow us to probe an even larger spacetime volume [162].

The current BBH detections are loud and individually resolvable [15, 270]. However, they are part of a much larger population [162] whose properties, such as the overall merger rate and the source mass distribution, can be inferred statistically [46, 174]. As new GW events are detected, this population can be constrained with increasing accuracy. The GW ensemble redshift distribution and correlations with source parameters constitute an important piece of evidence, allowing us to place tighter constraints on progenitors formation history and evolution channels [291, 178, 292, 293, 294, 295, 296, 297]. Ultimately, observing distinctive features in the population distribution would provide independent characterization of the expansion history of nearby universe [298]. Importantly, this population does not only consist of individually detectable BBH mergers but will contain many other distant, unresolved events [299]. Their emissions accumulate across all redshifts as a stochastic background of GWs (SGWB): an incoherent superposition of signals whose properties cannot be inferred individually [123, 300].

Broadly speaking, events are individually observable depending on the instrument sensitivities and the choice of search strategy [301, 156, 302, 303]. The majority of events that are not individually observed contribute instead to the SGWB. Current estimates predict a detection of a SGWB with a signal-to-noise ratio (SNR) of 3 after 40 months of observations [124, 120]. The observation of a stochastic background will complement individual detections, providing an integrated measure of the cosmological black holes' population history [126].

GWs from BBHs are generated by the dynamics of vacuum spacetime, as prescribed by general relativity. As a consequence, they carry information from an inherently scale-free physics. Additional assumptions on the formation mechanism, or observations of a counterpart are necessary to connect with weak or electromagnetic phenomena thereby introducing new length/energy scales and breaking the ubiquitous mass-distance degeneracy [298].

However, GWs are in principle affected by the intervening gravitational potential which influences the inferred spatial and temporal properties of the signals [304]. At the simplest level of description, the effect of lensing on a GW signal is to change its strain amplitude by a multiplicative magnification factor $\sqrt{\mu}$. As a consequence, and in absence of independent constraints on the lensing magnification, the mass-distance degeneracy is re-established even for chirping sources. Parameter estimation pipelines do not currently incorporate any lensing model, and therefore infer source properties agnostically of such a phenomenon. However, follow-up studies have addressed a number of questions: Is any detection actually magnified? Are there event couples originating from the same source emission, whose light-path has been altered to mimic independent events? Does lensing affect the population inference? [305, 306, 307, 308, 309, 310, 311, 312, 313, 314, 315].

In this letter we address one of the above questions, rephrasing it as a probabilistic statement. Given a set of observations, how likely is it for a fraction to be magnified by more than a certain μ? We show that by considering lensing of the entire population a significant amount of information can be leveraged from the SGWB; even the current non-detection has surprising astrophysical consequences. Significantly, we find the recent claims of lensed events to be statistically disfavoured [313, 314, 315].

6.1.2 Models

We now turn our attention to the modeling assumptions made. Firstly, we describe the effect of lensing on GW signals, and the parameterization of the lensing probability model. Secondly, we summarize the features of the population model for BBH mergers. Finally, we derive the associated energy density of the stochastic background including lensed events. Throughout this letter we use $G = c = 1$.

6.1.3 Lensing probability

Unlensed, chirping binaries provide a direct measurement of their luminosity distance d_L [316, 317]. If associated with electromagnetic counterparts, this gives an independent estimate of the source redshift z. Together, these constitute a point measurement in the expansion history of the universe [318, 319]

$$\frac{d_L(z)}{1+z} = \frac{1}{H_0} \int_0^z \mathrm{d}z' \frac{1}{E(z')} \,, \tag{6.1}$$

where H_0 is the local Hubble constant. $E(z)$ is a function of redshift, proportional to the time derivative of the logarithm of the scale factor, and encodes the information on the cosmological density parameters.

Alternatively, assuming a cosmological model breaks the mass-redshift degeneracy, thereby providing a redshift estimate for each observed event. However, this degeneracy is re-established by the addition of an *a priori* unknown lensing magnification μ.

Given a GW event, we focus on its true luminosity distance $d_L(z)$, chirp mass $\mathcal{M}$, and lensing magnification μ. Its strain amplitude is magnified by a multiplicative factor $\sqrt{\mu}$ [304]. Independent of the cosmology, the apparent mass $\tilde{\mathcal{M}}$, redshift $\tilde{z}$, and distance

$\tilde{d}_L$ are related to their true values by the following relationships:

$$\frac{d_L(\tilde{z})}{\sqrt{\tilde{\mu}}} = \frac{d_L(z)}{\sqrt{\mu}} \;, \qquad \tilde{\mathcal{M}}(1+\tilde{z}) = \mathcal{M}(1+z) \;. \tag{6.2}$$

The apparent parameters are those inferred by any pipeline that *assumes* a certain magnification $\tilde{\mu}$. Parameter estimates provided in published catalogues are computed under the assumption of no lensing, i.e. $\tilde{\mu} = 1$ [238, 288, 289, 290].

In order to incorporate the effect of lensing in the parameter reconstruction, additional independent information on the same transient would be required: e.g. the observation of electromagnetic counterparts, a detailed knowledge of the lensing potential along the GW travel path, or an association with a host galaxy. Another possibility is the association between two or more GW events, whose apparent properties can be referred back to a common source that has undergone multiple imaging [306]. In the absence of such additional information, prior knowledge on μ remains unaltered after any single detection, because of the above degeneracy.

In this letter we use a semi-analytic *lensing model* for the probability of a given magnification $\mathrm{d}P/\mathrm{d}\ln\mu$ from equation (B1) in [309] :

$$\frac{dP(\mu)}{d\ln\mu} = A\left(t_0\right) \int_0^{+\infty} dt \exp\left[\frac{\lambda}{t+t_0} - 2t\right] \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(\ln\mu - \delta - t)^2}{2\sigma^2}\right] \tag{6.3}$$

While A and δ are fixed by the probability distribution normalization and mean magnification to 1, the parameters σ, t_0, λ characterize the shape of the distribution and are matched to observations and large-scale simulations. We interpolate with cubic splines the fitted parameters from Table I in [309] across the redshift range $z \in [0, 20]$ We note that this model correctly captures the limiting behavior in both the strong and weak lensing regimes [320, 321], and is in agreement with recent hydrodynamical simulations [322]. Greater details on the lensing model are provided in [309].

6.1.4 Apparent merger distribution

Following [295, 46] (which are based on [238]), we parametrize the BBH differential merger rate R as a function of the binary masses $m_{1,2}$ and redshift z, as

$$\frac{\mathrm{d}^3 R}{\mathrm{d}m_1 \mathrm{d}m_2 \mathrm{d}z} = \mathcal{R}(z|\lambda,\gamma,z_\mathrm{P})p(m_1,m_2|m_\mathrm{min}, m_\mathrm{max}, \alpha), \tag{6.4}$$

where

$$p(m_1, m_2 \mid m_\mathrm{min}, m_\mathrm{max}, \alpha) \propto m_1^{-\alpha}\, \mathbf{I}(m_1 \mid m_\mathrm{min}, m_\mathrm{max})\, \mathbf{I}(m_2 \mid m_\mathrm{min}, m_1) \tag{6.5}$$

Here $\mathbf{I}(\cdot|a,b)$ are the indicator functions on the interval $[a,b]$, and throughout we adopt $\alpha = 2.3$ and $(m_\mathrm{min}, m_\mathrm{max}) = (5, 50)M_\odot$.

We set the cosmological merger rate to track the star formation rate (with no delay between formation and coalescence). This is modeled using a power law with index λ peaking at z_P and tapering off further in the past with index $(\lambda-\gamma)$ [323];

$$\mathcal{R}(z \mid \lambda, \gamma, z_\mathrm{P}) = R_0 \frac{(1+z)^\lambda}{1 + \left[\dfrac{1+z}{1+z_\mathrm{P}}\right]^\gamma}\,. \tag{6.6}$$

We tune R_0 to match the current estimate for the local merger rate from GW population analyses [46]. The uncertainty on the merger rate propagates to all redshifts affecting the entire population.

A straightforward consequence of fixing the mass and redshift distributions (Eqs. 6.5 and 6.6 respectively) is the amplitude of the stochastic background accumulated over

	R_0 (Gpc^{-3}yr^{-1})	λ	γ	$z_{\rm P}$
O1 + O2	57^{+40}_{-25}	$5.8^{+0.4}_{-0.4}$	5.6	1.9
Design	57^{+40}_{-25}	$3.4^{+0.6}_{-0.7}$	5.6	1.9

TABLE 6.1: Parameters modeling the merger rate density. R_0 is tuned to match the current estimate of the local merger rate, while $z_{\rm P}, \gamma$ capture the the star formation rate peak and decay further out in redshift. λ is adjusted to provide a stochastic background signal with a fixed SNR=2 at the two sensitivities considered. Fig. 6.1 shows the two resulting distributions.

the past light-cone of the observer. We have considered a number of merger rate models, varying both λ and γ, while keeping the star formation rate peak $z_{\rm P}$ fixed. In this letter, we present results for two choices of λ, with a fixed γ, that yield SGWB amplitudes consistent with current upper limits [126] (see Fig. 6.1, and Table 6.1, and the discussion in the following section).

For simplicity, we neglect in both models black hole spins. Depending on the spin properties of the BBHs, the enhancement on the overall rate can be significant, up to a factor of 3 in the mass range of interest for current detectors [296]. We leave a consistent inclusion of spin effects –i.e. on the intrinsic merger rate, on the spectral shape of the stochastic signal, and on individual event detectability– to future work.

We highlight that the local merger rate and mass distribution used here were obtained with hierarchical analyses on individual GW source parameters [49, 48, 50]. Therefore, in order to remain consistent with the prior assumptions therein, we have to consider z, m_1, m_2 as the *apparent* redshifts and masses with no intervening lensing, i.e. assuming $\tilde{\mu} = 1$. Henceforth, building on the notation in Eq. 6.2, we denote these parameters $\tilde{z}, \tilde{m}_1, \tilde{m}_2$, and related functions with a superscript tilde.

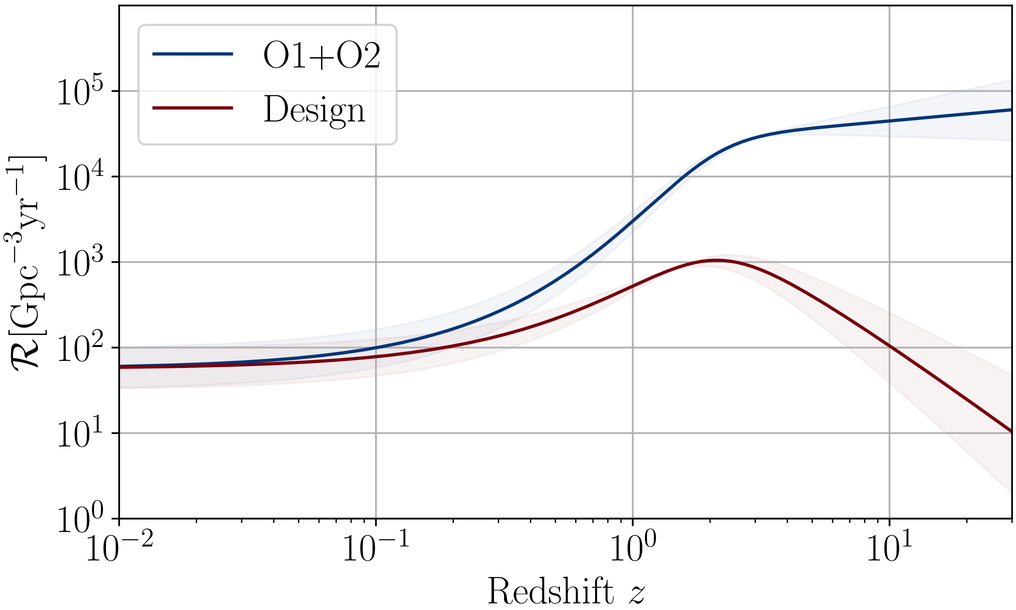

FIGURE 6.1: Cosmological merger rate density models considered, using the parameterization in Eq. 6.6. Parameter choices are listed in Table 6.1. Models are matched to the current local estimate for the BBH merger rate R_0. Assuming the functional form in Eq 6.6 we construct models providing a stochastic signal with SNR=2 at each given sensitivity (see Table 6.1 for the corresponding model parameters). The blue line refers to the sensitivity achieved after O1 and O2. The red line refers to the projected sensitivity after two years of observation at 50% duty cycle at design sensitivity. Shaded regions delimit similarly constructed models, matched to the upper and lower 90% confidence interval on the local merger rate estimate.

6.1.5 Lensed stochastic background

The stochastic background from BBH mergers is the incoherent superposition of individual GW events [121, 120]. We assume a flat ΛCDM cosmology and a simple leading order post-Newtonian expression for the GW energy spectrum from the inspiral of non-spinning BBHs [29]

$$\frac{\mathrm{d}E_{\mathrm{GW}}(m_1, m_2)}{\mathrm{d}f_r} = \frac{(\pi)^{2/3}}{3}\mathcal{M}(m_1, m_2)^{5/3} f_r^{-1/3}, \tag{6.7}$$

with $f_r = f(1+z)$ the GW frequency in the source rest frame. Integrating over the cosmological expansion history, gives a value for the energy density of GWs from BBH mergers expressed as a fraction of the critical density ρ_c: this is a standard result from the GW literature[121], and it reads

$$\Omega_{\rm BBH}(f) = \frac{1}{\rho_c}\int \mathrm{d}z \frac{f}{H_0(1+z)E(z)} \times \int \mathrm{d}m_1 \mathrm{d}m_2 \frac{\mathrm{d}^3 R}{\mathrm{d}m_1 \mathrm{d}m_2 \mathrm{d}z} \frac{\mathrm{d}E_{\rm GW}}{\mathrm{d}f_r}\bigg|_{f_r=f(1+z)} . \tag{6.8}$$

We stress here that Eq. 6.8 neglects the effects of lensing. Here, we seek to instead compute $\tilde{\Omega}_{\rm BBH}$ which accounts for the lensing model. In order to do this we must modify Eq. 6.8 by replacing $\{\mathrm{d}z, \mathrm{d}m_{1,2}\} \to \{\mathrm{d}\tilde{z}, \mathrm{d}\tilde{m}_{1,2}\}$, use the apparent differential merger rate $\mathrm{d}^3\tilde{R}$, and use the apparent redshifted frequency $f_r = f(1+\tilde{z})$.

We constrain the maximum allowed redshift evolution – λ in Eq. (6.6) – by considering upper-limits on a SGWB [124, 324], while keeping the local merger rate fixed to the observed value. We consider the current SGWB limit based on the O1 and O2 observing runs, using data from the two LIGO instruments only. As a limit for a non-detection we assume a signal-to-noise ratio smaller than 2 in a stochastic search [324]. Similarly, we forecast the projected limits after 2 years of observation at design sensitivity and 50% duty cycle of the network of the two LIGO instruments and Virgo. We denote the two scenarios *O1+O2* and *Design*, respectively.

As expected and clearly shown in Fig. 6.1 a non detection of a SGWB over longer integration time and with better sensitivities implies a lower merger rate outside the horizon for individual detections. The merger rate redshift evolution considered here is consistent with the results of [126].

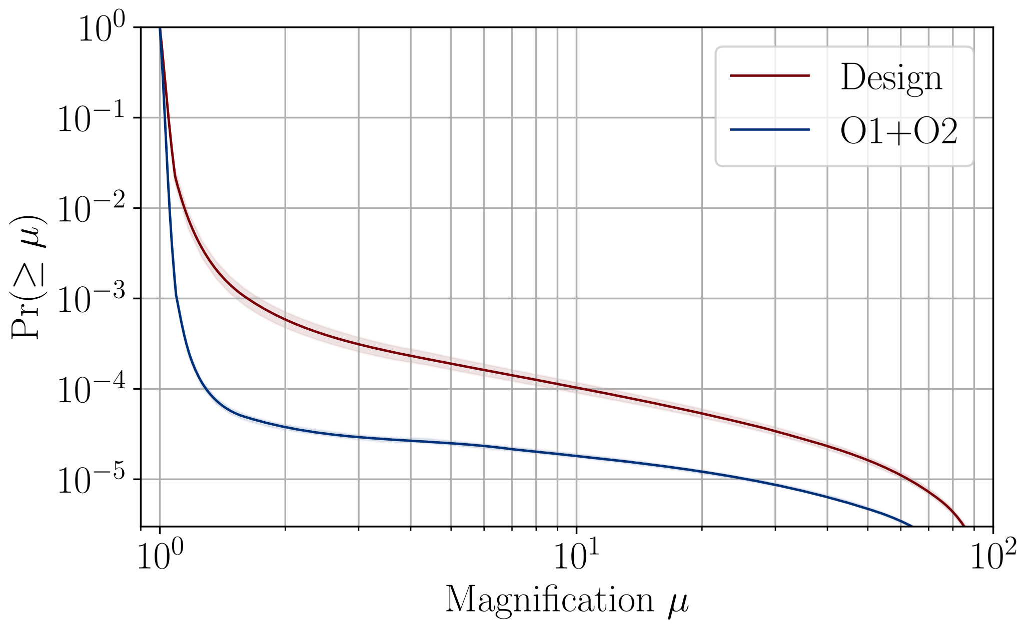

FIGURE 6.2: Complementary cumulative distribution for the lensing probability of detectable BBH mergers, constrained by the non-detection of the SGWB for two sensitivity scenarios. Solid lines and narrow shaded regions are obtained from corresponding models shown in Fig. 6.1. The fraction of lensed transients with $\mu > 2$ is less than $\sim 4 \times 10^{-5}$ after O1 and O2; a non detection of a SGWB after 2 years of operation at design sensitivity would yield a fraction a factor of 10 higher. The result depends very weakly on the local merger rate uncertainty, hence the light-blue shaded region has neglible width.

6.1.6 Lensed fraction

Having established the population models to be considered, we turn to our main task: quantifying the probability for an individual transient to be magnified with a particular magnification. For each apparent redshift shell $[\tilde{z}, \tilde{z} + \mathrm{d}\tilde{z}]$ we consider contributions from true redshifts shells up to $z = 20$, the maximum extent of the lensing model [309]. The pair $(z, \tilde{z})$ fixes uniquely the magnification and therefore the relationship between the two redshifts is given by Eq. 6.2 (with $\tilde{\mu} = 1$); this transformation, and its Jacobian, $|\partial \tilde{z}/\partial z|_{\mu}$, is further explored in the supplemental material.

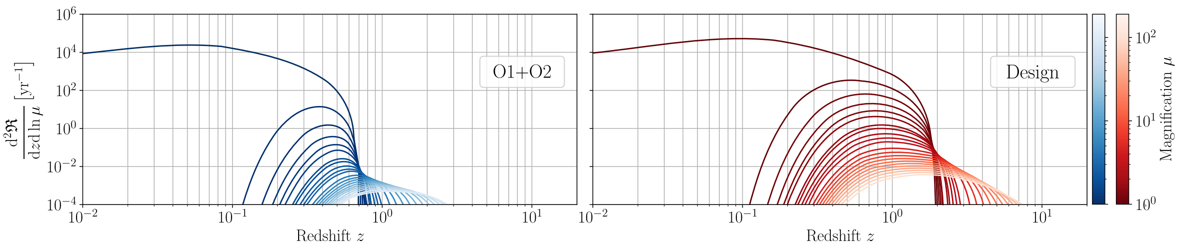

FIGURE 6.3: Differential rate of detectable lensed events for each redshift and logarithmic magnification bin. Results are shown as solid lines coloured according to magnification. Left (right) panels show results for the O1 and O2 (Design) population models respectively; see Table 6.1 and the accompanying discussion in the text. Moderately magnified events (e.g. $\mu < 10$) dominates the detected population of BBH mergers by at least three orders of magnitude. At Design sensitivity, a non-detection of a stochastic background will imply by itself a significant reduction of mergers at high redshift, as described by the model in Fig. 6.1. Concurrently, a better sensitivity enhances detections further out in redshift, at all magnifications. Predominantly, non magnified events will be observed out to $z \approx 2$. A very small fraction of strongly magnified ones will extend out $z \approx 6$. A comparison of the overall improvement integrated over redshift at each magnification, is presented in Fig. 6.2.

We use this relationship to write explicitly an expression for the differential rate of magnified events which is a proxy for the magnification probability,

$$\begin{aligned} \frac{\mathrm{d}^2\mathfrak{R}}{\mathrm{d}z\mathrm{d}\ln\mu} = \frac{\mathrm{d}P(\mu\mid z)}{\mathrm{d}\ln\mu}\frac{4\pi\chi^2\left(z\right)}{H_0\left(1+z\right)E\left(z\right)}\left|\frac{\partial\tilde{z}}{\partial z}\right|_\mu \\ \times\int\mathrm{d}\tilde{m}_1\mathrm{d}\tilde{m}_2\frac{\mathrm{d}^3\tilde{R}}{\mathrm{d}\tilde{m}_1\mathrm{d}\tilde{m}_2\mathrm{d}\tilde{z}}p_{\mathrm{det}}\left(\tilde{m}_1,\tilde{m}_2,\tilde{z}\right)\ . \end{aligned} \tag{6.9}$$

Additionally, using apparent masses and redshifts we filter events by their detectability. We use a fixed single detector threshold SNR $= 8$ for each given set of source parameters, and compute the observable fraction of the distribution in component masses, averaged over the source orientation [15]. We estimate selection effects $p_{\mathrm{det}}(\tilde{m}_1, \tilde{m}_2, \tilde{z})$ for both sensitivities using the publicly available code GWDET [325].

As discussed above, the SGWB should contain only unresolved events. To be consistent, the same selection effects should be added to Eq. 6.8 by the inclusion of a factor $(1-p_{\rm det}(\tilde{m}_1, \tilde{m}_2, \tilde{z}))$ in the innermost integral. However, we neglect this effect here because the region of parameter space where $p_{\rm det}$ is non-zero, i.e. at moderate masses and low redshift, is far from the peak of the intrinsic rate.

Results are shown in Figs. 6.2 and 6.3. Detections are dominated in both scenarios by low-redshift, unlensed events (i.e. $\mu \approx 1$). While at design sensitivity the detections will extend further out to $z \approx 2$, magnifications smaller than 2 will likely dominate the population by at least three orders of magnitude. This is clearly apparent in Fig. 6.2, where the contributions across redshifts are integrated out to a single magnification distribution. For ease of comparison we show both as cumulative distribution functions, i.e. factoring out the respective total rate of detections per year.

Remarkably, a better instrument sensitivity provides proportionally more events at larger magnification. This is the net result of a few competing factors. The assumed non-detection of a SGWB constrains the population to a shallower redshift distribution: as a consequence both lensed and unlensed events within the detection horizon are equally suppressed; however, the population of distant events at $z > 2$ is significantly depleted, therefore reducing their relative contribution to the apparent distribution.

We study the impact on our results of our modelling choices for (i) the lensing model, (ii) the redshift evolution of the merger rate, (iii) the BBH mass distribution. Overall we find our results to be robust; changing the mass distribution has the largest effect, increasing the fraction of lensed event by at most a factor of 2 (see details in Supplemental material).

6.1.7 Conclusions

A SGWB of astrophysical origin has not yet been observed. This constrains the redshift dependence of the BBH merger rate, particularly the number of mergers at high redshift. This in turn has consequences for the lensing probability of individual events. In this letter we exploit this surprising link between the non detection of a SGWB and the lensing probability to quantify the fraction of lensed BBH events. We provide estimates for the relative contribution of lensed BBHs to the total rate out to redshifts of $z \leq 20$ and magnifications of $\mu \leq 100$. Even the current non-detection of a SGWB already has interesting astrophysical implications; we find a fraction below $\sim 4 \times 10^{-5}$ of events to have a magnification $\mu \geq 2$. At design sensitivity, in the absence of a SGWB detection after two years of observation, this fraction increases by a factor of ~ 10.

If and when there is a detection of a SGWB, our argument will become even more informative. It can be applied to the BBH merger redshift distribution –constrained jointly from the mergers population and the SGWB detection– to predict the number of lensed events. For a detection of a SGWB in less than two years of observation at *Design* sensitivity, we expect the inferred lensing fraction to lie between the two curves shown in Fig. 6.2.

Simultaneously and independently, a similar study using complementary methods appeared [127] showing agreement with our results.

6.2 Constraining the lensing of BNSs from the SGWB

6.2.1 Introduction

Two binary binary neutron star (BNS) mergers have been detected so far [326, 327]. The event GW190814 [97] may also contain a neutron star. Forthcoming detector upgrades will provide better sensitivity, allowing us to probe ever larger spacetime volumes and detect more events [162]. Currently, the observed GW events involving neutron stars are loud and individually resolvable [15, 270]. However, many more events will lie below the threshold for detection, individually indistinguishable from the instrument noise. All these events are drawn from the same overall population. The unresolvable GW events, including those involving neutron stars, will pile up across the detector bandwidth and give a stochastic background of gravitational waves (SGWB). Such signal is subject to dedicated searches by current ground-based interferometers [124, 120, 328].

A fraction of BNS events will be gravitationally lensed; this has the effect of increasing the GW amplitude by a factor $\sqrt{\mu}$, where μ is the lensing magnification. The lensing of a GW depends on the intervening gravitational potential, and different events will experience different lensing magnifications. Multiply imaged GWs, phasing and wave-optics effects may also occur depending on the specific potential [304, 309, 311, 306, 329, 330].

In [4] we considered the analogous situation for binary black hole (BBH) GW events. That paper described in detail the formalism to quantify the impact of lensing on the amplitude and detection rate of individual events, as well as the amplitude of the associated SGWB (see also [127]). Subsequently, we leveraged the non detection of a stochastic background to get constraints on the probability of individual BBHs being lensed. In order to do so, we framed in a single statistical picture the observational data from GW detectors (i.e. individual events, stochastic background), their inferred properties, and the implication on the observation of lensed events.

In this paper we reapply the techniques from [4] to the analysis of BNS events. Using constraints on the BNS merger rate density after the first two observing runs [238], and the confirmed non-detection of a stochastic background [124], we report the implication on the expected number of lensed BNS observations, both in the weak and strong lensing regime. Throughout we follow the conventions of [4].

6.2.2 Models

We employ a semi-analytical model for the lensing probability $\mathrm{d}P(\mu \mid z)/\mathrm{d}\log\mu$ as a function of redshift out to redshifts $z \leq 20$. This model applies to magnifications μ up to $\mu \leq 200$, as described in [309] (i.e. it includes both strong and weak lensing). For details of our implementation of this lensing model we refer the interested reader to Appendix A of [4].

It is also necessary to model the BNS population. We neglect neutron star spins and matter effects (e.g. tides) as well as any orbital eccentricity in the binary. Under these simplifying assumptions a binary is described by the two component masses, m_1 and m_2. We model the distribution of component masses, $p(m_1, m_2)$, in three different scenarios. The first two match those employed in the rates analysis of GWTC-1 [238] (see Section VII.C). The third is included to investigate the effect of the width of the mass distribution.

- "Uniform"; the component masses m_1, m_2 are drawn independently from a uniform distribution in the range $[0.8M_{\odot}, 2.3M_{\odot}]$.
- "Gaussian"; the component masses are drawn from a Gaussian distribution with mean $1.33M_{\odot}$ and standard deviation $0.09M_{\odot}$.
- "Fixed"; all neutron star masses are equal to $1.4M_{\odot}$.

We choose to model the redshift evolution of the BNS merger rate $\mathcal{R}(z)$ by tracking the star formation rate. For details, see Equation 5 and Figure 1 of [4]. In a slight deviation from the previous study, we keep the population extinction (i.e. $\lambda - \gamma$) fixed at high redshifts (i.e. $z \gg z_{\mathrm{P}}$) (following Madau-Dickinson [323]) while increasing the

Mass	$R_0\,(\mathrm{Gpc}^{-3}\mathrm{yr}^{-1})$	Sensitivity	λ
Uniform	800^{+1970}_{-680}	O1 + O2 Design	$5.83^{+1.54}_{-1.06}$ $3.26^{+1.76}_{-1.41}$
Gaussian	1210^{+3230}_{-1040}	O1 + O2 Design	$5.686^{+1.6}_{-1.12}$ $3.078^{+1.85}_{-1.54}$
Fixed	1210^{+3230}_{-1040}	O1 + O2 Design	$5.685^{+1.6}_{-1.12}$ $3.077^{+1.85}_{-1.54}$

TABLE 6.2: Parameters modelling the merger rate density. The local rate R_0 for the "Uniform" and "Gaussian" cases are taken from the analysis in [238]. For the "Fixed" case we use the same values as for the "Gaussian" case. As described in the text, we increase the slope of the local merger rate λ to a value that gives a SGWB signal with a marginally detectable signal-to-noise ratio of 2 at the sensitivities considered.

slope of the local merger rate (λ). We fix the local rate (R_0) to the estimates provided in [238], while varying λ (see Table 6.2). We focus on results from one pipeline search only (pyCBC [156]); analogous results have been computed for other pipelines (e.g. GstLAL [155]) and differences are at the level of 1%.

By changing the value of λ, we effectively set the level of the BNS stochastic background (see Equation 7 in [4]), including the effect of lensing, to a marginally detectable level. As in [4], this is done for two different detector network sensitivities (namely the existing "O1+O2" sensitivity and a projected future "Design" sensitivity).

This section has briefly described the various modelling assumptions made for (i) the lensing probability, (ii) the properties of the BNS population, and (iii) the cosmological evolution of the BNS merger rate. We refer the reader to the Appendices in [4] for a more detailed discussion of the impact of these assumptions on the results of our analysis.

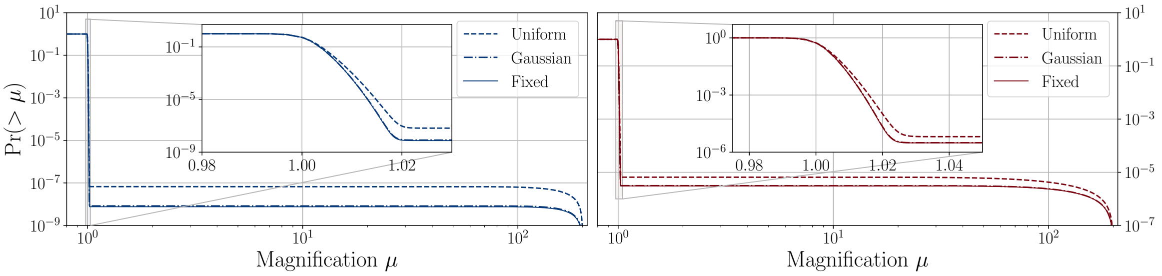

FIGURE 6.4: The complementary cumulative distribution for the lensing probability of BNS mergers. This is constrained by the current (left, blue) or future (right, red) non-detection of a SGWB. The fraction of lensed transients with $\mu > 1.02$ is less than $\sim 7 \times 10^{-8}$ after O1 and O2; a non detection of a SGWB after 2 years of operation at design sensitivity would increase the fraction of lensed event to $\sim 7 \times 10^{-6}$. The inset plots zoom in on a narrow range of magnification around $\mu \sim 1$ where the lensing probability steeply decreases. In our lensing model, there are no BNS events lensed with $\mu > 200$.

6.2.3 Results

The main result of our calculation is shown in Fig. 6.4. Here we plot the probability of a BNS event having a magnification above a certain value. Across all three models that we considered for the BNS masses and redshift distribution, the fraction of lensed events with magnification $\mu > 1.02$ is lower than 7×10^{-8} for the "O1+O2" sensitivity. We find it extremely unlikely on statistical ground that a significantly lensed binary neutron star will be observed in the near future. These complementary cumulative distributions shows a large drop just above $\mu = 1$ and then a very prominent plateau out to magnifications of $\mu \sim 200$ before dropping to zero. This behaviour can be understood by looking at the contribution of lensed events broken down by redshift as shown in Fig. 6.5. The main contribution to the observed BNS population (left blob in the plots of Fig. 6.5) gives the high probability near $\mu \sim 1$. The gap at intermediate redshifts visible is responsible for the plateau. And finally, in our model, there are no events with $\mu > 200$, and this is responsible for the final drop off.

At design sensitivity we expect a slightly higher fraction of lensed events (7×10^{-6}).

This can be seen in right panel of Fig. 6.4 where a non detection of a stochastic signal yields a higher value for the plateau.

6.2.4 Discussion

A SGWB from BNS events has not yet been observed and is expected to be subdominant with respect to the background from BBH events. The current non-detection places a constraint on the redshift evolution of the BNS merger rate; in particular it limits the rise in the rate at redshifts around $z \sim 2$. This in turn has important implications for the lensing of individual BNS events. Here, we have used the current non-detection of a SGWB to constrain the probability that an individual BNS event is magnified by more than a certain amount. In particular we find that the probability that $\mu > 1.02$ is less than $\sim 7 \times 10^{-8}$. This probability increases slightly for detectors upgraded towards design sensitivity, but remains small. Therefore significantly lensed BNS events should not be expected in the near future.

6.2.5 Updates after LIGO/Virgo third observing runs

Using data and results from the third observing run, I performed a similar analysis to the two presented above, with updated mass distribution, merger rate density, and selection effects constraints from the non-detection of a stochastic background. This study was performed as part of a broader collaboration work, searching for signatures of gravitational lensing of GWs in LIGO and Virgo data. Results are broadly consistent with the ones presented here, and the interested reader can compare with Sec.3 of [6].

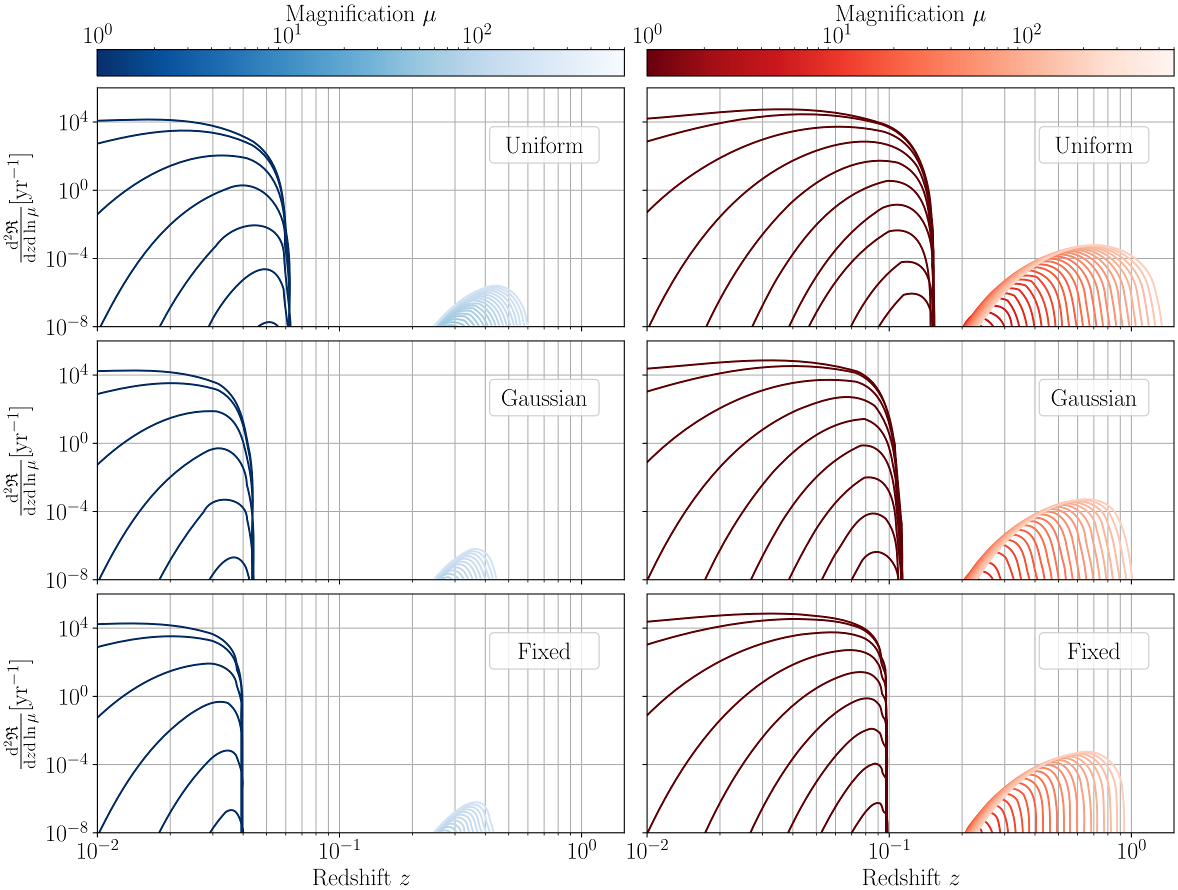

FIGURE 6.5: Differential rate of lensed events for each redshift and logarithmic magnification bin. Left column (blue) show results for the "O1+O2" sensitivity while the right column (red) shows results for "Design" sensitivity. The three rows shows the results for the different BNS mass distributions described in the text. Solid lines is colored according to magnification. Moderately magnified events (i.e. $\mu < 1.02$) dominate the detected population. For these moderately magnified events, there is a clear horizon redshift beyond which unlensed sources cannot be seen (e.g. around $z \sim 0.6 - 0.7$ in the top-left plot). To see more distant events we need them to be significantly magnified. However, there are few large lenses at low redshifts to provide this magnification. There is therefore a gap out to $z \sim 3$ where large lenses are plentiful and there is a secondary contribution to the observed BNS population. The secondary contribution is more significant at "Design" sensitivity.

Chapter 7

Conclusions and Prospects

This Chapter contains conclusions and prospects for future work. No original work is presented here.

7.1 Conclusions

In this thesis we have explored a number of topics at the intersection between GW astronomy and Bayesian statistics.

We discussed the problem of parameter estimation on multiple indistinguisable sources with LISA, and we have introduced a technique to significantly simplify the structure of the parameter space, drastically reducing the excess multimodality emerging from the likelihood, with no information loss. The solution we propose is efficient and simple: it acts as a transformation on the prior parameter space, hence from a computational point of view keeps unaltered the specific likelihood implementation. It turned out that an identical strategy can be applied on population inference with mixture models, whose components are intrinsically indistinguishable. We showed an example application of such application to the inference on the component masses distribution of BBHs observed by LIGO and Virgo. Ultimately, within the context of the "global-fit" in by LISA, the same strategy will be applicable to simplify each set of indistinguishable sources parameter estimation.

We also discussed the science case of a detection of DWD in known satellites of the Milky Way. We showed –through the development of a detailed Bayesian inference pipeline, and its application to a set of representative DWD sources informed by population synthesis models– that LISA will be able to either associate the DWDs to their host satellites, or discover new ones in regions otherwise unprobed by other observatories (e.g. the galactic disk and the bulge).

Going forward we expanded the capabilities of our simulation code, focusing on Bayesian parameter estimation of SmBBHs with LISA. By targeting precessing eccentric unequal mass binaries, we have challenged ourselves with the source class parameter space at full dimensionality. By showing that LISA will be able to predict the time to merger to within $\sim$ 1hr, we argued that a "multiband" detection will be possible in conjunction with ground-based detectors and quantified reliably the typical errors in the recovered parameters. We showed an example inference on the orbital eccentricity using the inspiral signal in the LISA sensitivity band, which would be otherwise unobserved to ground-based detectors. This will serve as a powerful probe among different formation channels.

Benchmarking LISA performances through extensive SmBBHs injections and recovery campaigns will be crucial to better explore the full parameter space, ultimately to understand correlations and degeneracies arising in realistic scenarios.

Exploring the middle ground between individually resolved SmBBHs and the stochastic foreground originating from the unresolved ones, we have investigated how the latter can constraints the former in the specific context of gravitationally lensed GWs. Highlighting how the merger rate density, spanning across redshifts in and out of the detector horizon, constitutes a link between resolved and unresolved sources, we showed a robust prediction of the probability of observing gravitationally lensed events compatible with a fixed maximum stochastic background energy density. Our results relied on the ineffectiveness of lensing magnification on the isotropic SGWB, which breaks when an angular dependence on its power spectrum is introduced. Further studies to explore this effect are required.

Bibliography

[1] Riccardo Buscicchio, Elinore Roebber, Janna M. Goldstein, and Christopher J. Moore. Label switching problem in Bayesian analysis for gravitational wave astronomy. PRD, 100(8):084041, October 2019.

[2] Elinore Roebber et al. Milky Way Satellites Shining Bright in Gravitational Waves. *Astrophys. J. Lett.*, 894(2):L15, 2020.

[3] Riccardo Buscicchio, Antoine Klein, Elinore Roebber, Christopher J. Moore, Davide Gerosa, Eliot Finch, and Alberto Vecchio. Bayesian parameter estimation of stellar-mass black-hole binaries with LISA. *Phys. Rev. D*, 104(4):044065, 2021.

[4] Riccardo Buscicchio, Christopher J. Moore, Geraint Pratten, Patricia Schmidt, et al. Constraining the Lensing of Binary Black Holes from Their Stochastic Background. PRL, 125(14):141102, October 2020.

[5] Riccardo Buscicchio, Christopher J. Moore, Geraint Pratten, Patricia Schmidt, and Alberto Vecchio. Constraining the lensing of binary neutron stars from their stochastic background. *Phys. Rev. D*, 102(8):081501, 2020.

[6] R. Abbott et al. Search for lensing signatures in the gravitational-wave observations from the first half of LIGO-Virgo's third observing run. *arXiv e-prints*, 5 2021.

[7] Christopher J. Moore, Eliot Finch, Riccardo Buscicchio, and Davide Gerosa. Testing general relativity with gravitational-wave catalogs: The insidious nature of waveform systematics. *iScience*, 24(6):102577, June 2021.

[8] Michele Maggiore. *Gravitational waves: Volume 1: Theory and experiments*, volume 1. Oxford university press, 2008.

[9] Albert Einstein and Marcel Grossmann. Entwurf einer verallgemeinerten relativitätstheorie und eine theorie der gravitation. *Zeitschrift für Mathematik und Physik*, 62:225–261, 1913.

[10] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Princeton University Press, 2017.

[11] F. Acernese, M. Agathos, K. Agatsuma, et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, January 2015.

[12] C. J. Moore, R. H. Cole, and C. P. L. Berry. Gravitational-wave sensitivity curves. *Classical and Quantum Gravity*, 32(1):015014, January 2015.

[13] LIGO Scientific Collaboration, J. Aasi, B. P. Abbott, R. Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, April 2015.

[14] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, and et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Reviews in Relativity*, 23(1):3, September 2020.

[15] Lee Samuel Finn and David F. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. PRD, 47(6):2198–2219, March 1993.

[16] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. , February 2017.

[17] Jun Luo et al. TianQin: a space-borne gravitational wave detector. *Class. Quant. Grav.*, 33(3):035010, 2016.

[18] Michele Vallisneri. Synthetic LISA: Simulating time delay interferometry in a model LISA. PRD, 71(2):022001, January 2005.

[19] Louis J. Rubbo, Neil J. Cornish, and Olivier Poujade. Forward modeling of space borne gravitational wave detectors. *Phys. Rev. D*, 69:082003, 2004.

[20] F. B. Estabrook and H. D. Wahlquist. Response of Doppler spacecraft tracking to gravitational radiation. *General Relativity and Gravitation*, 6(5):439–447, October 1975.

[21] A. Klein. The lisa response to gravitational waves. in preparation.

[22] Massimo Tinto and J. W. Armstrong. Cancellation of laser noise in an unequal-arm interferometer detector of gravitational radiation. *Phys. Rev. D*, 59:102003, 1999.

[23] Alberto Vecchio and Elizabeth D. L. Wickham. The Effect of the LISA response function on observations of monochromatic sources. *Phys. Rev. D*, 70:082002, 2004.

[24] Massimo Tinto and Sanjeev V. Dhurandhar. Time-delay interferometry. *Living Rev. Rel.*, 24(1):1, 2021.

[25] Michele Vallisneri. Synthetic LISA: Simulating time delay interferometry in a model LISA. *Phys. Rev. D*, 71:022001, 2005.

[26] Massimo Tinto and Sanjeev V. Dhurandhar. Time-delay interferometry. *Living Reviews in Relativity*, 24(1):1, December 2021.

[27] Albert Einstein. Über Gravitationswellen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin*, pages 154–167, January 1918.

[28] Benjamin P. Abbott et al. The basic physics of the binary black hole merger GW150914. *Annalen Phys.*, 529(1-2):1600209, 2017.

[29] P. C. Peters and J. Mathews. Gravitational Radiation from Point Masses in a Keplerian Orbit. *Physical Review*, 131(1):435–440, July 1963.

[30] Thibault Damour, Achamveedu Gopakumar, and Bala R. Iyer. Phasing of gravitational waves from inspiralling eccentric binaries. PRD, 70(6):064028, September 2004.

[31] Atsushi Nishizawa, Emanuele Berti, Antoine Klein, and Alberto Sesana. eLISA eccentricity measurements as tracers of binary black hole formation. PRD, 94(6):064020, September 2016.

[32] Dragoljub Marković. Possibility of determining cosmological parameters from measurements of gravitational waves emitted by coalescing, compact binaries. *Phys. Rev. D*, 48:4738–4756, Nov 1993.

[33] Curt Cutler and Éanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform? *Phys. Rev. D*, 49:2658–2697, Mar 1994.

[34] Samaya Nissanke, Daniel E. Holz, Scott A. Hughes, Neal Dalal, and Jonathan L. Sievers. Exploring short gamma-ray bursts as gravitational-wave standard sirens. *Astrophys. J.*, 725:496–514, 2010.

[35] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries. PRD, 49:6274–6297, Jun 1994.

[36] Lawrence E. Kidder. Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects. PRD, 52:821–847, 1995.

[37] Alejandro Bohé et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Phys. Rev.*, D95(4):044028, 2017.

[38] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93(4):044006, 2016.

[39] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D*, 93(4):044007, 2016.

[40] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.

[41] Cecilio García-Quirós, Marta Colleoni, Sascha Husa, Héctor Estellés, Geraint Pratten, Antoni Ramos-Buades, Maite Mateu-Lucena, and Rafel Jaume. Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries. *Phys. Rev. D*, 102(6):064002, 2020.

[42] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Leïla Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. PRL, 113(15):151101, Oct 2014.

[43] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. *Phys. Rev. D*, 102(4):044055, 2020.

[44] Nikolaos Karnesis, Stanislav Babak, Mauro Pieroni, Neil Cornish, and Tyson Littenberg. Characterization of the stochastic signal originating from compact binary populations as measured by LISA. PRD, 104(4):043019, August 2021.

[45] B.P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9(3):031040, 2019.

[46] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, and et al. Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo. ApJ, 882(2):L24, September 2019.

[47] R. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *Phys. Rev. X*, 11:021053, 2021.

[48] Ilya Mandel. Parameter estimation on gravitational waves from multiple coalescing binaries. PRD, 81(8):084029, April 2010.

[49] David W. Hogg, Adam D. Myers, and Jo Bovy. Inferring the Eccentricity Distribution. ApJ, 725(2):2166–2175, December 2010.

[50] Ilya Mandel, Will M. Farr, and Jonathan R. Gair. Extracting distribution parameters from multiple uncertain observations with selection biases. MNRAS, 486(1):1086–1093, June 2019.

[51] Christopher J. Moore and Davide Gerosa. Population-informed priors in gravitational-wave astronomy. *arXiv e-prints*, 8 2021.

[52] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, and et al. Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog. ApJ, 913(1):L7, May 2021.

[53] H A Bethe and G E Brown. Evolution of binary compact objects that merge. *Astrophys. J.*, 506:780, 1998.

[54] S. F. Portegies Zwart and L. R. Yungelson. Formation and evolution of binary neutron stars. *Astron. Astrophys.*, 332:173, 1998.

[55] Krzysztof Belczynski, Vassiliki Kalogera, and Tomasz Bulik. A comprehensive study of binary compact objects as gravitational wave sources: Evolutionary channels, rates, and physical properties. *Astrophys. J.*, 572:407, 2002.

[56] Steinn Sigurdsson and Lars Hernquist. Primordial black holes in globular clusters. *Nature*, 364:423, 1993.

[57] Simon F. Portegies Zwart and Stephen L. W. McMillan. Black hole mergers in the universe. *ApJ Lett.*, 528:L17, 2000.

[58] Carl L. Rodriguez, Sourav Chatterjee, and Frederic A. Rasio. Binary black hole mergers from globular clusters: Masses, merger rates, and the impact of stellar evolution. PRD, 93:084029, April 2016.

[59] Cristobal Petrovich and Fabio Antonini. Greatly enhanced merger rates of compact-object binaries in non-spherical nuclear star clusters. *Astrophys. J.*, 846(2):146, 2017.

[60] Giacomo Fragione and Bence Kocsis. Black hole mergers from an evolving population of globular clusters. *Phys. Rev. Lett.*, 121:161103, 2018.

[61] Imre Bartos, Bence Kocsis, Zolt Haiman, and Szabolcs Márka. Rapid and Bright Stellar-mass Binary Black Hole Mergers in Active Galactic Nuclei. *Astrophys. J.*, 835(2):165, 2017.

[62] Nicholas C. Stone, Brian D. Metzger, and Zoltan Haiman. Assisted inspirals of stellar mass black holes embedded in AGN discs: solving the final au problem. *Mon. Not. Roy. Astr. Soc.*, 464:946, 2017.

[63] M. Gröbner, W. Ishibashi, S. Tiwari, M. Haney, and P. Jetzer. Binary black hole mergers in AGN accretion discs: gravitational wave rate density estimates. A&A, 638:A119, June 2020.

[64] Hiromichi Tagawa, Zoltan Haiman, and Bence Kocsis. Formation and evolution of compact-object binaries in agn disks. *The Astrophysical Journal*, 898(1):25, Jul 2020.

[65] J. J. Eldridge and E. R. Stanway. BPASS predictions for Binary Black-Hole Mergers. *Mon. Not. Roy. Astron. Soc.*, 462(3):3302–3313, 2016.

[66] Nicola Giacobbo and Michela Mapelli. The progenitors of compact-object binaries: impact of metallicity, common envelope and natal kicks. *Mon. Not. Roy. Astron. Soc.*, 480(2):2011–2030, 2018.

[67] Simon Stevenson, Alejandro Vigna-Gómez, Ilya Mandel, Jim W. Barrett, Coenraad J. Neijssel, David Perkins, and Selma E. de Mink. Formation of the first three gravitational-wave observations through isolated binary evolution. *Nature Commun.*, 8:14906, 2017.

[68] Pablo Marchant, Norbert Langer, Philipp Podsiadlowski, Thomas M. Tauris, and Takashi J. Moriya. A new route towards merging massive black holes. *Astron. Astrophys.*, 588:A50, 2016.

[69] S. E. de Mink and I. Mandel. The chemically homogeneous evolutionary channel for binary black hole mergers: rates and properties of gravitational-wave events detectable by advanced LIGO. *Mon. Not. Roy. Astron. Soc.*, 460(4):3545–3553, 2016.

[70] Ilya Mandel and Selma E. de Mink. Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries. *Mon. Not. Roy. Astron. Soc.*, 458(3):2634–2647, 2016.

[71] Michal Dominik, Krzysztof Belczynski, Christopher Fryer, Daniel Holz, Emanuele Berti, Tomasz Bulik, Ilya Mandel, and Richard O'Shaughnessy. Double Compact Objects I: The Significance of the Common Envelope on Merger Rates. *Astrophys. J.*, 759:52, 2012.

[72] Nicola Giacobbo, Michela Mapelli, and Mario Spera. Merging black hole binaries: the effects of progenitor's metallicity, mass-loss rate and Eddington factor. *Mon. Not. Roy. Astron. Soc.*, 474(3):2959–2974, 2018.

[73] Ugo N. Di Carlo et al. Binary black holes in young star clusters: the impact of metallicity. *Mon. Not. Roy. Astron. Soc.*, 498(1):495–506, 2020.

[74] Fabio Antonini, Norman Murray, and Seppo Mikkola. Black hole triple dynamics: A breakdown of the orbit average approximation and implications for gravitational wave detections. *Astrophys. J.*, 781:45, 2014.

[75] Thomas O. Kimpson, Mario Spera, Michela Mapelli, and Brunetto M. Ziosi. Hierarchical black hole triples in young star clusters: impact of kozai-lidov resonance on mergers. *Mon. Not. Roy. Astr. Soc.*, 463:2443, 2016.

[76] Fabio Antonini, Silvia Toonen, and Adrian S. Hamers. Binary black hole mergers from field triples: Properties, rates, and the impact of stellar evolution. *Astrophys. J.*, 841:77, 2017.

[77] Davide Gerosa and Maya Fishbach. Hierarchical mergers of stellar-mass black holes and their gravitational-wave signatures. *Nature Astron.*, 5(8):8, 2021.

[78] B. J. Carr and S. W. Hawking. Black holes in the early universe. *Mon. Not. Roy. Astr. Soc.*, 168:399, 1974.

[79] Nicolas Fernandez and Stefano Profumo. Unraveling the origin of black holes from effective spin measurements with LIGO-Virgo. J. Cosmology Astropart. Phys., 2019(8):022, August 2019.

[80] Michael Zevin, Chris Pankow, Carl L. Rodriguez, Laura Sampson, Eve Chase, Vassiliki Kalogera, and Frederic A. Rasio. Constraining formation models of binary black holes with gravitational-wave observations. *Astrophys. J.*, 846:82, 2017.

[81] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev. D*, 49:6274–6297, 1994.

[82] K. Belczynski, G. Wiktorowicz, C. Fryer, D. Holz, and V. Kalogera. Missing Black Holes Unveil The Supernova Explosion Mechanism. *Astrophys. J.*, 757:91, 2012.

[83] Krzysztof Belczynski, Daniel E. Holz, Tomasz Bulik, and Richard O'Shaughnessy. The first gravitational-wave source from the isolated evolution of two 40-100 Msun stars. *Nature*, 534:512, 2016.

[84] K. Belczynski et al. The evolutionary roads leading to low effective spins, high black hole masses, and O1/O2 rates of LIGO/Virgo binary black holes. *Astron. Astrophys.*, 636:A104, 2020.

[85] Chase Kimball, Colm Talbot, Christopher P. L. Berry, Matthew Carney, Michael Zevin, Eric Thrane, and Vicky Kalogera. Black Hole Genealogy: Identifying Hierarchical Mergers with Gravitational Waves. *Astrophys. J.*, 900(2):177, 2020.

[86] Carl L. Rodriguez, Carl-Johan Haster, Sourav Chatterjee, Vicky Kalogera, and Frederic A. Rasio. Dynamical Formation of the GW150914 Binary Black Hole. *Astrophys. J. Lett.*, 824(1):L8, 2016.

[87] Carl L. Rodriguez, Michael Zevin, Chris Pankow, Vasilliki Kalogera, and Frederic A. Rasio. Illuminating Black Hole Binary Formation Channels with Spins in Advanced LIGO. *Astrophys. J. Lett.*, 832(1):L2, 2016.

[88] Carl L. Rodriguez and Abraham Loeb. Redshift Evolution of the Black Hole Merger Rate from Globular Clusters. *Astrophys. J. Lett.*, 866(1):L5, 2018.

[89] Will M. Farr, Simon Stevenson, M. Coleman Miller, Ilya Mandel, Ben Farr, and Alberto Vecchio. Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves. *Nature*, 548:426, 2017.

[90] Charles D. Bailyn, Raj K. Jain, Paolo Coppi, and Jerome A. Orosz. The Mass distribution of stellar black holes. ApJ, 499:367, 1998.

[91] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock. The Black Hole Mass Distribution in the Galaxy. ApJ, 725:1918–1927, December 2010.

[92] Will M. Farr, Niharika Sravan, Andrew Cantrell, Laura Kreidberg, and et al. The Mass Distribution of Stellar-mass Black Holes. ApJ, 741(2):103, November 2011.

[93] F. Özel, D. Psaltis, R. Narayan, and A. S. Villarreal. On the Mass Distribution and Birth Masses of Neutron Stars. ApJ, 757:55, 2012.

[94] B.P. Abbott et al. GW170817: Measurements of neutron star radii and equation of state. *Phys. Rev. Lett.*, 121(16):161101, 2018.

[95] B.P. Abbott et al. Properties of the binary neutron star merger GW170817. *Phys. Rev. X*, 9(1):011001, 2019.

[96] B.P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$. *Astrophys. J. Lett.*, 892:L3, 2020.

[97] R. Abbott, T. D. Abbott, S. Abraham, et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. ApJ, 896(2):L44, June 2020.

[98] R. Abbott et al. Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences. *Astrophys. J. Lett.*, 915(1):L5, 2021.

[99] Clifford E. Rhoades, Jr. and Remo Ruffini. Maximum mass of a neutron star. *Phys. Rev. Lett.*, 32:324–327, 1974.

[100] Vassiliki Kalogera and Gordon Baym. The maximum mass of a neutron star. *Astrophys. J. Lett.*, 470:L61–L64, 1996.

[101] John Antoniadis, Thomas M. Tauris, Feryal Ozel, Ewan Barr, David J. Champion, and Paulo C. C. Freire. The millisecond pulsar mass distribution: Evidence for bimodality and constraints on the maximum neutron star mass. *arXiv e-prints*, 5 2016.

[102] Philippe Landry and Jocelyn S. Read. The Mass Distribution of Neutron Stars in Gravitational-Wave Binaries. *arXiv e-prints*, 7 2021.

[103] Chris Pankow. On GW170817 and the Galactic Binary Neutron Star Population. *Astrophys. J.*, 866(1):60, 2018.

[104] Xingjiang Zhu, Eric Thrane, Stefan Oslowski, Yuri Levin, and Paul D. Lasky. Inferring the population properties of binary neutron stars with gravitational-wave measurements of spin. *Phys. Rev. D*, 98:043002, 2018.

[105] Nicholas Farrow, Xing-Jiang Zhu, and Eric Thrane. The mass distribution of Galactic double neutron stars. *Astrophys. J.*, 876(1):18, 2019.

[106] Ondrej Pejcha, Todd A. Thompson, and Christopher S. Kochanek. The observed neutron star mass distribution as a probe of the supernova explosion mechanism. *Mon. Not. Roy. Astron. Soc.*, 424:1570, 2012.

[107] T. M. Tauris et al. Formation of Double Neutron Star Systems. *Astrophys. J.*, 846(2):170, 2017.

[108] Ryan S. Lynch, Paulo C. C. Freire, Scott M. Ransom, and Bryan A. Jacoby. The Timing of Nine Globular Cluster Pulsars. *Astrophys. J.*, 745:109, 2012.

[109] Shanika Galaudage, Christian Adamcewicz, Xing-Jiang Zhu, Simon Stevenson, and Eric Thrane. Heavy Double Neutron Stars: Birth, Midlife, and Death. *Astrophys. J. Lett.*, 909(2):L19, 2021.

[110] G. Nelemans, L. R. Yungelson, S. F. Portegies Zwart, and F. Verbunt. Population synthesis for double white dwarfs . I. Close detached systems. A&A, 365:491–507, January 2001.

[111] Neil Cornish and Travis Robson. Galactic binary science with the new LISA design. In *Journal of Physics Conference Series*, volume 840 of *Journal of Physics Conference Series*, page 012024, May 2017.

[112] Martijn J. C. Wilhelm, Valeriya Korol, Elena M. Rossi, and Elena D'Onghia. The Milky Way's bar structural properties from gravitational waves. MNRAS, 500(4):4958–4971, January 2021.

[113] Shenghua Yu and C. Simon Jeffery. The influence of star formation history on the gravitational wave signal from close double degenerates in the thin disc. *Mon. Not. Roy. Astron. Soc.*, 429:1602, 2013.

[114] Valeriya Korol, Elena Maria Rossi, and Enrico Barausse. Constraining the Milky Way potential with Double White Dwarfs. *arXiv e-prints*, 10 2018.

[115] B. McKernan and K. E. S. Ford. On the resonant detonation of sub-Chandrasekhar mass white dwarfs during binary inspiral. *Mon. Not. Roy. Astron. Soc.*, 463(2):2039–2045, 2016.

[116] Simone Dall'Osso and Elena Maria Rossi. Constraining white dwarf viscosity through tidal heating in detached binary systems. *Mon. Not. Roy. Astron. Soc.*, 443(2):1057–1064, 2014.

[117] Valeriya Korol, Elena M. Rossi, Paul J. Groot, Gijs Nelemans, Silvia Toonen, and Anthony G. A. Brown. Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA. *Mon. Not. Roy. Astron. Soc.*, 470(2):1894–1910, 2017.

[118] Astrid Lamberts, Sarah Blunt, Tyson B Littenberg, Shea Garrison-Kimmel, Thomas Kupfer, and Robyn E Sanderson. Predicting the LISA white dwarf binary population in the Milky Way with cosmological simulations. *MNRAS*, 10 2019. stz2834.

[119] B. P. Abbott et al. GW150914: Implications for the stochastic gravitational wave background from binary black holes. *Phys. Rev. Lett.*, 116(13):131102, 2016.

[120] B. P. Abbott, R. Abbott, et al. GW170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences. PRL, 120(9):091101, March 2018.

[121] E. S. Phinney. A Practical Theorem on Gravitational Wave Backgrounds. *arXiv e-prints*, August 2001.

[122] Matthew R. Adams and Neil J. Cornish. Discriminating between a Stochastic Gravitational Wave Background and Instrument Noise. *Phys. Rev. D*, 82:022002, 2010.

[123] Joseph D. Romano and Neil. J. Cornish. Detection methods for stochastic gravitational-wave backgrounds: a unified treatment. *Living Reviews in Relativity*, 20(1):2, April 2017.

[124] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Search for the isotropic stochastic background using data from Advanced LIGO's second observing run. PRD, 100(6):061101, September 2019.

[125] R. Abbott et al. Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo's third observing run. *Phys. Rev. D*, 104(2):022004, 2021.

[126] Thomas Callister, Maya Fishbach, Daniel Holz, and Will Farr. Shouts and Murmurs: Combining Individual Gravitational-Wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers. *Astrophys. J. Lett.*, 896(2):L32, 2020.

[127] Suvodip Mukherjee, Tom Broadhurst, Jose M. Diego, Joseph Silk, and George F. Smoot. Inferring the lensing rate of LIGO-Virgo sources from the stochastic gravitational wave background. *Mon. Not. Roy. Astron. Soc.*, 501(2):2451–2466, 2021.

[128] Simone S. Bavera, Gabriele Franciolini, Giulia Cusin, Antonio Riotto, Michael Zevin, and Tassos Fragos. Stochastic gravitational-wave background as a tool to investigate multi-channel astrophysical and primordial black-hole mergers. *arXiv e-prints*, 9 2021.

[129] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. , 11 2021.

[130] Richard A. Redner and Homer F. Walker. Mixture densities, maximum likelihood and the em algorithm. *SIAM Review*, 26(2):195–239, 1984.

[131] Jean Diebolt and Christian P. Robert. Estimation of finite mixture distributions through bayesian sampling. *Journal of the Royal Statistical Society. Series B (Methodological)*, 56(2):363–375, 1994.

[132] Sylvia. Richardson and Peter J. Green. On bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(4):731–792, 1997.

[133] Matthew Stephens. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(4):795–809, 2000.

[134] A. Jasra, C. C. Holmes, and D. A. Stephens. Markov chain monte carlo methods and the label switching problem in bayesian mixture modeling. *Statist. Sci.*, 20(1):50–67, 02 2005.

[135] Merrilee Hurn, Ana Justel, and Christian P Robert. Estimating mixtures of regressions. *Journal of Computational and Graphical Statistics*, 12(1):55–79, 2003.

[136] Sylvia Fruhwirth-Schnatter. Markov chain monte carlo estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association*, 96(453):194–209, 2001.

[137] J. Skilling. Nested Sampling. In R. Fischer, R. Preuss, and U. V. Toussaint, editors, *American Institute of Physics Conference Series*, volume 735 of *American Institute of Physics Conference Series*, pages 395–405, November 2004.

[138] W. J. Handley, M. P. Hobson, and A. N. Lasenby. POLYCHORD: next-generation nested sampling. MNRAS, 453:4384–4398, November 2015.

[139] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts. Equation of state sensitivities when inferring neutron star and dense matter properties. MNRAS, 485(4):5363–5376, Jun 2019.

[140] W. J. Handley, M. P. Hobson, and A. N. Lasenby. Polychord 1.16. `https://github.com/PolyChord/PolyChordLite`, commit=cf2e888, 2019.

[141] Will J. Handley, Anthony N. Lasenby, Hiranya V. Peiris, and Michael P. Hobson. Bayesian inflationary reconstructions from Planck 2018 data. *arXiv e-prints*, page arXiv:1908.00906, Aug 2019.

[142] Edward Higson, Will Handley, Michael Hobson, and Anthony Lasenby. Bayesian sparse reconstruction: a brute-force approach to astronomical imaging and machine learning. MNRAS, 483(4):4828–4846, Mar 2019.

[143] Planck Collaboration, Y. Akrami, F. Arroja, M. Ashdown, et al. Planck 2018 results. X. Constraints on inflation. *arXiv e-prints*, page arXiv:1807.06211, Jul 2018.

[144] Richard D. Hall, Samantha J. Thompson, Will Handley, and Didier Queloz. On the Feasibility of Intense Radial Velocity Surveys for Earth-Twin Discoveries. MNRAS, 479(3):2968–2987, Sep 2018.

[145] Marius Millea and François Bouchet. Cosmic microwave background constraints in light of priors over reionization histories. A&A, 617:A96, Sep 2018.

[146] Malak Olamaie, Michael P. Hobson, Farhan Feroz, Keith J. B. Grainge, Anthony Lasenby, Yvette C. Perrott, Clare Rumsey, and Richard D. E. Saunders. Free-form modelling of galaxy clusters: a Bayesian and data-driven approach. MNRAS, 481(3):3853–3864, Dec 2018.

[147] F. Finelli, M. Bucher, A. Achúcarro, et al. Exploring cosmic origins with CORE: Inflation. J. Cosmology Astropart. Phys., 2018(4):016, Apr 2018.

[148] S. Hee, J. A. Vázquez, W. J. Handley, M. P. Hobson, and A. N. Lasenby. Constraining the dark energy equation of state using Bayes theorem and the Kullback-Leibler divergence. MNRAS, 466(1):369–377, Apr 2017.

[149] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, et al. Planck 2015 results. XX. Constraints on inflation. A&A, 594:A20, Sep 2016.

[150] S. Hee, W. J. Handley, M. P. Hobson, and A. N. Lasenby. Bayesian model selection without evidences: application to the dark energy equation-of-state. MNRAS, 455(3):2461–2473, Jan 2016.

[151] F. Feroz, M. P. Hobson, and M. Bridges. MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics. MNRAS, 398(4):1601–1614, Oct 2009.

[152] Riccardo Buscicchio, Elinore Roebber, J. M. Goldstein, and C. J. Moore. Hypertriangulation map: Python implementation. https://doi.org/10.5281/zenodo.3351629, July 2019.

[153] W.J. Gillard and E.D. Kvasov. Lipschitz optimization methods for fitting a sum of damped sinusoids to a series of observations. *Statistics and its interface*, 10:59–70, 01 2017.

[154] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, 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. *arXiv e-prints*, page arXiv:1811.12907, Nov 2018.

[155] Surabhi Sachdev, Sarah Caudill, Heather Fong, et al. The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs. *arXiv e-prints*, January 2019.

[156] Samantha A. Usman, Alexander H. Nitz, Ian W. Harry, et al. The PyCBC search for gravitational waves from compact binary coalescence. *CQG*, 33(21):215004, November 2016.

[157] Shaun Hooper, Shin Kee Chung, Jing Luan, David Blair, Yanbei Chen, and Linqing Wen. Summed Parallel Infinite Impulse Response (SPIIR) Filters For Low-Latency Gravitational Wave Detection. *Physical Review D*, 86(2):024012, aug 2011.

[158] T Adams, D Buskulic, V Germain, G M Guidi, F Marion, M Montani, B Mours, F Piergiovanni, and G Wang. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Classical and Quantum Gravity*, 33(17):175012, dec 2015.

[159] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher. A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 25(11):114029, Jun 2008.

[160] Gracedb: Gravitational-wave candidate event database, https://gracedb.ligo.org/.

[161] Sam Cooper, Aaron Jones, Sam Morrell, and George Smetana. Chirp, http://chirp.sr.bham.ac.uk/.

[162] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *LRR*, 21(1):3, April 2018.

[163] Alejandro Bohé, Lijing Shao, Andrea Taracchini, Alessandra Buonanno, Stanislav Babak, Ian W. Harry, Ian Hinder, Serguei Ossokine, Michael Pürrer, and Vivien Raymond. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. PRD, 95(4):044028, Feb 2017.

[164] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. PRD, 93(4):044006, Feb 2016.

[165] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. PRD, 93(4):044007, Feb 2016.

[166] Tito Dal Canton and Ian W. Harry. Designing a template bank to observe compact binary coalescences in Advanced LIGO's second observing run. *arXiv e-prints*, page arXiv:1705.01845, May 2017.

[167] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, and M. Coughlin. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. PRD, 91(4):042003, Feb 2015.

[168] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, et al. Tests of General Relativity with GW170817. *arXiv e-prints*, page arXiv:1811.00364, Nov 2018.

[169] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and V. B. Adya. Properties of the Binary Neutron Star Merger GW170817. *Physical Review X*, 9(1):011001, Jan 2019.

[170] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, et al. Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1. PRD, 100(10):104036, November 2019.

[171] Maximiliano Isi, Matthew Giesler, Will M. Farr, Mark A. Scheel, and Saul A. Teukolsky. Testing the no-hair theorem with GW150914. *arXiv e-prints*, page arXiv:1905.00869, May 2019.

[172] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and V. B. Adya. A gravitational-wave standard siren measurement of the Hubble constant. Nature, 551(7678):85–88, Nov 2017.

[173] Shasvath J. Kapadia, Sarah Caudill, Jolien D. E. Creighton, Will M. Farr, Gregory Mendell, Alan Weinstein, Kipp Cannon, Heather Fong, Patrick Godwin, and Rico K. L. Lo. A self-consistent method to estimate the rate of compact binary coalescences with a Poisson mixture model. *arXiv e-prints*, page arXiv:1903.06881, Mar 2019.

[174] Colm Talbot and Eric Thrane. Measuring the Binary Black Hole Mass Spectrum with an Astrophysically Motivated Parameterization. ApJ, 856(2):173, April 2018.

[175] Daniel Wysocki. Inferences about the distribution, merger rate, and evolutionary processes of compact binaries from gravitational wave observations. *arXiv e-prints*, page arXiv:1712.02643, Dec 2017.

[176] Daniel Wysocki, Jacob Lange, and Richard O'Shaughnessy. Reconstructing phenomenological distributions of compact binaries via gravitational wave observations. *arXiv e-prints*, page arXiv:1805.06442, May 2018.

[177] Simon Stevenson, Christopher P. L. Berry, and Ilya Mandel. Hierarchical analysis of gravitational-wave measurements of binary black hole spin-orbit misalignments. MNRAS, 471(3):2801–2811, Nov 2017.

[178] Javier Roulet and Matias Zaldarriaga. Constraints on binary black hole populations from LIGO-Virgo detections. MNRAS, 484(3):4216–4229, April 2019.

[179] Yang Bai, Vernon Barger, and Sida Lu. Measuring the Black Hole Mass Spectrum from Redshifts of aLIGO Binary Merger Events. *arXiv e-prints*, page arXiv:1802.04909, Feb 2018.

[180] W. Del Pozzo, C. P. L. Berry, A. Ghosh, T. S. F. Haines, L. P. Singer, and A. Vecchio. Dirichlet process Gaussian-mixture model: An application to localizing coalescing binary neutron stars with gravitational-wave observations. MNRAS, 479(1):601–614, Sep 2018.

[181] Jade Powell, Simon Stevenson, Ilya Mandel, and Peter Tino. Unmodelled Clustering Methods for Gravitational Wave Populations of Compact Binary Mergers. *arXiv e-prints*, page arXiv:1905.04825, May 2019.

[182] David Keitel. Galactic double neutron star total masses and Gaussian mixture model selection. MNRAS, 485(2):1665–1674, May 2019.

[183] N.L. Hjort, C. Holmes, P. Müller, and S.G. Walker. *Bayesian Nonparametrics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2010.

[184] P. Orbanz and Y. W. Teh. Bayesian nonparametric models. In *Encyclopedia of Machine Learning*. Springer, 2010.

[185] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, and R. X. Adhikari. Observation of Gravitational Waves from a Binary Black Hole Merger. PRL, 116(6):061102, Feb 2016.

[186] Michele Vallisneri, Jonah Kanner, Roy Williams, Alan Weinstein, and Branson Stephens. The LIGO open science center. *Journal of Physics: Conference Series*, 610:012021, may 2015.

[187] Brendon J. Brewer and Daniel Foreman-Mackey. DNest4: Diffusive Nested Sampling in C++ and Python. *arXiv e-prints*, page arXiv:1606.03757, Jun 2016.

[188] Joshua S Speagle. dynesty: A Dynamic Nested Sampling Package for Estimating Bayesian Posteriors and Evidences. *arXiv e-prints*, page arXiv:1904.02180, Apr 2019.

[189] Daniel Foreman-Mackey, David W. Hogg, and Timothy D. Morton. Exoplanet Population Inference and the Abundance of Earth Analogs from Noisy, Incomplete Catalogs. ApJ, 795(1):64, Nov 2014.

[190] Will M. Farr, Ilya Mandel, Chris Aldridge, and Kirsty Stroud. The Occurrence of Earth-Like Planets Around Other Stars. *arXiv e-prints*, page arXiv:1412.4849, Dec 2014.

[191] John Veitch, Walter Del Pozzo, and Cody Matt Pitkin. johnveitch/cpnest: pypi release. https://doi.org/10.5281/zenodo.825456, July 2017.

[192] Riccardo Buscicchio, Elinore Roebber, J. M. Goldstein, and C. J. Moore. A novel solution to the label switching problem in bayesian analysis - supplementary material. https://doi.org/10.5281/zenodo.3243319, July 2019.

[193] A. Stroeer and A. Vecchio. The LISA verification binaries. *Classical and Quantum Gravity*, 23(19):S809–S817, Oct 2006.

[194] Massimo Tinto and Sanjeev V. Dhurandhar. Time-Delay Interferometry. *Living Rev. Relativ.*, 17(1):6, Aug 2014.

[195] Antoine Petiteau, Gérard Auger, Hubert Halloin, Olivier Jeannin, Eric Plagnol, Sophie Pireaux, Tania Regimbau, and Jean-Yves Vinet. LISACode: A scientific simulator of LISA. PRD, 77(2):023002, Jan 2008.

[196] Peter D. Welch. The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *Audio and Electroacoustics, IEEE Transactions on*, 15:70 – 73, 07 1967.

[197] M. S. Bartlett. Periodogram Analysis and continuous spectra. *Biometrika*, 37(1-2):1–16, 06 1950.

[198] Alexander Stroeer and John Veitch. Bayesian approach to the study of white dwarf binaries in LISA data: The application of a reversible jump Markov chain MonteCarlo method. PRD, 80(6):064032, Sep 2009.

[199] Jeff Crowder and Neil J. Cornish. Extracting galactic binary signals from the first round of Mock LISA Data Challenges. *Classical and Quantum Gravity*, 24(19):S575–S585, Oct 2007.

[200] Neil J. Cornish and Shane L. Larson. Lisa data analysis: Source identification and subtraction. *Phys. Rev. D*, 67:103001, May 2003.

[201] Elinore Roebber, Riccardo Buscicchio, Alberto Vecchio, C. J. Moore, Antoine Klein, Valeriya Korol, Silvia Toonen, Davide Gerosa, T. E. Woods, S. M. Gaebel, and J. M. Goldstein. Milky way satellites shining bright in gravitational waves: dataset release. https://doi.org/10.5281/zenodo.3668905, February 2020.

[202] James S. Bullock and Michael Boylan-Kolchin. Small-Scale Challenges to the ΛCDM Paradigm. ARA&A, 55(1):343–387, Aug 2017.

[203] Harlow Shapley. A Stellar System of a New Type. *Harvard College Observatory Bulletin*, 908:1–11, Mar 1938.

[204] Gaia Collaboration. The Gaia mission. A&A, 595:A1, November 2016.

[205] Joshua D. Simon. The Faintest Dwarf Galaxies. ARA&A, 57:375–415, Aug 2019.

[206] Željko Ivezić, Steven M. Kahn, J. Anthony Tyson, et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. ApJ, 873(2):111, Mar 2019.

[207] Charlie Conroy and James S. Bullock. Beacons in the Dark: Using Novae and Supernovae to Detect Dwarf Galaxies in the Local Universe. ApJ, 805(1):L2, May 2015.

[208] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. , Feb 2017.

[209] Valeria Korol, Orlin Koop, and Elena M. Rossi. Detectability of Double White Dwarfs in the Local Group with LISA. ApJ, 866(2):L20, Oct 2018.

[210] J. Greiner, G. Hasinger, and H. C. Thomas. Another supersoft X-ray source in the LMC. A&A, 281:L61–L64, Jan 1994.

[211] A. van Teeseling, K. Reinsch, F. V. Hessman, and K. Beuermann. RX J0439.8-6809: a double-degenerate supersoft X-ray source? A&A, 323:L41–L44, Jul 1997.

[212] K. Werner and T. Rauch. Analysis of HST/COS spectra of the bare C-O stellar core <ASTROBJ>H1504+65</ASTROBJ> and a high-velocity twin in the Galactic halo. A&A, 584:A19, Dec 2015.

[213] Alan W. McConnachie. The Observed Properties of Dwarf Galaxies in and around the Local Group. AJ, 144(1):4, Jul 2012.

[214] David O. Cook, Daniel A. Dale, Benjamin D. Johnson, Liese Van Zee, Janice C. Lee, Robert C. Kennicutt, Daniela Calzetti, Shawn M. Staudaher, and Charles W. Engelbracht. The Spitzer Local Volume Legacy (LVL) global optical photometry. MNRAS, 445(1):881–889, Nov 2014.

[215] V. Korol, S. Toonen, A. Klein, V. Belokurov, F. Vincenzo, R. Buscicchio, D. Gerosa, C. J. Moore, E. Roebber, E. M. Rossi, and A. Vecchio. Prospects for detecting milky way satellites with lisa: Modelling gw population. A&A under review, 2020.

[216] ESA. LISA science requirements document. Technical Report ESA-L3-EST-SCI-RS-001, ESA, May 2018. www.cosmos.esa.int/web/lisa/lisa-documents/.

[217] S. F. Portegies Zwart and F. Verbunt. Population synthesis of high-mass binaries. A&A, 309:179–196, May 1996.

[218] S. Toonen, G. Nelemans, and S. Portegies Zwart. Supernova Type Ia progenitors from merging double white dwarfs. Using a new population synthesis model. A&A, 546:A70, October 2012.

[219] S. Toonen, M. Hollands, B. T. Gänsicke, and T. Boekholt. The binarity of the local white dwarf population. A&A, 602:A16, June 2017.

[220] Thomas M. Brown, Jason Tumlinson, Marla Geha, Joshua D. Simon, Luis C. Vargas, Don A. VandenBerg, Evan N. Kirby, Jason S. Kalirai, Roberto J. Avila, Mario Gennaro, Henry C. Ferguson, Ricardo R. Muñoz, Puragra Guhathakurta, and Alvio Renzini. The Quenching of the Ultra-faint Dwarf Galaxies in the Reionization Era. ApJ, 796(2):91, Dec 2014.

[221] Daniel R. Weisz, Andrew E. Dolphin, Evan D. Skillman, Jon Holtzman, Karoline M. Gilbert, Julianne J. Dalcanton, and Benjamin F. Williams. The Star Formation Histories of Local Group Dwarf Galaxies. I. Hubble Space Telescope/Wide Field Planetary Camera 2 Observations. ApJ, 789(2):147, Jul 2014.

[222] Daniel R. Weisz, Nicolas F. Martin, Andrew E. Dolphin, Saundra M. Albers, Michelle L. M. Collins, Annette M. N. Ferguson, Geraint F. Lewis, Dougal Mackey, Alan McConnachie, R. Michael Rich, and Evan D. Skillman. Comparing the Quenching Times of Faint M31 and Milky Way Satellite Galaxies. ApJ, 885(1):L8, Nov 2019.

[223] P. Kroupa, C. A. Tout, and G. Gilmore. The distribution of low-mass stars in the Galactic disc. MNRAS, 262:545–587, June 1993.

[224] Kevin B. Burdge, Michael W. Coughlin, Jim Fuller, Thomas Kupfer, Eric C. Bellm, Lars Bildsten, Matthew J. Graham, David L. Kaplan, Jan van Roestel, Richard G. Dekany, Dmitry A. Duev, Michael Feeney, Matteo Giomi, George Helou, Stephen Kaye, Russ R. Laher, Ashish A. Mahabal, Frank J. Masci, Reed Riddle, David L. Shupe, Maayane T. Soumagnac, Roger M. Smith, Paula Szkody,

Richard Walters, S. R. Kulkarni, and Thomas A. Prince. General relativistic orbital decay in a seven-minute-orbital-period eclipsing binary system. Nature, 571(7766):528–531, July 2019.

[225] Valeriya Korol, Elena M. Rossi, Paul J. Groot, Gijs Nelemans, Silvia Toonen, and Anthony G. A. Brown. Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA. MNRAS, 470(2):1894–1910, Sep 2017.

[226] Travis Robson and Neil Cornish. Impact of galactic foreground characterization on a global analysis for the LISA gravitational wave observatory. *Classical and Quantum Gravity*, 34(24):244002, December 2017.

[227] Antoine Petiteau, Gérard Auger, Hubert Halloin, Olivier Jeannin, Eric Plagnol, Sophie Pireaux, Tania Regimbau, and Jean-Yves Vinet. LISACode: A scientific simulator of LISA. PRD, 77(2):023002, January 2008.

[228] Stanislav Babak, Jonathan Gair, Alberto Sesana, Enrico Barausse, Carlos F. Sopuerta, Christopher P. L. Berry, Emanuele Berti, Pau Amaro-Seoane, Antoine Petiteau, and Antoine Klein. Science with the space-based interferometer LISA. V: Extreme mass-ratio inspirals. *Phys. Rev.*, D95(10):103012, 2017.

[229] Thomas A. Prince, Massimo Tinto, Shane L. Larson, and J. W. Armstrong. Lisa optimal sensitivity. *Phys. Rev. D*, 66:122002, Dec 2002.

[230] Kyle Kremer, Katelyn Breivik, Shane L. Larson, and Vassiliki Kalogera. Accreting Double White Dwarf Binaries: Implications for LISA. ApJ, 846(2):95, Sep 2017.

[231] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart. Short-period AM CVn systems as optical, X-ray and gravitational-wave sources. MNRAS, 349(1):181–192, March 2004.

[232] Jeff Crowder and Neil J. Cornish. Solution to the galactic foreground problem for lisa. *Phys. Rev. D*, 75:043008, Feb 2007.

[233] Arkadiusz Błaut, Stanislav Babak, and Andrzej Królak. Mock lisa data challenge for the galactic white dwarf binaries. *Phys. Rev. D*, 81:063008, Mar 2010.

[234] Valeriya Korol, Elena M. Rossi, and Enrico Barausse. A multimessenger study of the Milky Way's stellar disc and bulge with LISA, Gaia, and LSST. MNRAS, 483(4):5518–5533, Mar 2019.

[235] G. Iorio, V. Belokurov, D. Erkal, S. E. Koposov, C. Nipoti, and F. Fraternali. The first all-sky view of the Milky Way stellar halo with Gaia+2MASS RR Lyrae. MNRAS, 474(2):2142–2166, Feb 2018.

[236] J. Ted Mackereth, Jo Bovy, Henry W. Leung, Ricardo P. Schiavon, Wilma H. Trick, William J. Chaplin, Katia Cunha, Diane K. Feuillet, Steven R. Majewski, Marie Martig, Andrea Miglio, David Nidever, Marc H. Pinsonneault, Victor Silva Aguirre, Jennifer Sobeck, Jamie Tayar, and Gail Zasowski. Dynamical heating across the Milky Way disc using APOGEE and Gaia. MNRAS, 489(1):176–195, Oct 2019.

[237] G. Torrealba, V. Belokurov, S. E. Koposov, T. S. Li, M. G. Walker, J. L. Sanders, A. Geringer-Sameth, D. B. Zucker, K. Kuehn, N. W. Evans, and W. Dehnen. The hidden giant: discovery of an enormous Galactic dwarf satellite in Gaia DR2. MNRAS, 488(2):2743–2766, Sep 2019.

[238] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. PRX, 9(3):031040, July 2019.

[239] R. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run. *Physical Review X*, 11(2):021053, April 2021.

[240] Alberto Sesana. Prospects for Multiband Gravitational-Wave Astronomy after GW150914. PRL, 116(23):231102, June 2016.

[241] Koutarou Kyutoku and Naoki Seto. Concise estimate of the expected number of detections for stellar-mass binary black holes by eLISA. MNRAS, 462(2):2177–2183, October 2016.

[242] Kaze W. K. Wong, Ely D. Kovetz, Curt Cutler, and Emanuele Berti. Expanding the LISA Horizon from the Ground. PRL, 121(25):251102, December 2018.

[243] Christopher J. Moore, Davide Gerosa, and Antoine Klein. Are stellar-mass black-hole binaries too quiet for LISA? MNRAS, 488(1):L94–L98, September 2019.

[244] Katelyn Breivik, Carl L. Rodriguez, Shane L. Larson, Vassiliki Kalogera, and Frederic A. Rasio. Distinguishing between Formation Channels for Binary Black Holes with LISA. ApJ, 830(1):L18, October 2016.

[245] Johan Samsing and Daniel J. D'Orazio. Black Hole Mergers From Globular Clusters Observable by LISA I: Eccentric Sources Originating From Relativistic N-body Dynamics. MNRAS, 481(4):5445–5450, December 2018.

[246] Davide Gerosa, Sizheng Ma, Kaze W. K. Wong, Emanuele Berti, Richard O'Shaughnessy, Yanbei Chen, and Krzysztof Belczynski. Multiband gravitational-wave event rates and stellar physics. PRD, 99(10):103004, May 2019.

[247] Katie Chamberlain and Nicolás Yunes. Theoretical physics implications of gravitational wave observation with future detectors. PRD, 96(8):084039, October 2017.

[248] Giuseppe Gnocchi, Andrea Maselli, Tiziano Abdelsalhin, Nicola Giacobbo, and Michela Mapelli. Bounding alternative theories of gravity with multiband GW observations. PRD, 100(6):064024, September 2019.

[249] Rhondale Tso, Davide Gerosa, and Yanbei Chen. Optimizing LIGO with LISA forewarnings to improve black-hole spectroscopy. PRD, 99(12):124043, June 2019.

[250] Alexandre Toubiana, Sylvain Marsat, Stanislav Babak, Enrico Barausse, and John Baker. Tests of general relativity with stellar-mass black hole binaries observed by LISA. PRD, 101(10):104038, May 2020.

[251] Salvatore Vitale. Multiband Gravitational-Wave Astronomy: Parameter Estimation and Tests of General Relativity with Space- and Ground-Based Detectors. PRL, 117(5):051102, July 2016.

[252] Sean McGee, Alberto Sesana, and Alberto Vecchio. Linking gravitational waves and X-ray phenomena with joint LISA and Athena observations. *Nature Astronomy*, 4:26–31, January 2020.

[253] Alexandre Toubiana, Sylvain Marsat, Stanislav Babak, John Baker, and Tito Dal Canton. Parameter estimation of stellar-mass black hole binaries with LISA. PRD, 102(12):124037, December 2020.

[254] Curt Cutler. Angular resolution of the LISA gravitational wave detector. PRD, 57(12):7089–7102, June 1998.

[255] Travis Robson, Neil J. Cornish, and Chang Liu. The construction and use of LISA sensitivity curves. CQG, 36(10):105011, May 2019.

[256] Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. LRR, 17(1):2, December 2014.

[257] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries. PRD, 49(12):6274–6297, June 1994.

[258] Antoine Klein. EFPE: Efficient fully precessing eccentric gravitational waveforms for binaries with long inspirals. *arXiv e-prints*, page arXiv:2106.10291, June 2021.

[259] Neil J. Cornish. Detection strategies for extreme mass ratio inspirals. *Classical and Quantum Gravity*, 28(9):094016, May 2011.

[260] Alvin J. K. Chua, Christopher J. Moore, and Jonathan R. Gair. Augmented kludge waveforms for detecting extreme-mass-ratio inspirals. PRD, 96(4):044005, August 2017.

[261] Stanislav Babak, Jonathan Gair, Alberto Sesana, Enrico Barausse, and et al. Science with the space-based interferometer LISA. V. Extreme mass-ratio inspirals. PRD, 95(10):103012, May 2017.

[262] Adam Pound, Barry Wardell, Niels Warburton, and Jeremy Miller. Second-Order Self-Force Calculation of Gravitational Binding Energy in Compact Binaries. PRL, 124(2):021101, January 2020.

[263] Alvin J. K. Chua, Michael L. Katz, Niels Warburton, and Scott A. Hughes. Rapid Generation of Fully Relativistic Extreme-Mass-Ratio-Inspiral Waveform Templates for LISA Data Analysis. PRL, 126(5):051102, February 2021.

[264] Maarten van de Meent and Niels Warburton. Fast self-forced inspirals. *Classical and Quantum Gravity*, 35(14):144003, July 2018.

[265] Thomas A. Prince, Massimo Tinto, Shane L. Larson, and J. W. Armstrong. LISA optimal sensitivity. PRD, 66(12):122002, December 2002.

[266] Jean-Baptiste Bayle, Marc Lilley, Antoine Petiteau, and Hubert Halloin. Effect of filters on the time-delay interferometry residual laser noise for LISA. PRD, 99(8):084023, April 2019.

[267] Michele Vallisneri, Jean-Baptiste Bayle, Stanislav Babak, and Antoine Petiteau. Time-delay interferometry without delays. PRD, 103(8):082001, April 2021.

[268] Jean-Baptiste Bayle, Olaf Hartwig, and Martin Staab. Adapting time-delay interferometry for LISA data in frequency. PRD, 104(2):023006, July 2021.

[269] Olaf Hartwig and Jean-Baptiste Bayle. Clock-jitter reduction in LISA time-delay interferometry combinations. PRD, 103(12):123027, June 2021.

[270] Curt Cutler and Éanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform\? PRD, 49(6):2658–2697, March 1994.

[271] John Veitch, Walter Del Pozzo, et al. johnveitch/cpnest: v0.11.3. , January 2021.

[272] Antoine Klein, Yannick Boetzel, Achamveedu Gopakumar, Philippe Jetzer, and Lorenzo de Vittori. Fourier domain gravitational waveforms for precessing eccentric binaries. PRD, 98(10):104043, November 2018.

[273] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. Comparison of search templates for gravitational waves from binary inspiral. PRD, 63(4):044023, February 2001.

[274] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. Erratum: Comparison of search templates for gravitational waves from binary inspiral: 3.5PN update [Phys. Rev. D 66, 027502 (2002)]. PRD, 72(2):029901, July 2005.

[275] Alberto Mangiagli, Antoine Klein, Alberto Sesana, Enrico Barausse, and Monica Colpi. Post-Newtonian phase accuracy requirements for stellar black hole binaries with LISA. PRD, 99(6):064056, March 2019.

[276] Sylvain Marsat and John G. Baker. Fourier-domain modulations and delays of gravitational-wave signals. , June 2018.

[277] Daniel A. Shaddock. Operating LISA as a Sagnac interferometer. PRD, 69(2):022001, January 2004.

[278] Louis J. Rubbo, Neil J. Cornish, and Olivier Poujade. Forward modeling of space-borne gravitational wave detectors. PRD, 69(8):082003, April 2004.

[279] C. W. Clenshaw and A. R. Curtis. A method for numerical integration on an automatic computer. *Numer. Math.*, 2(1):197–205, December 1960.

[280] Sylvain Marsat, John G. Baker, and Tito Dal Canton. Exploring the Bayesian parameter estimation of binary black holes with LISA. PRD, 103(8):083011, April 2021.

[281] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. A&A, 641:A6, September 2020.

[282] Adam G. Riess, Stefano Casertano, Wenlong Yuan, J. Bradley Bowers, et al. Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with ΛCDM. ApJ, 908(1):L6, February 2021.

[283] Lawrence E. Kidder, Clifford M. Will, and Alan G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. PRD, 47(10):R4183–R4187, May 1993.

[284] R. Abbott et al. GW190521: A Binary Black Hole Merger with a Total Mass of 150 $M_\odot$. PRL, 125(10):101102, September 2020.

[285] Patricia Schmidt, Frank Ohme, and Mark Hannam. Towards models of gravitational waveforms from generic binaries: II. Modelling precession effects with a single effective precession parameter. PRD, 91(2):024043, January 2015.

[286] Davide Gerosa, Matthew Mould, Daria Gangardt, Patricia Schmidt, Geraint Pratten, and Lucy M. Thomas. A generalized precession parameter χ_p to interpret gravitational-wave data. PRD, 103(6):064067, March 2021.

[287] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, and et al. GW190412: Observation of a binary-black-hole coalescence with asymmetric masses. PRD, 102(4):043015, August 2020.

[288] Alexander H. Nitz, Thomas Dent, Gareth S. Davies, et al. 2-OGC: Open Gravitational-wave Catalog of Binary Mergers from Analysis of Public Advanced LIGO and Virgo Data. ApJ, 891(2):123, March 2020.

[289] Tejaswi Venumadhav, Barak Zackay, Javier Roulet, Liang Dai, and Matias Zaldarriaga. New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo. PRD, 101(8):083030, April 2020.

[290] Barak Zackay, Liang Dai, Tejaswi Venumadhav, et al. Detecting Gravitational Waves With Disparate Detector Responses: Two New Binary Black Hole Mergers. *arXiv e-prints*, October 2019.

[291] Maya Fishbach and Daniel E. Holz. Where Are LIGO's Big Black Holes? ApJ, 851(2):L25, December 2017.

[292] Davide Gerosa and Emanuele Berti. Are merging black holes born from stellar collapse or previous mergers? PRD, 95(12):124046, June 2017.

[293] Chase Kimball, Colm Talbot, Christopher P. L. Berry, et al. Black hole genealogy: Identifying hierarchical mergers with gravitational waves. *arXiv e-prints*, April 2020.

[294] Colm Talbot and Eric Thrane. Determining the population properties of spinning black holes. PRD, 96(2):023012, July 2017.

[295] Maya Fishbach, Daniel E. Holz, and Will M. Farr. Does the Black Hole Merger Rate Evolve with Redshift? ApJ, 863(2):L41, August 2018.

[296] Davide Gerosa, Emanuele Berti, Richard O'Shaughnessy, Krzysztof Belczynski, and et al. Spin orientations of merging black holes formed from the evolution of stellar binaries. PRD, 98(8):084036, October 2018.

[297] Vishal Baibhav, Davide Gerosa, Emanuele Berti, Kaze W. K. Wong, and et al. The mass gap, the spin gap, and the origin of merging binary black holes. *arXiv e-prints*, page arXiv:2004.00650, April 2020.

[298] Will M. Farr, Maya Fishbach, Jiani Ye, et al. A Future Percent-level Measurement of the Hubble Expansion at Redshift 0.8 with Advanced LIGO. ApJ, 883(2):L42, October 2019.

[299] Nelson Christensen. Stochastic gravitational wave backgrounds. *Reports on Progress in Physics*, 82(1):016903, January 2019.

[300] Rory J. E. Smith, Colm Talbot, Francisco Hernand ez Vivanco, and Eric Thrane. Inferring the population properties of binary black holes from unresolved gravitational waves. MNRAS, 496(3):3281–3290, June 2020.

[301] Cody Messick, Kent Blackburn, Patrick Brady, et al. Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data. PRD, 95(4):042001, February 2017.

[302] S. Klimenko, G. Vedovato, M. Drago, et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. PRD, 93(4):042004, February 2016.

[303] Sebastian M. Gaebel, John Veitch, Thomas Dent, et al. Digging the population of compact binary mergers out of the noise. MNRAS, 484(3):4008–4023, April 2019.

[304] Tommaso Treu. Strong Lensing by Galaxies. ARA&A, 48:87–125, September 2010.

[305] Graham P. Smith, Christopher Berry, Matteo Bianconi, et al. Strong-lensing of Gravitational Waves by Galaxy Clusters. In *IAU Symposium*, volume 338, pages 98–102, October 2018.

[306] O. A. Hannuksela, K. Haris, K. K. Y. Ng, et al. Search for Gravitational Lensing Signatures in LIGO-Virgo Binary Black Hole Events. ApJ, 874(1):L2, March 2019.

[307] G. P. Smith, M. Bianconi, M. Jauzac, et al. Deep and rapid observations of strong-lensing galaxy clusters within the sky localization of GW170814. MNRAS, 485(4):5180–5191, June 2019.

[308] Graham P. Smith, Mathilde Jauzac, John Veitch, et al. What if LIGO's gravitational wave detections are strongly lensed by massive galaxy clusters? MNRAS, 475(3):3823–3828, April 2018.

[309] Liang Dai, Tejaswi Venumadhav, and Kris Sigurdson. Effect of lensing magnification on the apparent distribution of black hole mergers. PRD, 95(4):044011, February 2017.

[310] Ken K. Y. Ng, Kaze W. K. Wong, Tom Broadhurst, et al. Precise LIGO lensing rate predictions for binary black holes. PRD, 97(2):023012, January 2018.

[311] Shun-Sheng Li, Shude Mao, Yuetong Zhao, and Youjun Lu. Gravitational lensing of gravitational waves: a statistical perspective. MNRAS, 476(2):2220–2229, May 2018.

[312] Alvin K. Y. Li, Rico K. L. Lo, Surabhi Sachdev, C. L. Chan, et al. Targeted Subthreshold Search for Strongly-lensed Gravitational-wave Events. *arXiv e-prints*, April 2019.

[313] Tom Broadhurst, Jose M. Diego, and George F. Smoot. Reinterpreting Low Frequency LIGO/Virgo Events as Magnified Stellar-Mass Black Holes at Cosmological Distances. *arXiv e-prints*, February 2018.

[314] Tom Broadhurst, Jose M. Diego, and III Smoot, George F. Twin LIGO/Virgo Detections of a Viable Gravitationally-Lensed Black Hole Merger. *arXiv e-prints*, page arXiv:1901.03190, January 2019.

[315] T. Broadhurst, J. M. Diego, and G. F. Smoot. A Distant Origin For Magnified LIGO/Virgo Black Holes Implied By Binary Component Masses. *arXiv e-prints*, February 2020.

[316] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. Nature, 323(6086):310–311, September 1986.

[317] Daniel E. Holz and Scott A. Hughes. Using Gravitational-Wave Standard Sirens. ApJ, 629(1):15–22, August 2005.

[318] Hsin-Yu Chen, Maya Fishbach, and Daniel E. Holz. A two per cent Hubble constant measurement from standard sirens within five years. Nature, 562(7728):545–547, October 2018.

[319] B. P. Abbott, R. Abbott, et al. A gravitational-wave measurement of the Hubble constant following the second observing run of Advanced LIGO and Virgo. *arXiv e-prints*, August 2019.

[320] Stefan Hilbert, Simon D. M. White, Jan Hartlap, and Peter Schneider. Strong-lensing optical depths in a ΛCDM universe - II. The influence of the stellar mass in galaxies. MNRAS, 386(4):1845–1854, June 2008.

[321] Ryuichi Takahashi, Masamune Oguri, Masanori Sato, and Takashi Hamana. Probability Distribution Functions of Cosmological Lensing: Convergence, Shear, and Magnification. ApJ, 742(1):15, November 2011.

[322] Andrew Robertson, Graham P. Smith, Richard Massey, et al. What does strong gravitational lensing? The mass and redshift distribution of high-magnification lenses. MNRAS, May 2020.

[323] Piero Madau and Mark Dickinson. Cosmic Star-Formation History. ARA&A, 52:415–486, August 2014.

[324] Eric Thrane and Joseph D. Romano. Sensitivity curves for searches for gravitational-wave backgrounds. PRD, 88(12):124032, December 2013.

[325] D. Gerosa. Dgerosa/Gwdet: V0.1, September 2017.

[326] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. PRL, 119(16):161101, October 2017.

[327] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass ~ 3.4 $M_{\odot}$. ApJ, 892(1):L3, March 2020.

[328] Francisco Hernandez Vivanco, Rory Smith, Eric Thrane, et al. Accelerated detection of the binary neutron star gravitational-wave background. PRD, 100(4):043023, August 2019.

[329] Liang Dai, Barak Zackay, Tejaswi Venumadhav, et al. Search for Lensed Gravitational Waves Including Morse Phase Information: An Intriguing Candidate in O2. *arXiv e-prints*, July 2020.

[330] Peter T. H. Pang, Otto A. Hannuksela, Tim Dietrich, et al. Lensed or not lensed: determining lensing magnifications for binary neutron star mergers from a single detection. MNRAS, 495(4):3740–3750, May 2020.