



University
of Glasgow

Linley, Jethro (2022) *Probing the general dynamics of compact binary systems with multi-band gravitational-wave observations*. PhD thesis.

<https://theses.gla.ac.uk/83388/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

Probing the General Dynamics of Compact Binary Systems with Multi-Band Gravitational-Wave Observations



by Jethro Linley

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

Submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy

Supervisors: Prof. Graham Woan & Dr. John Veitch

July 2022

Abstract

There are two main goals of this thesis. The first is developing our so-called ‘downsampling’ procedure: an approximation technique for reproducing the *likelihood function* (in simulation) in relatively short time, of the sort of signals we expect to observe with the LISA Gravitational Wave (GW) detector, such as the inspiral part of binary black-hole mergers. The procedure is tested on a variety of signals and shown to provide highly accurate reproductions of the likelihood function up to a few thousand times faster than our estimates of convergence time considering the fastest type of analysis one could perform (frequency domain, stationary noise) using all the data points.

The second goal is to understand the features and effectiveness of using multi-band GW data in data analysis, in particular on expansion of the signal model to include anticipated physical effects that modify the low-frequency part of the waveform, such as acceleration and relativistic time-delays. In addition, we extend the model to cover two well-known modified gravity theories, to test the assertion that multi-band GW data analysis will provide strong theory constraints. Since the degree to which the various time-delays (which have their own parameter space) modify the waveform ranges from being negligible to significant and essential to model, we devise a formalism for splitting the time-delay parameter space into distinct regions that either fully model, ‘dimensionally reduce’, or neglect the time-delays from the model. The advantages of doing this are that one acquires posteriors that are considerably more informative, since the waveform is not being ‘over-modelled’.

The posteriors we obtain confirm what one would expect to see (particularly in terms of parameter degeneracies) after considering and comparing the effects on the GW phase that arise from varying different combinations of parameters. In the fully modelled region, it is possible to recover well-constrained time-delay parameters (i.e., the Keplerian orbital parameters, including supermassive black hole mass), but it is very difficult to

derive general results about the expected behaviours of posteriors since the time-delay functions are complicated functions of time. The behaviour of posteriors of the ‘parameter reduced’ models are far easier to predict and understand, but the new parameters are relegated to nuisance parameters. The modified gravity theories we analyse are not well-constrained simply by multi-band observations alone; very high SNR appears to be the more important factor, but even then, different aspects of the nature of the waveform modifications (depending on the theory) can lead to significant bias, or render the effects too weak for GW astronomy to provide any useful constraints.

By inspecting one-parameter families of posterior distributions (treating the posteriors as functions of some model parameter or signal property) and observing their structural evolution as that parameter is varied, we uncover and discuss some interesting features and behaviours of the distributions. Topics for future study are highlighted and include extensions and refinements of the waveform, population, and detector models, further studies of the Kepler parameter space division scheme, and posterior sampling issues to be addressed.

Declaration

The material presented in this thesis is the result of my own work, undertaken within the Institute for Gravitational Research, in the School of Physics and Astronomy at the University of Glasgow, and produced under the supervision of Prof. Graham Woan and Dr. John Veitch. This work has not been submitted for any other degree at any institution. Material from Chapters 3, 7 & 8 is planned to be reformatted as journal style papers and submitted for publication in the near future.

This research has been kindly supported by the Science and Technology Facilities Council (STFC). We are grateful for computational resources provided by Cardiff University, and funded by an STFC grant supporting UK Involvement in the Operation of Advanced LIGO.

Acknowledgements

I would like to extend a very warm thank you to all of the lovely IGR folk for being so friendly and welcoming. Thanks to Joe, Simon and Hunter of 427a, for a lot of help in the early days in understanding LSC, general PhD and computer related things, and for the office entertainment before the lockdown sent us home to work. Thanks to Matt Pitkin and Chris Messenger for some useful discussions along the way and especially Michael Williams for patiently helping me a great deal with configuring and operating his sampler, Nessai.

Extra special thanks of course go to my brilliant supervisors, Graham Woan and John Veitch, for their efforts in developing my understanding of the science and how to go about finding anchors to hang useful results from, in the seemingly endless set of possibilities that remains to be explored in this field. They have always been exceptionally considerate, interested (and interesting!), and have been an abundant source of guidance in more aspects than just the technical side of doing research. I have thoroughly enjoyed our varied discussions, both on the serious work side of things and on the more conceptual musings on the machinery of nature. I have learned a tremendous amount from them both and shall be eternally grateful.

Thank you to my family, especially my mother, who has consistently supported me in every way she possibly could, through all of my wild ventures into different disciplines. Finally, a big thank you to my partner, Rūta, for having kept me levelheaded and making sure I stepped outside every now and then.

Contents

1	Vacuum Wave Solutions in GR	21
1.1	Theoretical foundations	21
1.2	Wave equation in GR	23
1.3	Sources of GWs	28
1.3.1	GWs from accelerated mass	29
1.3.2	A GW classification formalism	31
1.3.2.1	Compact binaries	32
1.3.2.2	Continuous waves	35
1.3.2.3	Burst sources	36
1.3.2.4	Stochastic background	37
1.4	BHB waveform models	38
1.4.1	Approximants	39
1.4.2	BHB model parameters	40
2	GW Data & Analysis	42
2.1	Laser interferometers	43
2.1.1	Time-delay interferometry	44
2.1.2	Detector response and antenna patterns	46
2.1.3	Basic overview of detectors	47
2.1.3.1	LIGO	47
2.1.3.2	LISA	47
2.1.3.3	DECIGO	48
2.2	Noise	48
2.2.1	Sources	49
2.2.1.1	LISA noise sources	49

2.2.1.2	LIGO noise sources	49
2.2.2	Power spectral density	50
2.3	Parameter estimation	51
2.3.1	Bayesian inference	51
2.3.2	Inner products and transformations	52
2.3.3	Noise covariance matrix	56
2.3.4	Inverting large autocorrelation matrices	57
2.4	Sampling	60
2.4.1	Markov chain Monte-Carlo sampling	61
2.4.2	Nested sampling	61
3	The Large Dataset Problem	63
3.1	Introduction	63
3.2	Preliminaries	66
3.2.1	Whitening transformation matrix	66
3.2.2	Data vector spaces	67
3.2.3	Note on time vs. frequency domain	69
3.2.4	Divergences	70
3.3	Defining a new likelihood function	72
3.3.1	Downsampling	72
3.3.1.1	Downsampling schemes	72
3.3.1.2	Residuals and metrics	73
3.3.1.3	Time domain sample correlations	74
3.3.1.3.1	Maximum correlated samples	76
3.3.2	Likelihood approximation methods	77
3.3.3	Exact preservation of Fisher information	79
3.3.4	An optimal noise reduction factor	81
3.3.4.1	Deriving the optimal noise reduction factor	83
3.3.4.1.1	Primitive method	84
3.3.4.1.2	General Gaussian approximations	84
3.3.5	Summary	91
3.4	Testing the downsampling procedure	92
3.4.1	Posterior distribution noise sources	93

3.4.1.1	Sampler noise	93
3.4.1.2	Downsampling data sample selection noise	94
3.4.1.3	Physical noise	94
3.4.2	Determining closeness of distributions	95
3.4.2.1	Combined distance between marginal distributions	95
3.4.3	Convergence criteria	96
3.4.3.1	Noisy posterior distributions	96
3.4.3.2	Summary	98
3.4.4	The distances evaluated	98
3.5	Speed comparison and further remarks	104
3.5.1	Posterior estimate evaluation times	104
3.5.2	Discussion	107
4	The Low-Frequency Landscape	109
4.1	Waveforms and modifiers	110
4.1.1	Known-knowns	111
4.1.2	Known-unknowns	112
4.1.3	Unknown-unknowns	112
4.2	Establishing a foundational signal model	113
4.2.1	Binary formation scenarios	113
4.2.2	Orbits & time-delays	116
4.2.2.1	Orbital parameters	116
4.2.2.2	Roemer delay	119
4.2.2.3	Periastron precession	120
4.2.2.4	Relativistic time dilation	122
4.2.2.5	Shapiro delay	123
4.2.3	Summary and examples	124
4.2.4	Reduced dimensionality time-delay approximants	127
5	Quadratic Gravity Theories	138
5.1	Einstein-dilaton Gauss-Bonnet	139
5.2	Dynamical Chern-Simons	140
5.3	Small deviations From GR	141

5.4	Schematic frequency domain waveform	142
5.5	Phase shifted Fourier waveforms in time domain	144
5.5.1	Validity of time-domain approximants	152
5.6	Acceleration/modified gravity bias	154
6	Analysis Framework Details	157
6.1	Multi-band data analysis	157
6.1.1	Multi-band detector networks	158
6.1.2	Multi-band parameter estimation	159
6.2	Exploring the framework	160
6.2.1	Comparison of a posterior with its PCM approximation	160
6.2.2	Threshold SNRs	162
6.2.3	The large dimensionality/sampling problem	162
6.2.3.1	Reduced spin parameter model	162
6.2.3.2	Numerical marginalisation of phase in time-domain	163
6.2.4	Real-world data analysis: selection of priors	166
6.3	Modelling caveats	166
6.3.1	Missing parameters	167
6.3.1.1	Spin parameters	167
6.3.1.2	Eccentricity	168
6.3.1.3	SMBH spin	168
6.3.2	Omitted physical effects	169
6.3.2.1	Hubble expansion	169
6.3.2.2	Shapiro and geometric time-delay	169
6.3.3	Redshift and mass/distance scale degeneracy	169
7	Multi-band GW Astronomy Part I:	
	Isolated Binaries	172
7.1	Data analysis pipeline	174
7.2	‘Low-mass’ binary	175
7.2.1	GR binary	175
7.2.2	Quadratic gravity	182
7.2.2.1	Einstein-dilaton Gauss-Bonnet	182

7.2.2.2	Dynamical Chern-Simons	189
7.3	‘High-mass’ binary	192
7.3.1	GR binary	192
7.3.2	Quadratic gravity	194
7.3.2.1	Einstein-dilaton Gauss-Bonnet	194
7.3.2.2	Dynamical Chern-Simons	195
7.4	Prospects for multi-band GW data analysis	197
7.4.1	Population models, detection rates	197
7.4.2	Detector network summary	198
7.4.2.1	LISA	198
7.4.2.2	DECIGO	199
7.4.2.3	LIGO	199
7.4.2.4	LISA & LIGO	199
7.4.2.5	LISA & DECIGO	200
7.4.2.6	DECIGO & LIGO	200
7.4.2.7	LISA, DECIGO & LIGO	201
7.4.3	Outlook	201
8	Multi-band GW Astronomy Part II:	
	Binaries in Galactic Orbits	203
8.1	‘Low-mass’ binary with acceleration	204
8.1.1	Constant acceleration	206
8.1.2	Linear acceleration	208
8.2	‘High-mass’ binary with acceleration	210
8.2.1	Constant acceleration	210
8.2.2	Linear acceleration	212
8.3	Complete time-delay model	213
8.3.1	Nearby system	215
8.3.2	Distant system	217
8.4	Carrier signals modified at low-frequencies	219
8.5	Approximating the time-delays	223
8.6	Modified gravity with acceleration	228
8.7	Discussion	229

9 Final Remarks and Outlook	232
A Closed-form approximations of IADs	239
B Shapiro Delay	245

List of Figures

1.1	“Plus”-polarised plane gravitational wave	28
1.2	“Cross”-polarised plane gravitational wave	29
1.3	An example gravitational wave of a binary black-hole coalescence	33
1.4	Example continuous wave	35
1.5	A simple example of a possible burst GW.	36
1.6	An example stochastic GW background signal.	37
2.1	A simple, equal arm laser interferometer schematic diagram.	44
2.2	Schematic diagram of time-delay interferometer.	45
2.3	Various detector amplitude spectral densities.	51
2.4	An example power spectral density (PSD) function for LISA	59
2.5	The (one-sided) <i>inverse autocorrelation function</i> for the LISA PSD	60
3.1	The (one-sided) <i>whitening function</i> derived from the LISA PSD	68
3.2	Results of CJMS distances for example systems as function of downsam- pling rate.	101
3.3	Converging posteriors as time domain sample number is increased.	102
3.4	Converging Fisher matrices as fraction of time domain samples is increased.	103
3.5	Parameter estimation convergence times	105
4.1	Galaxy structure diagram	114
4.2	Depiction of the Tait-Bryan Euler angles	116
4.3	Depiction of the example orbits studied.	125
4.4	The various time-delays of the example orbits.	126
4.5	The integrated (ascribed) acceleration difference for example orbits as a function of semi-major axis.	130

4.6	Acceleration phase diagram for circular, Newtonian orbits	132
4.7	Approximate boundaries between proposed regions of Keplerian element space requiring different minimal models of time-delays	137
5.1	Modified gravity parameter phase diagram	144
6.1	Overlay of true posterior and ‘Fisher information matrix posterior’.	161
6.2	Example of potential variation in posterior estimation results	165
7.1	The marginal precision and accuracy plots for the low-mass, isolated GR binary.	179
7.2	Marginals of the posterior for the low-mass GR binary in LISA+LIGO as a function of the SNR in LISA.	180
7.3	The marginal precision and accuracy plots for the low-mass, isolated EdGB binary.	184
7.4	Corner-plot for the low-mass EdGB binary in LISA+LIGO.	185
7.5	Detailed evolution of the low-mass EdGB binary as a function of the LISA SNR.	187
7.6	Corner-plot for the low-mass EdGB binary in LISA+LIGO, marginalised onto $(\mathcal{M}_c, q, \chi_{\text{eff}}, \zeta_{\text{EdGB}})$, with LISA SNR=300.	188
7.7	Marginals of the posterior for the low-mass dCS binary in LISA+LIGO as a function of the SNR in LISA.	190
7.8	Marginal precision and accuracy plots for the low-mass dCS binary.	191
7.9	The marginal precision and accuracy plots for the high-mass, isolated GR binary.	193
7.10	The marginal precision and accuracy plots for the high-mass, isolated EdGB binary.	195
7.11	The marginal precision and accuracy plots for the high-mass, isolated dCS binary.	196
8.1	Marginals of the posterior of the low-mass system in LISA+LIGO as a function of (constant) line-of-sight acceleration of the binary.	205
8.2	The marginal precision and accuracy plots for the low-mass, constantly accelerated GR binary.	207

8.3	Corner plot for binary showing that DECIGO is only sensitive to a single acceleration parameter	208
8.4	Marginal precision and accuracy plots for the low-mass, linearly accelerated GR binary.	209
8.5	Marginals of the posterior of the high-mass system in LISA+LIGO as a function of (constant) line-of-sight acceleration.	211
8.6	Marginal precision and accuracy plots for the high-mass, constantly accelerated GR binary.	212
8.7	Marginal precision and accuracy plots for the high-mass, linearly accelerated GR binary.	213
8.8	Corner plot of fully modelled high-mass system on ‘nearby’ variant of ‘example orbit 1’.	216
8.9	Corner plot of fully modelled high-mass system on ‘distant’ variant of ‘example orbit 1’.	218
8.10	Comparison of the ‘confirmed future detector network’ Kepler parameter posterior with a potential future network with an extremely high SNR in ground-based detectors.	220
8.11	Superimposed posteriors (marginalised over period) of systems with ground-based detector measurable parameter posteriors approximated as a delta function, for orbits of various periods.	221
8.12	Marginal precision and accuracy plots for LISA posteriors of Kepler parameters (from known carrier wave signals) for different ‘pitch’, α	222
8.13	Marginals of the posteriors produced with fully modelled time-delay data and model.	224
8.14	Marginals of the posteriors produced with fully modelled time-delay data but with no time-delays present in the template model for parameter estimation	225
8.15	Marginals of the posteriors produced with fully modelled time-delay data, with time-delays in parameter estimation modelled as constant accelerations.	226
8.16	Marginals of the posteriors produced with fully modelled time-delay data, with time-delays in parameter estimation modelled as linear accelerations.	226

8.17 Corner plot of the high-mass, isolated GR binary posterior, generated with model searching simultaneously over a constant acceleration, A , and EdGB gravity, ζ_{EdGB} 228

A.1 First stage in approximating a section of an orbit with linear acceleration 240

A.2 Examples of linear approximations to portions of orbits 241

A.3 The partial, approximate boundaries between regions of parameter space of Kepler orbits requiring different minimal models to describe the time-delays 242

A.4 First stage in approximating a section of an orbit with constant acceleration 243

B.1 Shapiro’s original problem. 245

B.2 Time-delay due to situation of emitter on ‘near-side’ of SMBH. 245

B.3 Path difference between near-side and far-side position. 246

List of Tables

1.1	Details of a subset of the binary black-hole parameters.	41
2.1	References for details of GW detectors that feature in the results of Chapters 7 & 8.	48
3.1	BHB test systems for assessing the downsampling procedure.	99
3.2	Best fit model of posterior evaluation times	106
4.1	Details of the Keplerian elements.	117
4.2	Keplerian elements for example systems.	124
4.3	Coefficients of best fit linear approximations to example time-delays . . .	133
7.1	Injection parameters for low-mass and high-mass binaries	173
7.2	Details of the ‘low-mass’ fiducial GR system signal in the various detector networks.	176
7.3	Standard deviations of the marginals of posteriors from zero-noise realisation data, for the different detector networks, using the isolated low-mass system.	177
7.4	Precision-to-accuracy ratios for low-mass, isolated GR binary	177
7.5	Details of the ‘high-mass’ fiducial system signal in the various detector networks.	192
8.1	Updates to Keplerian and distance parameters for two distinct variations (<i>nearby</i> and <i>distant</i>) of the ‘Orbit 1’ example orbit of Table 4.2.	215
8.2	Maximum log likelihoods sampled by <i>Nessai</i> for all time-delay models. .	227

Acronyms

ACF autocorrelation function.

ASD amplitude spectral density.

BH black hole.

BHB black hole binary.

CBC compact binary coalescence.

CE Cosmic Explorer.

CMB cosmic microwave background.

CMJS combined marginalised Jensen-Shannon.

CW continuous wave.

dCS dynamical Chern-Simons.

DECIGO DECi-hertz Interferometer Gravitational wave Observatory.

DFT discrete Fourier transform.

DOF degree of freedom.

EdGB Einstein-dilaton Gauss-Bonnett.

EFEs Einstein field equations.

ET Einstein Telescope.

FFT fast Fourier transform.

FIM Fisher information matrix.

GB Gauss-Bonnett.

GC globular cluster.

GR general relativity.

GRB gamma ray burst.

GW gravitational wave.

GWB gravitational wave background.

IAD integrated (ascribed) acceleration difference.

IMR inspiral-merger-ringdown.

ISCO innermost stable circular orbit.

JS Jensen-Shannon.

KDE kernel density estimate.

KL Kullback-Liebler.

LFN laser frequency noise.

LIGO Laser Interferometer Gravitational-Wave Observatory.

LISA Laser Interferometer Space Antenna.

MCMC Markov chain Monte Carlo.

MCS maximum correlated samples.

MLE maximum likelihood estimate.

MPA marginal precision and accuracy.

NSC nuclear star cluster.

PAR precision to accuracy ratio.

PCM parameter covariance matrix.

PE parameter estimation.

PN post-Newtonian.

PSD power spectral density.

ROQ reduced order quadrature.

RPM reduced parameter model.

SMBH supermassive black hole.

SNR signal to noise ratio.

SPA stationary phase approximation.

TDI time delay interferometry.

TT transverse traceless.

WSS wide sense stationary.

Introduction

As a functional scientific pursuit, gravitational wave (GW) astronomy is in its early stages. With the first confirmed detection (and still one of the most widely used model examples) of a GW event occurring only a few years ago in late 2015 (named *GW150914*), this relatively new endeavour received an enormous burst of energy and excitement, not to mention financial support, and we are witnessing its burgeoning into a hugely popular and influential field of research. The excitement is appropriate; the great potential of using GWs to measure many physical phenomena and to understand nature in regimes that would otherwise have remained mysterious (perhaps forever) suddenly became a reality. Apart from this, the validity of our current best theory of gravity, general relativity (GR), to some degree relied on such detections; the existence of GWs is of course predicted by the theory. Moreover, the GR waveform itself agreed with predictions to within the error margins imposed by the detector [1]. Prior to this first direct detection, evidence for GWs was scarce; perhaps the best hope of finding GWs was carried by observations of the *Hulse-Taylor* binary system [2] discovered in 1974, which was shown to contain a pulsar producing a signal consistent with an orbital decay, precisely matching that expected due to the emission of GWs in GR. The detection of GW150914 proved that the universe admits GWs, and thus that we have access to a means of further examining theories of gravity and of matter fields, and probing extreme physical conditions and regions of the universe, such as its earliest moments, strong gravity conditions, and indeed the very properties of matter fields and of spacetime itself.

So far, we have detected ~ 100 confirmed GW events with the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo detectors [3] based on the ground. Concepts and plans for future observatories and other novel methods for measuring GWs abound, some of which will be detailed in this thesis. One detector of particular interest, for a number of reasons, but not least because its mission is confirmed [4] is the Laser

Interferometer Space Antenna (LISA), a space-borne observatory planned to launch and commence operations in 2037. LISA will open up a new, low-frequency bandwidth to observation, which promises to unveil exquisitely detailed recordings of low-frequency GWs from high-mass systems, and to give unprecedented constraints on theories of gravity, cosmology, conditions at the big-bang and topological defects. There is also scope to measure, with a somewhat less clear but still very informative set of data, relatively low-mass systems, allowing tests of a wide range of hypotheses including measuring the cosmic expansion rate, lensing studies and testing dark matter theories, binary formation scenarios, informing galaxy structure models, as well as precision parameter estimation of ‘ordinary’ binary systems, which will also help to constrain population models. For the present investigation, this low-mass regime is the focus of the LISA dataset.

This thesis presents new results in the field of *multi-band GW astronomy*, making precise some of the expectations of data analysis and parameter estimation in future eras where detector networks consist of both ground, and space-based (high, and low-frequency, respectively) GW detectors. Analysing LISA (compact binary) data is generally prohibitively computationally expensive (as shall be discussed in detail in this thesis) so previous studies (see Chapter 3 and references therein) have used approximation techniques or formalisms to investigate these expectations¹. In order to deal with this problem here (and to examine the quality of approximation techniques already in use) we first introduce a new method of approximation by which to, relatively rapidly, define a posterior probability distribution in very close resemblance to that which we should expect to derive from real LISA data, at least for the broad class of signals we are focused on in this thesis (compact binary coalescences).

After introductory material in Chapters 1 & 2, the definition and testing of the down-sampling procedure is presented in Chapter 3. The remaining original work details some of the more notable findings that have been discovered by employing this tool, especially in the context of multi-band GW astronomy. In particular, we investigate the comparative feasibility of detector networks of various strain sensitivities in their abilities to constrain source parameters of a given signal. In Chapter 4, we begin to expand the param-

¹It is probably worth briefly remarking that many problems presently being studied in the field of GW astronomy, including those tackled in this thesis, are sufficiently highly structured that ‘simple’, closed-form solutions of exact problems are rarely afforded, and we are more or less bound to the necessity of investigating and solving those problems by numerical methods.

parameter space defining multi-band signals by considering a binary, modelled itself as a point source emitter, in orbit around a supermassive black-hole. Supermassive black-holes are generally to be found at the centres of galaxies, which cause the nearby environment to become relatively crowded with stellar objects [5]; binaries are therefore expected to form there and these signals may be present in the LISA dataset.

Additionally, we present results on the analysis of two modified gravity theories, to test the claim that multi-band GW data analysis will provide particularly stringent tests of General Relativity; the theoretical background is introduced in Chapter 5. Before we examine the posteriors for these systems, an overview of the data analysis framework and all of the adopted approximations/assumptions is given in Chapter 6. The results of multi-band GW data analysis of non-accelerated/non-time-delayed systems is provided with discussion in Chapter 7, with some interesting behaviour of posteriors pointed out. Results for the time-delayed systems are given in Chapter 8, including confirmations of the predictions of Chapter 4 on when these orbital parameters can be recovered from the signal, when they cannot be recovered but must be modelled (at least approximately), and when they can be neglected. A summary of the thesis is given in Chapter 9.

Vacuum Wave Solutions in GR

1.1 Theoretical foundations

It was found soon after the development of GR in 1915 that the vacuum Einstein field equations (EFEs) appeared to admit wave solutions for the metric [6, 7], referred to as GWs. There is a complicated and interesting story behind these early attempts at formulating a rigorous approach to understanding and confirming the existence of vacuum waves, a nice account of which can be found in Ref. [8]. The diffeomorphism invariance of the theory, not being so well understood at the time, caused a high degree of suspicion regarding the authenticity of wavelike propagation of the gravitational field; for example, one can often find a coordinate chart that comoves with the propagating waves in which the spacetime appears to be flat, or, conversely, produce a ‘wavy’ coordinate system on a flat spacetime, apparently demonstrating gravitational waves where in fact there are none.

Determining the correct perspective for describing the physical reality of the wave solution (with which to confirm or deny the existence of GWs) would require fresh insights. In 1952, Choquet-Bruhat’s seminal work on the Cauchy problem in GR showed that if one has some initial data set $(\mathcal{H}, \bar{g}, \mathcal{K})$ and an embedding $f : \mathcal{H} \hookrightarrow \mathcal{M}$, for (semi-Riemannian) manifold \mathcal{M} with metric g , spacelike 3-manifold \mathcal{H} with induced metric \bar{g} , and where \mathcal{K} is a smooth tensor field (the *second fundamental form* of f) on \mathcal{H} that satisfies the ‘vacuum constraint equations’ [9] (whereby the initial data set is said to be a *vacuum initial data set*), then (\mathcal{M}, g) is its unique *development*, and is such that g has

vanishing Ricci curvature. That is, the (globally hyperbolic, and thus well-behaved) development (\mathcal{M}, g) is a unique vacuum solution of the EFEs of which $f(\mathcal{H})$ is a Cauchy surface [10]. Thus Choquet-Bruhat’s work was the first to demonstrate that local solutions of vacuum spacetimes both exist and are unique; it follows immediately that GWs are real physical phenomena in GR, in a setting in which the gravitational field propagates at a finite speed [11]. This result was applicable to the local case of existence and uniqueness, and was extended to the global case (where there exists a unique *maximal* development) by Geroch and Choquet-Bruhat in 1969 [12]; Dafermos goes as far as referring to this theorem of uniqueness of the global development as the *fundamental theorem of General Relativity* [13].

Pirani formulated wave solutions in terms of the physically measurable Riemann curvature tensor in 1956 [14], finally removing remaining doubts about the existence of potentially *observable* GWs that propagate at the speed of light in GR, that had persisted after Einstein constructed rigorous cylindrical wave solutions in 1937 with his long-time collaborator Rosen [15]. Pirani’s work was important for studying GWs because of this, and because of the practical coordinate system that he used and thought experiment posed, which showed test particles could physically be moved in space by the effect of passing GWs [16]. We now understand this motion very well for planar waves, which informs GW detector design as we shall see. The metric is a rank-2 symmetric tensor field, which, for waves propagating at the speed of light, gives rise to two degrees of freedom manifest in the two GW polarisations that shall also shortly be derived. These properties of the metric give rise to the fact that, in a hypothetical quantum field theory for gravity, the quanta known as *gravitons* are excitations of a (ordinarily massless) *spin-2* field, deriving from the π -rotational symmetry of the polarisation states. The two polarisations are conventionally called the “*plus*” and “*cross*” polarisations, since one can choose a particularly useful gauge in which the GW is decomposed into two independent parts, one of which alternately dilates and contracts space in a “horizontal” and “vertical” direction in a plane with normal vector in the direction of the wave (evoking a ‘plus’ pattern, +) and the other which does so along the two ‘diagonals’ in that plane (evoking a ‘cross’ pattern, ×). The understanding of these components is often aided by portraying their effects on a ring of test particles, as can be seen in Figures 1.1 and 1.2. We now step through the modern derivation of the GR wave equation.

1.2 Wave equation in GR

We now derive the wave equation in *linearised* gravity. The EFEs are (neglecting a possible cosmological constant):

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \kappa T_{\mu\nu}, \quad (1.1)$$

where $R_{\mu\nu}$ is the *Ricci tensor*, $g_{\mu\nu}$ is the *metric tensor*, $R \equiv R_{\mu\nu}g^{\mu\nu}$ is the *Ricci scalar*, $T_{\mu\nu}$ is the *stress-energy tensor*, and κ is *Einstein's constant*.

The Ricci tensor is the contraction $R_{\mu\nu} \equiv R^{\sigma}{}_{\mu\sigma\nu}$ of the *Riemann tensor*, $R^{\rho}{}_{\alpha\sigma\beta}$, which is defined in terms of the *Christoffel symbols* (these are the *Levi-Civita connection coefficients*)

$$\Gamma^{\rho}{}_{\alpha\beta} = \frac{1}{2}g^{\delta\rho} (\partial_{\alpha}g_{\delta\beta} + \partial_{\beta}g_{\delta\alpha} - \partial_{\delta}g_{\alpha\beta}), \quad (1.2)$$

as

$$R^{\rho}{}_{\alpha\sigma\beta} \equiv \partial_{\sigma}\Gamma^{\rho}{}_{\alpha\beta} - \partial_{\beta}\Gamma^{\rho}{}_{\alpha\sigma} + \Gamma^{\delta}{}_{\alpha\beta}\Gamma^{\rho}{}_{\delta\sigma} - \Gamma^{\delta}{}_{\alpha\sigma}\Gamma^{\rho}{}_{\delta\beta}. \quad (1.3)$$

This rank-4 tensor field mediates parallel transport of data around the spacetime and encodes local curvature by measuring the degree to which second covariant derivatives fail to commute (i.e., deviate from flat spacetime).

The EFEs are a set of non-linear, second-order, hyperbolic partial differential equations in the metric (sourced, in general) and are notoriously extremely difficult to solve. A pragmatic approach of finding vacuum GW solutions is to study their linearised form by considering small perturbations from flat spacetime, which significantly reduces the difficulty (we shall discuss how this simplification is often reasonable in realistic situations in the following subsections). A nice discussion of a more technically rigorous (less prone to misinterpretations regarding the gauge invariance) viewpoint¹ can be found in Ref. [17]. We shall take a more pedestrian approach here since we only require a derivation of the linearised wave equations. We begin then by taking the deviation from flat spacetime to be very small, where, given a basis, one may write the metric components

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.4)$$

where $\eta_{\mu\nu} \equiv \text{diag}(-1, 1, 1, 1)$ is the metric on Minkowski spacetime, and the $|h_{\mu\nu}| \ll 1$. Thus any second-degree (or higher) terms in the metric and its derivatives only contain

¹One can formulate solutions by considering a diffeomorphism and corresponding pull-back on tensors between a physical spacetime and a background spacetime.

relatively significant contributions at first-order in $h_{\mu\nu}$, i.e. the field equations shall be effectively linearised, as we will see.

The connection coefficients are now given by

$$\Gamma^{\rho}_{\alpha\beta} = \frac{1}{2}\eta^{\delta\rho}(\partial_{\alpha}h_{\delta\beta} + \partial_{\beta}h_{\delta\alpha} - \partial_{\delta}h_{\alpha\beta}) + \mathcal{O}(h^2), \quad (1.5)$$

and the Riemann tensor can be written

$$R^{\rho}_{\alpha\sigma\beta} \equiv \frac{1}{2}(\partial_{\alpha}\partial_{\sigma}h^{\rho}_{\beta} - \partial^{\rho}\partial_{\sigma}h_{\alpha\beta} + \partial^{\rho}\partial_{\beta}h_{\alpha\sigma} - \partial_{\alpha}\partial_{\beta}h^{\rho}_{\sigma}) + \mathcal{O}(h^2). \quad (1.6)$$

The Ricci tensor is then

$$R_{\mu\nu} = \frac{1}{2}(\partial^{\alpha}\partial_{\mu}h_{\alpha\nu} + \partial^{\alpha}\partial_{\nu}h_{\alpha\mu} - \square h_{\mu\nu} - \partial_{\mu}\partial_{\nu}h) + \mathcal{O}(h^2), \quad (1.7)$$

where $h \equiv \eta^{\mu\nu}h_{\mu\nu}$ and $\square = \partial_{\mu}\partial^{\mu}$ is the (flat space) d'Alembertian. The Ricci scalar becomes

$$R = \partial^{\mu}\partial^{\nu}h_{\mu\nu} - \square h + \mathcal{O}(h^2), \quad (1.8)$$

Bringing this all together into (1.1) and dropping $\mathcal{O}(h^2)$ terms, we arrive at the 'linearised EFES', given by:

$$\partial^{\alpha}\partial_{\mu}h_{\alpha\nu} + \partial^{\alpha}\partial_{\nu}h_{\alpha\mu} - \square h_{\mu\nu} - \partial_{\mu}\partial_{\nu}h - \eta_{\mu\nu}\partial^{\alpha}\partial^{\beta}h_{\alpha\beta} + \eta_{\mu\nu}\square h = 2\kappa T_{\mu\nu}. \quad (1.9)$$

Let the *trace-reversed* metric perturbation be defined as

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

so-called since $h = -\bar{h}$, in terms of which equation (1.9) becomes

$$\square\bar{h}_{\mu\nu} + \eta_{\mu\nu}\partial^{\alpha}\partial^{\beta}\bar{h}_{\alpha\beta} - \partial^{\alpha}\partial_{\nu}\bar{h}_{\nu\alpha} - \partial^{\alpha}\partial_{\mu}\bar{h}_{\nu\alpha} = -2\kappa T_{\mu\nu}. \quad (1.11)$$

Now there is some residual gauge freedom (the choice of a coordinate system) that can be put to very good use. First, note the tensor transformation law, which, for a tensor S may be expressed as

$$S_{\mu\nu}(x') = \frac{\partial x^{\alpha}}{\partial x'^{\mu}} \frac{\partial x^{\beta}}{\partial x'^{\nu}} S_{\alpha\beta}(x), \quad (1.12)$$

and consider the coordinate transformation

$$x^{\mu} \rightarrow x'^{\mu} = x^{\mu} + \xi^{\mu}(x), \quad (1.13)$$

for vector field ξ^μ with $\mathcal{O}(\xi) \sim \mathcal{O}(h)$, under which, to $\mathcal{O}(\xi)$,

$$\begin{aligned} g_{\mu\nu}(x) &\rightarrow g'_{\mu\nu}(x) \equiv g_{\mu\nu}(x') = g_{\mu\nu}(x) - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu) \\ \Rightarrow h_{\mu\nu}(x) &\rightarrow h_{\mu\nu}(x) - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu), \end{aligned} \quad (1.14)$$

and thus the trace-reversed metric perturbation transforms as

$$\bar{h}_{\mu\nu}(x) \rightarrow \bar{h}'_{\mu\nu}(x) \equiv \bar{h}_{\mu\nu}(x') = \bar{h}_{\mu\nu}(x) - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu - \eta_{\mu\nu} \partial_\rho \xi^\rho). \quad (1.15)$$

We are now in a position to begin fixing the gauge. Note that

$$\partial^\nu \bar{h}_{\mu\nu}(x) \rightarrow (\partial^\nu \bar{h}'_{\mu\nu})'(x) \equiv \partial^\nu \bar{h}_{\mu\nu}(x') = \partial^\nu \bar{h}_{\mu\nu}(x) - \square \xi_\mu(x). \quad (1.16)$$

The freedom in the vector field ξ allows us to apply the *de Donder gauge condition* in the new, primed coordinates: a choice of ξ such that

$$(\partial^\nu \bar{h}'_{\mu\nu})' = 0, \quad (1.17)$$

by fixing the $\square \xi_\mu = \partial^\nu \bar{h}_{\mu\nu}$. Thus expressing the linearised field equations (1.11) in the new coordinates, gauge fixed using the de Donder condition leads to the simple set of equations

$$\square \bar{h}'_{\mu\nu} = -2\kappa T'_{\mu\nu}, \quad (1.18)$$

which has reduced the original ten degree of freedom (DOF)s in the EFEs to six. Thus in vacuum, these equations reduce to

$$\square \bar{h}'_{\mu\nu} = 0, \quad (1.19)$$

a set of wave equations for the gravitational field, describing waves propagating at the speed of light.

For plane GWs, there is remaining gauge freedom that may be exploited to further simplify the expression of the metric however, for a given GW propagation direction. To do this, suppose now that the vector field

$$\xi_\mu = \xi_\mu^1 + \xi_\mu^2, \quad (1.20)$$

where the vector fields $\xi^{1,2}$ are individually $\mathcal{O}(h)$ and where $\square \xi_\mu = \square \xi_\mu^1 = \partial^\nu \bar{h}_{\mu\nu}$, with $\square \xi_\mu^2 = 0$; the de Donder gauge condition still holds. The vector field ξ^1 thus satisfies

a wave equation sourced by the divergence of $\bar{h}_{\mu\nu}$, and ξ^2 satisfies an unsourced wave equation. Now define

$$A_{\mu\nu}^{1,2} \equiv \partial_\mu \xi_\nu^{1,2} + \partial_\nu \xi_\mu^{1,2} - \eta_{\mu\nu} \partial_\alpha (\xi^{1,2})^\alpha. \quad (1.21)$$

Then inserting (1.20) and (1.21) into (1.15), we can write

$$\bar{h}'_{\mu\nu}(x) = \bar{h}_{\mu\nu}(x) - A_{\mu\nu}^1(x) - A_{\mu\nu}^2(x). \quad (1.22)$$

The (flat space) d'Alembertian and the partials ∂_ν commute, and since $\square \xi_\mu^2 = 0$, we have that

$$\square A_{\mu\nu}^2 = 0, \quad (1.23)$$

so that (1.19) continue to hold. However, we may now impose four further conditions on $\bar{h}'_{\mu\nu}$ using the four arbitrary components of ξ^2 . In particular, ξ^2 can be chosen to set the metric perturbation in the so-called *transverse-traceless gauge*, which shall be explained in further detail after the following description of the transformation.

The first step in achieving this gauge is choosing ξ_0^2 such that $\bar{h}'_{\mu\nu}$ is traceless. Hence we demand

$$\begin{aligned} 0 = \bar{h}' &= g'^{\mu\nu} \bar{h}'_{\mu\nu} = \bar{h} - A^1 + A_{00}^2 - \sum_{1,2,3}^i A_{ii}^2 + \mathcal{O}(h^2) \\ \Rightarrow A_{00}^2 &= \partial_0 \xi_0^2 - \sum_{1,2,3}^i \partial_i \xi_i^2 = A^1 + \sum_{1,2,3}^i A_{ii}^2 - \bar{h} + \mathcal{O}(h^2). \end{aligned} \quad (1.24)$$

where $A^1 = g'^{\mu\nu} A_{\mu\nu}^1$. But note that

$$\sum_{1,2,3}^i A_{ii}^2 = 3\partial_0 \xi_0^2 - \sum_{1,2,3}^i \partial_i \xi_i^2, \quad (1.25)$$

so that to $\mathcal{O}(h)$ we have

$$\begin{aligned} -2\partial_0 \xi_0^2 &= A^1 - \bar{h} \\ \xi_0^2 &= -\frac{1}{2} \int dx^0 (A^1 - \bar{h}). \end{aligned} \quad (1.26)$$

The next step is to choose the remaining ξ_i^2 such that $\bar{h}'_{0i} = 0$. Using (1.22), we set

$$\begin{aligned} 0 &= \bar{h}_{0i} - A_{0i}^1 - \partial_0 \xi_i^2 - \partial_i \xi_0^2 \\ \Rightarrow \xi_i^2 &= \int dx^0 (\bar{h}_{0i} - A_{0i}^1 - \partial_i \xi_0^2). \end{aligned}$$

After having demanded that $\bar{h}' = 0$, we have $\bar{h}'_{\mu\nu} = h'_{\mu\nu}$, so from (1.17) we obtain

$$\partial^0 h'_{00} + \partial^i h'_{0i} = \partial^0 h'_{00} = 0, \quad (1.27)$$

since we chose the ξ_i^2 such that $\bar{h}'_{0i} = h'_{0i} = 0$. Then $h'_{00} = \text{const.}$, which we can set equal to zero for convenience because we are only interested in the time dependent part of the solution (a non-zero value would correspond to the static Newtonian potential). We have thus arrived at the *transverse traceless (TT)* gauge. Now that we have completed all the coordinate transformations, let us drop the primes, and write the metric perturbation in the TT gauge as $h_{\mu\nu}^{\text{TT}}$. In this gauge, all of $h_{0\mu} = 0$, the trace $h = 0$, and the de Donder gauge condition reduces to $\partial_i h^{ij} = 0$.

A useful set of solutions to equations (1.19) are the well-known plane wave solutions, with which one can write the TT gauge metric perturbation as

$$h_{\mu\nu}^{\text{TT}}(x) = p_{\mu\nu}(\mathbf{k}) e^{ik_\alpha x^\alpha}, \quad (1.28)$$

where the $p_{\mu\nu}$ are constants that we refer to as the *polarisation tensor*, which is dependent on $k_\mu = (-\omega, \mathbf{k})$, the *wave vector*, $\omega = |\mathbf{k}| = \sqrt{\delta_{ij} k^i k^j}$ is the frequency of the wave, and noting that we will take the real part in the end. Inserting those solutions into (1.19) gives

$$0 = \square h_{\mu\nu}^{\text{TT}} = \square p_{\mu\nu}(\mathbf{k}) e^{ik_\alpha x^\alpha} = \eta^{\rho\sigma} \partial_\rho \partial_\sigma p_{\mu\nu} e^{ik_\alpha x^\alpha} = -k_\sigma k^\sigma h_{\mu\nu}^{\text{TT}}, \quad (1.29)$$

so that for any non-trivial solution where $h_{\mu\nu}^{\text{TT}}$ are not all everywhere zero, we must have $k_\sigma k^\sigma = 0$, that is, the wave vector k is null and propagates at the speed of light.

Now if we choose, per convention, a wave vector $k^\mu = (\omega, 0, 0, 1)$, i.e., a solution of plane waves propagating along the ‘z-axis’, then $e^{ik_\alpha x^\alpha} = e^{i\omega(t-z)}$ and the condition $0 = \partial^j h_{jl}^{\text{TT}} = -ik^j h_{jl}^{\text{TT}}$ sets all of $h_{3\mu}^{\text{TT}} = 0$. Then tracelessness and symmetry of $h_{\mu\nu}^{\text{TT}}$ allows only two independent components of the polarisation tensor/metric perturbation that we call h_+ and h_\times , which, after taking the real part, reads

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos(\omega(t-z)). \quad (1.30)$$

The TT gauge metric perturbation solution for this particular wave vector therefore has two DOFs only. In the linear regime, the gravitational field is decoupled from itself

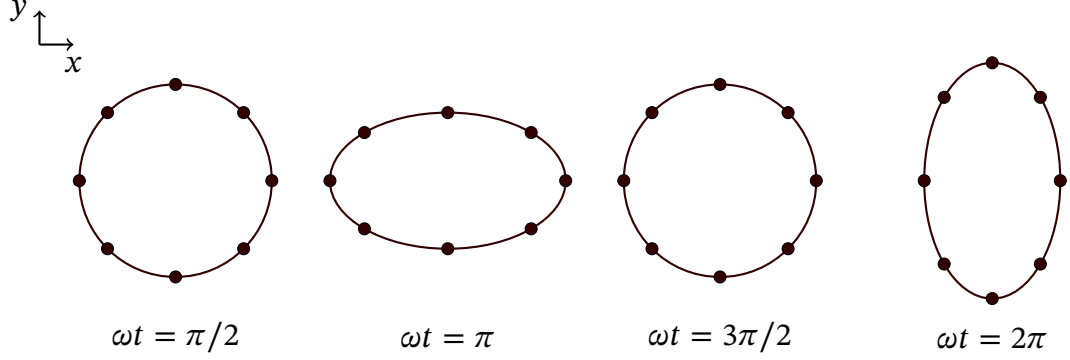


Figure 1.1: “Plus”-polarisation effect of a GW propagating into the page on a ring of test particles for $z = 0$ (for positive constant h_+).

and one is allowed to superpose other solutions, each with their own two DOFs and polarisation tensors. For multiple solutions with the same wave vector, one can therefore generate more interesting waveforms, for example, as one does with a Fourier series. The TT gauge thus gives a particularly simple representation of a linearised plane wave solution, in which the GW is a perturbation of the metric acting on space transversely to its direction of propagation, and in which the perturbation is traceless.

We can now see wherefrom the ‘+’ and ‘×’ labelling convention arises by applying the above metric perturbation to a ring of test particles resting in the xy -plane at some fixed z as one varies t ; we see that h_{xx}^{TT} sinusoidally expands and contracts space in the x -direction over time, where h_{yy}^{TT} simultaneously acts in the opposite manner on the y -direction. This introduces a characteristic ‘plus’ type oscillation, as viewed in the xy -plane, to the ring of test particles as portrayed in Figure 1.1. For the off-diagonal components on the other hand, the metric perturbation describes a simultaneous ‘shear’ of the x -direction of space in the y -direction by h_{xy}^{TT} along with a shear of the y -direction of space in the x -direction by h_{yx}^{TT} . The net result of the two oscillating shears is contraction and expansion, like the ‘+’ mode, but rotated in the xy -plane by an angle of $\frac{1}{4}\pi$ as in Figure 1.2

1.3 Sources of GWs

In GR, the metrics are solutions of the EFEs given in (1.1). Similarly, in linearised GR, the metric perturbations are solutions of the linearised EFEs in (1.18). The EFEs are a relationship between two objects defined on a spacetime: the metric tensor and the




Figure 1.2: “Cross”-polarisation effect on ring of test particles at $z = 0$.

stress-energy tensor. Ordinarily one takes some stress-energy tensor (defined perhaps everywhere on the manifold or just on a Cauchy surface) then goes about trying to solve for some (set of) solutions for the metric satisfying the EFEs. Conversely, one could just as well choose a metric and see which, if any, stress-energy tensors satisfy the EFEs (along with covariant energy conservation). In GW science we are of course interested in those solutions that exhibit GWs. One of the things we would like to know is the part of the solution that is well approximated by the linearised GR wave equations we derived above, and in particular which is generated by some ‘localised’ source. The *source* of a GW is generally meant to refer to a relatively isolated system of mass or curvature, in a particular region of the universe in which the dynamics are not significantly affected by conditions outside that region. In astronomy and astrophysics, one is often interested in isolated systems that evolve on relatively short timescales, independent of ‘cosmic time’. We shall now briefly describe a means of generating GWs by mass acceleration in Section 1.3.1, then point out a commonly used classification of GWs in Section 1.3.2.

1.3.1 GWs from accelerated mass

The EFEs were originally sought and found by Einstein with the intention of having the spacetime curvature sourced by matter, but it has turned out that the vacuum part of spacetime has a complicated life of its own. There are very interesting properties that can be found in the pure vacuum region, notably the recent results of Christodoulou, Klainerman, Rodnianski and others [18, 19] showing that, given certain initial vacuum data, a *closed, trapped surface* can arise; this results in a gravitational curvature singularity as a result of the singularity theorem of Penrose [20], i.e., formation of a black-hole. In other words, those results say that starting purely from ‘GWs in vacuum’, the GWs

can propagate such that they focus in a region (under certain constraints as detailed in the references), whereby a black-hole can form directly out of the GWs. Non-charged black-hole spacetimes are *Kerr* spacetimes, which take two parameters, which we call the *mass* and *spin* of the black-hole. From an outside observer's viewpoint, a black-hole formed from GWs can, for all intents and purposes, appear completely indistinguishable from one formed from the collapse of matter (although whether this is generally true in fact depends on the so-called *no-hair theorem*; the idea that the only information an astrophysically formed black-hole spacetime contains are the two aforementioned parameters), which the GW community are presently attempting tests of [21] and hope to be able to shed more light on in the decades to come. Nevertheless, such black-holes constitute isolated systems of curvature, and may comprise one or both members of a pair of compact bodies that itself could be a source of GWs (see Section 1.3.2.1).

With this caveat pointed out, we will now focus on the more grounded and familiar usage of the EFEs, which is considering an ordinary matter distribution as the source of GWs. It is less enlightening to derive the dependence of the metric perturbation on the mass distribution as it is to see the derivation of the metric perturbation itself, so we refer the interested reader to [22] for those details and provide a brief summary here. The sourced, linear EFEs are given in (1.18), with which one can deduce that the metric perturbation is given by:

$$\bar{h}_{\mu\nu}(t, \mathbf{x}) = 4 \int d^3x' \frac{T_{\mu\nu}(t - |\mathbf{x} - \mathbf{x}'|, \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|}, \quad (1.31)$$

where $T_{\mu\nu}$ is taken to be the stress-energy tensor in the Newtonian limit, i.e. where the velocity of matter particles/fluid elements $v/c \ll 1$, such that the linearised field equations hold. For astrophysical sources, we can simplify this further by restricting valid solutions to hold only in regions at large distances r from the source, where $|\mathbf{x} - \mathbf{x}'| \approx |\mathbf{x}| = r$. Further recasting of the stress-energy tensor allows the above integrand to be written in terms of a time derivative term, and some spatial divergence terms that can be seen to vanish by applying the divergence theorem, taking the boundary of the integration region outside of the source. One is then left with

$$\bar{h}^{ij}(t, \mathbf{x}) = \frac{2}{r} \frac{d^2}{dt^2} \int d^3x' T^{tt}(t - r, \mathbf{x}') x'^i x'^j. \quad (1.32)$$

Defining the projection tensor

$$P_{ij}(\mathbf{n}) \equiv \delta_{ij} - n_i n_j, \quad (1.33)$$

where \mathbf{n} is the local direction of propagation, we can project the \bar{h}_{ij} onto the local TT gauge. That is, we can write our wave vector $k^\mu = (\omega, \mathbf{n})$, so that for propagation along the z-axis, $\mathbf{n} = (0, 0, 1)$. We obtain the TT gauge solution firstly by projecting perpendicularly to \mathbf{n} to leave transverse components, thus

$$\bar{h}_{ij}^{\text{T}} \equiv \bar{h}^{kl} P_{ik} P_{jl}, \quad (1.34)$$

then ensuring the trace vanishes, thus

$$h_{ij}^{\text{TT}} = \bar{h}_{ij}^{\text{TT}} \equiv \bar{h}_{ij}^{\text{T}} - \frac{1}{2} \bar{h}^{kl} P_{ij} P_{kl} = \bar{h}^{kl} (P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl}). \quad (1.35)$$

One often defines the second moment of the mass distribution from (1.32) as

$$I_{ij}(t, \mathbf{x}) \equiv \int d^3 x' T^{tt}(t - r, \mathbf{x}') x'^i x'^j, \quad (1.36)$$

and the quadrupole moment tensor as

$$\mathcal{J}_{ij} \equiv I_{ij} - \frac{1}{3} \delta_{ij} I, \quad (1.37)$$

where $I \equiv I_{ij} \delta^{ij}$. The TT gauge metric perturbation can then be written in terms of the transverse traceless part of the quadrupole moment tensor as

$$h_{ij}^{\text{TT}}(t, \mathbf{x}) = \frac{2}{r} \frac{d^2 \mathcal{J}_{kl}(t - r, \mathbf{x})}{dt^2} \left[P_{ik}(\mathbf{n}) P_{jl}(\mathbf{n}) - \frac{1}{2} P_{ij}(\mathbf{n}) P_{kl}(\mathbf{n}) \right]. \quad (1.38)$$

This is the leading order term in the multipolar expansion of the TT gauge metric perturbation.

1.3.2 A GW classification formalism

It is helpful to begin to organise the study of GWs. One important step in doing this is to invent some sort of classification formalism to categorise GWs, depending on general properties of the waveform. Rather than trying to solve the overwhelming problem of finding the solution for a part of the metric containing GWs that result from general initial conditions and subsequent dynamical evolution, restricting ourselves to working on a particular ‘type’ of GW allows us to make some progress in, for instance, understanding details of their production in astrophysical situations that are known to occur. Many sources are simple enough that one can model them parametrically given a certain type of GW. A common way to categorise is to consider GWs as belonging to one of the following four groups, which is somewhat oversimplified but is nonetheless a very useful foundation.

1.3.2.1 Compact binaries

It is certainly fair to say that the category of GWs that has so far received the majority of attention in the scientific literature are those generated by a so-called *compact binary coalescence (CBC)*. Appropriately named, these GWs originate from a *binary system* (a pair of stellar objects) where each member of the pair is *compact* in the ordinary sense of being relatively small and very dense. When each object is a black hole (BH), a CBC is referred to as a black hole binary (BHB). The pair are assumed to be gravitationally bound and orbiting their common centre of mass in a quasi-circular fashion. All binaries emit gravitational radiation due to this orbital acceleration, however the compactness of the members of CBCs allows for their close proximity (i.e. a small orbital radius) before they ‘merge’, and thus a far greater amplitude than one might otherwise see, allowing for easier detection. If the binary is truly isolated, this closeness is guaranteed; the acceleration of one object due to the other generates GWs, the energy of which are taken from the orbital energy (angular momentum) of the binary. The pair then fall closer together, thus emitting higher energy GWs and falling inwards at a greater rate. This runaway process of a CBC is referred to as the *inspiral*, which goes on emitting GWs with gradually ramping up frequency and amplitude (called a ‘chirp’) approximately until the pair reach a proximity at which the pair no longer orbit, but instead plunge into each other, becoming a single object. This very short duration plunge is known as the *merger* and the resulting highly non-spherical object radiates a great deal of energy very rapidly as it settles into a quasi-spherical state: often modelled as a perturbed, *Kerr* (spinning) black-hole. The final stage, known as the *ringdown*, radiates away the perturbations, tending towards a Kerr black-hole, where curvature acceleration vanishes and no more GWs are emitted. The whole process is often referred to as *inspiral-merger-ringdown (IMR)*. One then talks about the *IMR waveform* (see Figure 1.3) when discussing parameterised GWs of CBCs, especially in the context of simulating waveforms. More detail on these waveforms is given in Section 1.4.1. Members of CBCs are ordinarily BHs, but other massive, compact astrophysical objects may constitute one or both members of a binary system emitting detectable GWs, including *neutron stars*, *white dwarfs* and other exotic objects, such as *quark stars* [23] of various sorts.

The GW amplitude is proportional to the total mass of the binary, and inversely proportional to the distance, as one can see from equation (1.38). Thus binaries that are

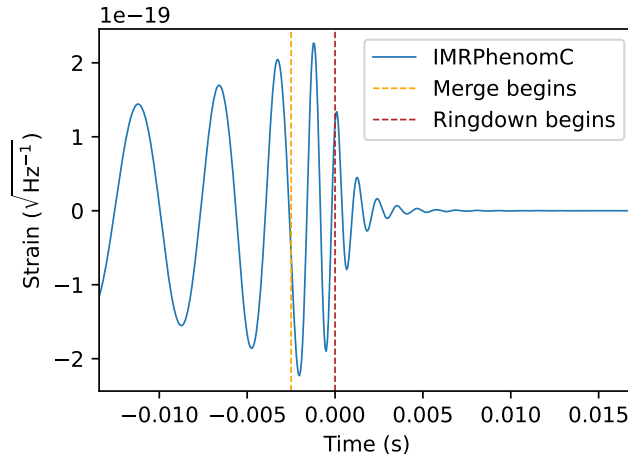


Figure 1.3: An example gravitational wave of a binary black-hole merger, depicting all three stages of its evolution: the inspiral, merger and ringdown.

more massive and closer to a detector are, on the whole, louder than the less massive and more distant ones. The prevalence of these sources has secured them a category of their own; so far, all of our detections of GW events have been sourced from CBCs.

The detectable part of the GW usually contains the portion of the IMR waveform occurring at and around merger where the amplitude and frequency are greatest. The latest catalogue of GW events can be found in Ref. [24]. Often, the detectable portion of the signal in ground based detectors (see Section 2.1) containing any useful information is of short duration, lasting from a few seconds to perhaps a few tens of seconds. As a result, these signals are sometimes placed in a category referred to as ‘transients’ (which itself contains *bursts*, discussed in Section 1.3.2.3), although the inspiral part may in practice last many millions of years. This point is pertinent for this thesis; we hope to see how multi-band detector networks (see Section 6.1.2) can benefit inferences about the source by obtaining and analysing signal from different regions of the GW.

Suppose the detector observes a binary with a circular orbit at merger, the loudest part of the waveform. This occurs approximately when the orbital radius of the binary drops to around the *innermost stable circular orbit*, r_{ISCO} , of one of the objects; the radius within which test particles can no longer stably orbit (circularly) and which then fall inwards towards the object. Beyond this point, the high non-linearity of the EFEs dominates over ‘weak field approximations’ of the orbit. Up to this point however, one can approximate the orbit and the gravitational waveform using Newtonian and *post-Newtonian* (*PN*) terms in expressions of those quantities that are written as terms of increasing or-

ders of v^2/c^2 , where v is the characteristic speed of the binary [25]. One usually marks the innermost stable circular orbit (ISCO) as the end of the inspiral and beginning of the merger [26]. The PN approximation, truncated at some reasonable number of terms (see Section 1.4.1), is generally accurate [27] until this point for systems consisting of compact objects that do not deviate much from Newtonian systems (i.e. slowly moving with little spin). Using a Kepler orbit approximation for the inspiral then, with orbital frequency $\omega_s^2 = GM/R^3$ and where $R_{\text{ISCO}} = 6GM/c^2$, where M is the mass of the object (total mass of the binary) [26], the ISCO GW frequency can be estimated, where one finds that

$$(f_s)_{\text{ISCO}} \approx 2.2\text{kHz} \left(\frac{M_\odot}{M} \right), \quad (1.39)$$

where M_\odot is the solar mass unit. Note that the gravitational wave frequency f_{gw} (the measured signal) is related to the orbital frequency f_s via $f_{\text{gw}} = 2f_s$. The ISCO is computed assuming a test particle of negligible mass orbiting a massive body, but under the Newtonian approximation, the force between the bodies is the same if the mass were shared amongst them, so we can approximate the above to be equivalent to the ISCO for binaries of comparable mass, where we take M to be the *total mass* of the binary.

The current ground-based detectors are most sensitive to the frequency band from around 20 Hz to 5 kHz, so for binaries that merge in that band, this corresponds to a total mass range of between around 0.8 to 200 solar masses. As mentioned, however, the amplitude depends on the mass also, so the lower mass systems are harder to detect (when comparing sources at the same distance). The future space-based detectors will be sensitive to much lower frequencies. In particular, LISA will be most sensitive to the band from around 1 mHz to 100 mHz, corresponding to CBC *mergers* in the mass range from around 4×10^4 to 4×10^6 solar masses.

However, for the inspiral part of a CBC waveform, consider the lowest order part of the post-Newtonian expression of the frequency (which dominates the frequency value far from merger where the weak field approximation becomes accurate) [26]

$$f_{\text{gw}}(\tau) \approx 134 \text{ Hz} \left(\frac{1.21 M_\odot}{\mathcal{M}_c} \right)^{5/8} \left(\frac{1 \text{ s}}{\tau} \right)^{3/8}, \quad (1.40)$$

where $\tau = t_c - t$ is the *time to coalescence*, t_c is the coalescence time, and \mathcal{M}_c is a useful reparameterisation of the binary component masses m_1 and m_2 that characterises the ‘chirping’ behaviour of the waveform in the inspiral stage, defined by

$$\mathcal{M}_c \equiv \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad (1.41)$$

and is known as the *chirp mass*. As one increases the time to coalescence, the frequency drops; since (1.40) is a reasonable approximation for all $\tau \gtrsim \tau((f_{\text{gw}})_{\text{ISCO}})$ each CBC thus covers a wide frequency range (the lower bound depends on the formation scenarios, see Section 4.2.1). Figure 1.4 shows such an earlier section (of the signal displayed in Figure 1.3) only a few tenths of a second before merger. It is then possible that a low frequency part of the inspiral of a CBC may be detected in space-based, low-frequency detectors, and a high-frequency part of the same signal may be detected in ground-based, high-frequency detectors. CBCs are thus good examples of *multi-band* GW sources, since their signals appear in multiple detector bands, but whether a CBC is classed as ‘multi-band’ does not have a strict classification formalism; it depends on the *detector network* (see Chapter 2).

1.3.2.2 Continuous waves

Differentiating equation (1.40) with respect to τ , one finds that $\dot{f}_{\text{gw}}(\tau) \propto \tau^{-11/8}$. As one increases τ to look at earlier stages of the inspiral, one finds very soon that $\dot{f}_{\text{gw}}(\tau) \rightarrow 0$; $f_{\text{gw}}(\tau)$ tends to more closely resemble a constant (over a given finite duration). Earlier stages of the signal thus tend to become increasingly more manifestly sinusoidal. If one takes an early enough segment of a given length of an inspiral and observes for a short enough time, it is possible that one would be unable to detect any variation from an ordinary sinusoidal wave. This is one example of a so-called *continuous wave (CW)*.

Other possible sources of CWs include *neutron star mountains*. It is thought that




Figure 1.4: An earlier section of the inspiral of the CBC merger shown in Figure 1.3. The waveform is already rather slowly evolving just seconds before merger. The evolution is not obviously noticeable by eye over this short duration; it resembles a ‘continuous’, sinusoidal wave.

neutron stars (which have radii of around 11km) may have ‘mountains’: small bumps, with heights above the surface on the order of millimetres, on the surfaces of the otherwise highly spherically symmetric mass of neutron matter. As a neutron star with a mountain rotates, the acceleration of the subtly asymmetric mass distribution generates GWs at an extremely stable frequency; the energy loss of these systems to GWs is very slow and causes little evolution of the state the source. Neutron star mountains are the primary search target of CWs, since the current understanding is that they are the most likely source of detectable CWs. There is also some interest in studying CWs from other types of dense, exotic stars [28] such as quark stars. Neutron stars are also thought to be able to generate CWs via *oscillation* processes [29].

1.3.2.3 Burst sources

The so-called *burst* signals are short duration, ‘unmodelled’ (of unknown waveform) because either the source itself is unknown, or perhaps that the waveform of a known source is presently too difficult to reproduce and is not yet known. Potential sources of bursts include (but not limited to) gamma ray burst (GRB) sources [30], supernovae core collapse [31], collisions of cosmic strings [32], and hyperbolic encounters of compact objects [33].

There is also work being done on considering interactions between black-holes and cosmic strings [34]. Searching for unmodelled signals requires different sorts of methods than those we shall consider here, since our focus is on IMR waveforms. Further details on burst searches can be found in Ref. [35].




Figure 1.5: A simple example of a possible burst GW. The waveforms are unmodelled and may derive from collapsing supernovae, or cosmic string collisions, for example.

1.3.2.4 Stochastic background

The (*stochastic*) *gravitational wave background (GWB)* is the ‘background noise’ of gravitational waves generated by all sources located exactly on the *past light-cone* [17] of the detector, which are all linearly superposed, resulting in an essentially stochastic signal. The sources that contribute are understood to be those quiet enough such that they cannot be individually recognised from the cacophony of sources, but which together can be expected to constitute a signal with specific statistical properties.

Sources are ordinarily categorised as being either of cosmological or astrophysical origin. Cosmological and astrophysical models therefore allow predictions of the statistical properties of the GWB, measurements of which can then be used to place constraints on the models [36].

Cosmologically, GWs are thought to originate from inflation [37] and other processes, such as electroweak phase transitions [38], occurring at the big-bang; this part of the GWB is analogous to the cosmic microwave background (CMB): the earliest electromagnetic source of information in our universe. The astrophysical part of the GWB is the superposed combination of all other sources (such as those described above); one might think of this part as (weakly) analogous to the foreground effects that partially obscure the CMB, though foreground GWs add to the GWB (since matter is weakly coupled to gravity and does not impede propagation of GWs), whereas foreground matter obscurations generally subtract from the CMB. GWs from primordial black-hole mergers are an interesting potential source that may be said to lie on the boundary between of cosmo-




Figure 1.6: An example stochastic GW background signal. Contributions to this signal can include background compact binary confusion noise, and GWs from early universe phase transitions.

logical and astrophysical origin, since they can be technically thought of as astrophysical objects, but originate from the big-bang and may well provide good constraints on the conditions there [39].

1.4 BHB waveform models

Equation (1.38), describing the gravitational waveform sourced by non-relativistic matter, is derived from assuming that GR is the theory of gravity governing the relationship between spacetime (the metric tensor) and energy (stress-energy tensor). However, given a different theory embodied by a different set of equations governing that relationship, one is likely (/certain) to arrive at a different gravitational waveform in the linear (/fully non-linear) regime, given the same initial data. We shall indeed consider two other (one-parameter families of) gravity theories in Chapters 5–8, examining how the waveform is affected, and thus how the waveforms may be used as probes of these theories. The theories are formulated such that as the parameters mediating the strength of the modifications go to zero, the theories reduce to GR, and thus one can measure the closeness of fit of gravitational waveforms to those theory parameters to physically test them.

The ordinary EFEs on their own are extremely difficult to solve however, despite being the most simple, and indeed, the only 4-dimensional metric tensor theory of gravity that can be derived from (the variation of) an action with at most second-order derivatives of the metric: a consequence of *Lovelock’s theorem*, [40]. The modified gravity theories considered later have therefore been used to derive waveforms *in the small coupling limit only* (see Section 5.3); that is, the theories considered are assumed to be small deviations from GR. In that case, one can start with the framework used to derive waveforms in GR (the so-called *PN* framework [41]) and include small corrections from the additional couplings to gravity defined by the theory [42] throughout the derivation of the PN waveform. Ultimately, these corrections appear in the form of extra terms that arise at specific PN orders in the waveform phase and amplitude, in addition to the ordinary GR PN expansion terms. Hence we shall in this section briefly introduce some archetypal (unmodified) GR waveform ‘approximants’ that have been developed and are in use by the GW community (see [43] and references therein for details). Details of waveform modifications we have modelled for this thesis are given in Section 4.2.

1.4.1 Approximants

Gravitational waves from BHBs, despite the simplicity of solutions of single black-hole spacetimes, are extremely difficult to model. This is due to the gravitational self-interactions embodied by the highly non-linear nature of the EFEs. No neat, closed-form expression of the global solution of the metric is known, so we must solve the equations either numerically, or by making approximations. Numerical solutions are so expensive (and difficult to initialise) that it is not usually feasible to do *parameter estimation* (a common data analysis technique, see Section 2.3) with numerically generated waveforms, although progress is being made in that area [44]. We are thus generally limited to using waveform approximations, commonly referred to as *approximants*.

There are various ways one can go about formalising the framework within which to generate approximants. The details are beyond the scope of this thesis, but we note two prominent approximants. The *spinning effective-one-body calibrated to numerical relativity* (SEOBNR) waveform is generated by transforming the relevant equations into equations describing a single object [45, 46], and the *phenomenological inspiral, merger and ringdown* (IMRphenom), which splices together portions of the waveform produced using different methods; the early inspiral is written as a PN expansion, and the late inspiral, merger and ringdown can be derived from numerical relativity solutions [47]. Black-hole perturbation theory is also often employed to model the ringdown, which, since one can inspect sets of equations describing the solution (which is lacking in numerical solutions) also offers interesting insights into the behaviour [48].

Approximants are of course not perfect waveform models. As such, one will always have some systematic error. For the detection of GW150914, this error has been estimated, and it is believed that the systematic error is less than the error that is present due to the noise in the signal [49, 50]. One can expect that the PN expansions become more accurate as one looks to earlier points of the signal; since the PN expands in terms of orders of v^2/c^2 , where v is a characteristic speed of the binary, at earlier times one will have lower speeds. This is confirmed by inspecting the differences between the *TaylorT3* approximant and the numerical relativity waveforms in Ref. [27]. Indeed, the waveform approximant of Ref. [27] (without the ‘twisting-up’ procedure required to model spin precession) shall be implemented and employed to provide the results for this thesis.

1.4.2 BHB model parameters

Kerr black-hole spacetimes are completely characterised by a mass and a spin parameter. Introducing another Kerr black-hole means that to specify the spacetime, one now has to consider the relative positions, velocities and spin directions of both black-holes. This can be done at some reference coordinate time, t_0 , when coordinates have been specified. If one assumes the black-holes are gravitationally bound, their relative positions and velocities can be reparameterised as Keplerian parameter *eccentricity* at t_0 , and, since the binary's evolution is completely determined by its properties at t_0 , the *time to coalescence*. Thus there remain two mass, six spin and one eccentricity parameters that are *intrinsic* to the binary [51]. The time to coalescence is coordinate/observer dependent (if approximating the background spacetime to be Minkowski, i.e., no imprint of expansion of space on the GW) and thus can be classed as *extrinsic*.

We do not observe the GW solution over the entire spacetime of course, and observers are located essentially at a point in space, far from the source. The waveform appears differently depending on this location and relative orientation, which we require further extrinsic parameters to describe. These are: distance, phase, inclination, orientation (polarisation), right ascension, declination, pericentre angle and coalescence time [51].

The commonly used effective spin and spin precession parameters, χ_{eff} and χ_{p} , respectively, are given by [51]

$$\chi_{\text{eff}} \equiv \frac{\chi_1^{\parallel} + q\chi_2^{\parallel}}{1 + q}$$

$$\chi_{\text{p}} \equiv \max\{\chi_1^{\perp}, \chi_2^{\perp}q(3q + 4)/(4q + 3)\},$$

where χ_i^{\parallel} and χ_i^{\perp} are, respectively, the parallel and perpendicular components of the i^{th} black-hole's spin relative to the orbital angular momentum. These two parameters are useful approximate reparameterisations of the *six* spin components of the waveform, able to capture the majority (but clearly not all, since there are fewer degrees of freedom, see [52]) of the spin-precession effects of coalescing BHBs. More details on the IMR parameters detectable by LIGO can be found in Ref. [51]. See Table 1.1 for further details on the reduced set of parameters we will use in our model. See Section 6.3.1 for discussion on the effects on analysis of the omission of parameters from the model.

Note that all the parameters besides the sky position parameters (right ascension and declination) directly affect the waveform that arrives at the detector, either through the

intrinsic parameters that generate the GW, or the relative location and orientation of the source on the sky. The sky parameters affect only the part of the waveform that is *measured by the detector*, as they describe detector orientation with respect to the incident GW, on which the detector response depends. Thus, detectability of sky position itself is detector dependent; one can better measure the location of short signals using multiple detectors with different orientations (such as LIGO), or of long signals with detectors that change spatial orientation or position (or both) as functions of time (such as LISA).

Name	Symbol	Units	Range
Chirp Mass	\mathcal{M}_c	M_\odot	$[\mathcal{M}_c^{\min}, \mathcal{M}_c^{\max}]$
Mass Ratio	q	-	$(0, 1]$
Effective Spin	χ_{eff}	-	$[0, 1)$
Luminosity Distance	d_L	Mpc	$[d_L^{\min}, d_L^{\max}]$
Coalescence Phase	ϕ_c	rad	$[0, 2\pi)$ (cyclic)
Inclination	ι	rad	$[0, \pi)$ (not cyclic)
Orientation	ψ	rad	$[0, \pi)$ (not cyclic)
Right Ascension	α_\oplus	rad	$[0, 2\pi)$ (cyclic)
Declination	δ_\oplus	rad	$[-\pi/2, \pi/2)$ (not cyclic)
Coalescence Time	t_c	s	$[t_c^{\min}, t_c^{\max}]$

Table 1.1: Details of a subset of the BHB parameters. The constants \mathcal{M}_c^{\min} , \mathcal{M}_c^{\max} , d_L^{\min} , d_L^{\max} , t_c^{\min} & t_c^{\max} depend on detector characteristics and operating time, and limits imposed by cosmological models/astrophysical processes. This subset will be used to define our BHB signal model. It is chosen to capture the main dependencies of the waveform and as a partial simplification of the problem; we will later introduce waveform modifications with extra parameters, increasing the difficulty of the problem. Omitted parameters include spin precession and orbital eccentricity, see Section 6.3.1 for more discussion. A full list of parameters can be found in Ref. [51].

GW Data & Analysis

GWs can be observed indirectly, by, for example, observing systems that evolve according to predictions consistent with generation or interaction with GWs, such as Hulse-Taylor-like pulsars [2] and *pulsar timing arrays* (GWs passing between a pulsar and Earth will change the arrival time of pulses; the change in pulse arrival time from an array of pulsars could be used to reconstruct GW source parameters, [53]). In big-bang cosmology, gravitational waves in the early universe would be indirectly observable as a manifestation of ‘B-mode’ polarisation in the CMB.

We now live in the era of direct GW detection however, with instruments capable of measuring the distortions in spacetime themselves, caused by passing GWs. Previous attempts at building *resonant bar detectors*, large metal cylinders that resonate as GWs transit, had unconfirmed effectiveness; despite initial claims of detections by Weber, his inferences were not consistent with calculations of the maximum rate of energy loss by mass to gravitational radiation conversion allowed for stability of galaxy structure, and later observation attempts with higher quality resonant bar detectors were not able to reproduce Weber’s claims [54]. LIGO, LISA and most other serious current conceptions of GW detectors with very high sensitivity are *interferometers* [55], which we introduce in Section 2.1. It is essential to understand and properly treat inherent *noise* in the detector’s output, detailed in Section 2.2, which could otherwise be confused with signal data and which directly informs *parameter estimation*, discussed in Section 2.3.

2.1 Laser interferometers

Interferometers use the principle of superposition of electromagnetic (EM) fields that results from the linearity of Maxwell’s equations. Superposition allows for the simplest realisation of *interference*, as the electric and magnetic fields at some time and location are simply given by the sum of values from different ‘wave-packets’ at that point. This phenomenon is exploited in interferometers by splitting a coherent light beam into multiple beams with a *beam splitter* and sending those beams in different directions. The beams are guided along a path by a set of mirrors and recombined at another location, where interference patterns can arise. There are various possible configurations for such a device. A simple interferometer design, on which LIGO [56] is based, is shown in Figure 2.1. Other (existing or planned) GW interferometers of this or similar type include Virgo, Kagra, Einstein Telescope (ET), Cosmic Explorer (CE), DECI-hertz Interferometer Gravitational wave Observatory (DECIGO), and others [55]. If the cavity length changes in one arm differently to that in the other arm due to an incident GW (see Figures 1.1 & 1.2, considering the mirrors to be the test particles) one would observe an interference pattern in the output at the photodetector, since the extra distance the light is required to travel results in a phase offset between the beams where they are recombined.

One can then very precisely derive the proper distance between the mirrors and thus the metric perturbation, since it is encoded in the phase difference in the light, $\Delta\phi_L$, which is computed from the interference observed at the photodetector. In the TT-gauge (see Section 1.2), and when the GW frequency is much lower than the photon round trip frequency in the cavity [57], the phase difference is given by

$$\Delta\phi_L(t) = h_{11}^{\text{TT}}(t) \frac{4\pi L}{\lambda_L}, \quad (2.1)$$

where $h_{11}^{\text{TT}} = -h_{22}^{\text{TT}}$ is the metric perturbation given in equation (1.38), L is the length of the cavity, and λ_L is the wavelength of the laser light. We can therefore scale the sensitivity to incoming GWs by increasing L or decreasing λ_L . However, various sources of noise and practical constraints prevent one from arbitrarily increasing L and/or decreasing λ_L ; some of these noise sources are mentioned in Section 2.2.

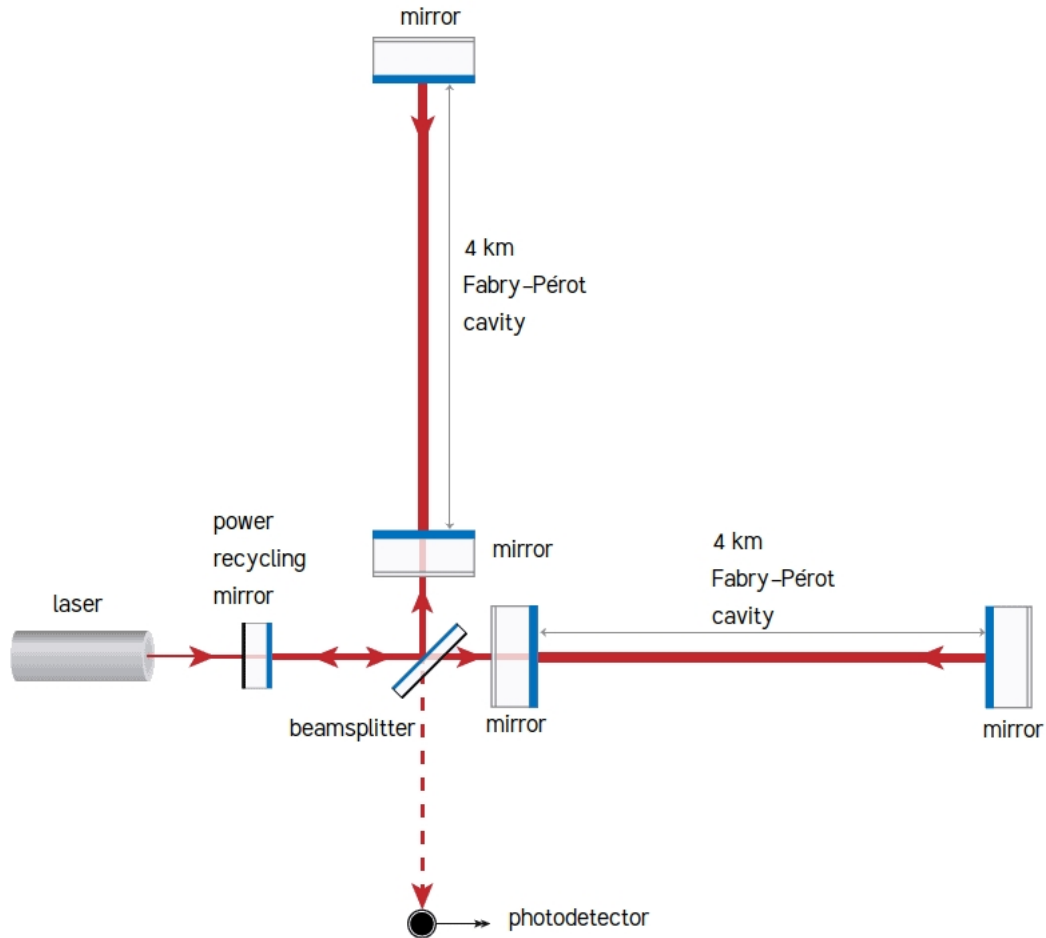


Figure 2.1: An equal arm laser interferometer schematic diagram. This is a vast simplification of the LIGO detectors, with more detail available in Ref. [56]. The power recycling mirror increases the detector sensitivity, as do the Fabry-Pérot cavities, which cause the light to spend more time in the arm, reflecting back and forth multiple times and thus accruing more phase offset when the arm length changes due to a passing GW. Figure adapted from source: IOP publishing.

2.1.1 Time-delay interferometry

In ground-based detectors, *laser frequency noise (LFN)*, which results from the uncertainty in the frequency of the laser light, can be relatively easily cancelled by having two arms, where the second arm contains light that is π -radians out of phase with the light in the first. When the beams are combined, the fields destructively interfere, resulting in a ‘dark fringe’ at the photodetector, independent of laser frequency and thus its uncertainty. However, space-based detectors consist of free falling base stations with different and time-varying distances between the stations, as depicted in Figure 2.2. It is therefore not possible to remove the LFN in such space-based detectors in the same way as is done

in ground-based detectors.

Suppressing LFN in LISA can be achieved by the now well-established method of *time delay interferometry (TDI)* [58]. This works by first obtaining time differences for each arm (see Figure 2.2), each generated by interference of transmitted light with returned light, then writing a linear combination of the phases that are artificially delayed in such a way as to be completely independent of the LFN. This is essential since the LFN is orders of magnitude louder than the expected signals and would otherwise completely drown out the signal [59]. However, to do this, we need to know the delays. In 2005, *TDI ranging (TDIR)* was proposed as a technique to achieve this, which minimises the power in the combined output as a function of time-delays (derived from spacecraft positions resulting from their orbits) for each TDI combination [60]. More recent work shows promise in including the time-delays as parameters in a "global-fit" model [59], in which one estimates the arm lengths as parameters, simultaneously with estimating GW parameters (see Section 2.3).

Other space-based interferometers that require TDI approaches of suppressing LFN include the TianQin and Taiji detectors [55]. We will not be modelling TDI in our simulation of the LISA observatory.




Figure 2.2: Example schematic diagram of time-delay interferometer (TDI). A laser beam is split into two beams and directed into two arms of different lengths, L_1 and L_2 . Each beam is made to interfere with the injected light and recorded in its own photodetector (P.D). Figure source: [58].

2.1.2 Detector response and antenna patterns

Single arm interferometers could (in principle) measure the projection onto the axis parallel to the arm of the expansion/contraction of space caused by passing GWs. Dual arm detectors *could* measure spatial distortions projected onto the planes spanned by the arms (LIGO does not achieve this exactly however, since the arm measurements are combined to remove LFN as mentioned in Section 2.1.1, leading to blind spots, see, for example, Ref. [61]).

The metric perturbation measured by a detector is thus dependent on the relative position and orientation of the detector and source (and the subtleties of the workings of the detector). The part of the output of a detector, $h_{\text{out}}(t)$, generated by metric perturbation $h_{ij}^{\text{TT}}(t)$ can be written [62]

$$h_{\text{out}}(t) = h_+(t)F^+(t) + h_\times(t)F^\times(t), \quad (2.2)$$

where $F^+(t) = F^+(t; \psi, \alpha_\oplus, \delta_\oplus)$ and $F^\times(t) = F^\times(t; \psi, \alpha_\oplus, \delta_\oplus)$ are the *antenna pattern functions*, usually different for each detector; see Table 2.1 for references.

For our modelling of detector response in LIGO, we use the complete model as provided by `lalSuite` [63], which treats the problem properly. For modelling the LISA and DECIGO detectors, we ignore the antenna patterns and use the *sky-averaged* sensitivity [62]. This is acceptable for general simulation studies; these detectors have no (fixed) blind spots because of the motions of the base stations through space and the long durations of the signals. Omission of the antenna patterns from the model can have an effect on the sky localisation capabilities of the analysis of the detector output. This is a bigger issue for DECIGO, with signals lasting on the order of days (see Table 7.2, for example); the analyses would make more use of the changing spacecraft constellation configuration and consequent changing signal arrival times. However, since we will model *Roemer delay* (the large changing signal delay throughout a year as LISA and DECIGO orbit the Sun; see Section 4.2.2.2) and signals will be present in LISA for multiple years where the Roemer delay overwhelmingly dominates the time-delays, the effect of omitting the LISA antenna patterns on the LISA sky localisation will be minimal.

2.1.3 Basic overview of detectors

In this short section, we give a brief introduction to the detectors to be modelled and whose data are to be analysed in the final chapters of this thesis. See Chapter 6 for more general discussion on multi-band detector networks. To select the detectors to study, we begin with the currently existing LIGO observatory and, since the LISA mission is confirmed, we ask: how would the results of the data analysis change if LISA were already existing in the network prior to the LIGO detections? We then ask the same question of DECIGO, despite its mission status being unconfirmed, since it straddles the frequency band between LISA and LIGO and we may acquire a better understanding of, for instance, how parameter covariances evolve as functions of frequency.

2.1.3.1 LIGO

LIGO consists of two laser interferometer GW detectors (see schematic in Figure 2.1), based in the USA, one in Livingston and one in Hanford, separated by a distance of 3,002 km. Their arms are oriented at a relative angular offset so as to provide optimal simultaneous sensitivity to a particular polarisation of GW. Whilst the arms are each 4 km long, the Fabry-Pérot cavities trap the laser light such that photons make around 300 round trips, accruing phase offset with each reflection before exiting the cavity, giving the arms an effective path length of 1,200 km.

2.1.3.2 LISA

LISA is a space-based, future GW detector with plans to commence operations in 2037. It shall consist of a constellation of three spacecraft in tilted heliocentric orbits (with respect to the Earth's) trailing Earth at a mean distance of around 5×10^7 km, such that the constellation maintains an equilateral triangular arrangement, with side lengths of 2.5×10^6 km. Due to the large distances between the spacecraft, laser interferometry is not feasible; only a small fraction of the power of each laser beam is received, hence TDI is required (see Section 2.1.1). Each spacecraft contains a free floating test mass and surrounding housing that absorbs non-gravitational forces (such as solar radiation) using small thrusters to keep the housing centred on the test mass. The required hardware, such as lasers and communications devices, are attached to the housing.

2.1.3.3 DECIGO

DECIGO is a proposed space-based detector, designed to be most sensitive in the deci-hertz regime, plugging the gap between LISA and LIGO-like detectors; see Figure 2.3. DECIGO will consist of four clusters of spacecraft, each of which, like LISA shall be a group of three spacecraft in an equilateral triangular arrangement, following a heliocentric orbit. However, the interferometer arms are proposed to be 1000 km long Fabry-Pérot cavities. Two clusters will overlap in a regular hexagram configuration. The hexagram and the remaining two clusters will then be positioned in an equilateral triangle formation with the Sun at the centre, to form the entire DECIGO constellation.

Detectors	Shorthand	PSD & antenna patterns
LISA 4 yr / LISA 10 yr	L_s^4 / L_s^{10}	In Ref. [62]
LIGO	L_g	In Ref. [63]
DECIGO	D	In Ref. [64]

Table 2.1: References for details of GW detectors that feature in the results of Chapters 7 & 8.

2.2 Noise

Detectors are not perfect machines and do not operate in ideal environments. Various physical effects deriving from electronic components and/or mechanical operations and properties of the instrument, and perturbations directly from the environment introduce an unwanted contribution of *noise* to the detector output, which must be appropriately treated in order to acquire correct inferences from the data. From the instrumentation side, the job is to minimise noise in the output of a detector, and from the data characterisation side, the job is to precisely *statistically* model the noise to facilitate the proper treatment of the data.

We briefly introduce some noise sources in Section 2.2.1, and look at the statistical properties of the noise in Section 2.2.2. We will consider the noise to be (*weak-sense*) *stationary* (meaning the noise sample correlations and noise mean are constant in time) and *Gaussian*, however, note that these are not necessarily realistic assumptions for real detector data; non-stationarity and non-Gaussianity are expected properties of LISA data.

2.2.1 Sources

There are a multitude of noise sources in laser interferometers; some unique to a given instrument, and some common in all laser interferometry. We have already mentioned one common noise source: the laser frequency noise. As discussed however, this is able to be effectively cancelled in both ground and space based observatories and thus does not significantly contribute to the detector output, and for correct accounting of noise in data analysis, only precise modelling of noise present in the detector output is required. We give brief introductions to some LISA and LIGO noise sources here. More detail on LISA noise sources can be found in Ref. [65], and on LIGO noise sources in Ref. [66].

2.2.1.1 LISA noise sources

- Galactic binary confusion noise: millions [67] of unresolvable galactic (i.e., in the Milky Way) binaries are thought to be present in the LISA data, made up of white dwarfs, neutron star, and stellar mass black-holes. Since the signals are not expected to be able to be teased out of the data, they constitute a source of noise.
- Acceleration noise: the LISA spacecraft shield and housing attempt to follow free-falling test masses. However, various disturbances are experienced by the test masses, such as cosmic ray deposition leading to electrostatic forces and Brownian force by stray particles in the test-mass housing.

2.2.1.2 LIGO noise sources

- Seismic noise: this includes the Earth's seismic activity, earthquakes, ocean waves, and noise from human activity (by transportation or operation of heavy machinery for example). This causes ground movement at the LIGO sites and is the dominant noise source at low frequencies.
- Thermal noise: the mirrors/test masses at LIGO are suspended on thin silica threads that vibrate with thermal noise, causing the arm lengths to change. Thermal noise is also introduced into the mirror coatings by heating from the laser.
- Quantum noise: the two main sources of quantum noise are the so-called *shot noise* and the radiation pressure noise. The shot noise is (Heisenberg's) uncertainty of

the arrival time of photons at the photodetector, and the radiation pressure noise derives from the recoil of the mirrors from photon momentum transfer.

2.2.2 Power spectral density

The (one-sided) *power spectral density* (PSD), $S_n(f)$, of the detector noise, $\tilde{n}(f)$, describes the power of frequency components of the noise and can be defined as [68]

$$\mathbb{E}[\tilde{n}(f')^* \tilde{n}(f)] = \frac{1}{2} \delta(f - f') S_n(f), \quad (2.3)$$

where \mathbb{E} is the expectation value operator and δ is the Dirac delta function. If two noise sources are independent and generated by distinct noise generating processes, P_1 & P_2 , generating noise $\tilde{n}_{P_1}(f)$ & $\tilde{n}_{P_2}(f)$, respectively, then the total noise can be written

$$\tilde{n}(f) = \tilde{n}_{P_1}(f) + \tilde{n}_{P_2}(f). \quad (2.4)$$

The total PSD is simply the sum of the PSDs of the two processes, since from equation (2.3) we can write

$$\begin{aligned} \mathbb{E}[\tilde{n}(f')^* \tilde{n}(f)] &= \mathbb{E}[(\tilde{n}_{P_1} + \tilde{n}_{P_2})(f')^* (\tilde{n}_{P_1} + \tilde{n}_{P_2})(f)] \\ &= \mathbb{E}[\tilde{n}_{P_1}(f')^* \tilde{n}_{P_1}(f) + \tilde{n}_{P_1}(f')^* \tilde{n}_{P_2}(f) + \tilde{n}_{P_2}(f')^* \tilde{n}_{P_1}(f) + \tilde{n}_{P_2}(f')^* \tilde{n}_{P_2}(f)] \\ &= \mathbb{E}[\tilde{n}_{P_1}(f')^* \tilde{n}_{P_1}(f)] + \mathbb{E}[\tilde{n}_{P_2}(f')^* \tilde{n}_{P_2}(f)] \\ &= \frac{1}{2} \delta(f - f') [S_{P_1}(f) + S_{P_2}(f)], \end{aligned}$$

since we have $\mathbb{E}[\tilde{n}_{P_1}(f')^* \tilde{n}_{P_2}(f)] = \mathbb{E}[\tilde{n}_{P_1}(f')^*] \mathbb{E}[\tilde{n}_{P_2}(f)]$ by independence, and where we have assumed $\mathbb{E}[\tilde{n}_{P_1}(f')^*] = \mathbb{E}[\tilde{n}_{P_2}(f)] = 0$. Thus we can see that

$$S_n(f) = S_{P_1}(f) + S_{P_2}(f). \quad (2.5)$$

This argument extends to any number of (independent) noise sources, where the total PSD is the sum of the PSDs from different processes. In GW laser interferometers, as for many other sorts of signal detectors, the noise processes often add together in such a way as to produce a PSD ‘noise bucket’ or ‘sensitivity bucket’, with loud (relatively) low and high frequency noise overpowering signals at those frequencies, and quiet mid-frequency noise where the detector is sensitive to signals. Some approximate PSDs for detectors relevant to this thesis are shown in Figure 2.3.

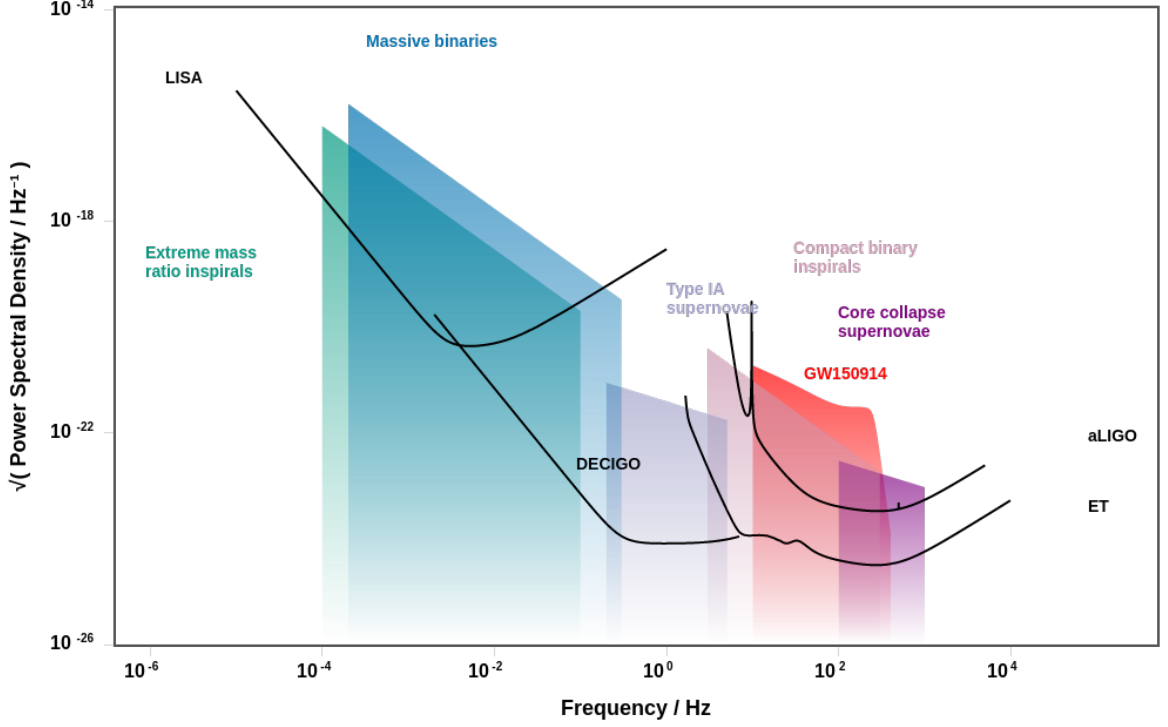


Figure 2.3: Approximate, sky averaged amplitude spectral densities (ASDs: the square root of the PSD) for LISA, DECIGO, ET & LIGO, along with some typical example GW sources. Source: gwplotter.com [68].

2.3 Parameter estimation

If a signal contains a waveform generated by some parameterised model, like that of our BHBs in Section 1.4, we can estimate those generating parameters by Bayesian methods. This section presents a slightly more in depth introduction than is usual, detailing and deriving some essential results for the methods introduced in Chapter 3.

2.3.1 Bayesian inference

As is commonplace in GW data analysis, we use Bayesian inference methods to define the *posterior* probability distribution of model parameters using Bayes' theorem,

$$p(\boldsymbol{\theta} | \mathbf{d}) = \frac{p(\boldsymbol{\theta}) p(\mathbf{d} | \boldsymbol{\theta})}{p(\mathbf{d})}, \quad (2.6)$$

where $p(\boldsymbol{\theta})$ is called the *prior* probability of the parameters, written as the vector $\boldsymbol{\theta}$, $p(\mathbf{d} | \boldsymbol{\theta})$ is the *likelihood*; the probability of obtaining data \mathbf{d} given parameters $\boldsymbol{\theta}$ and given a waveform model, $p(\mathbf{d})$ is the *evidence* for the model given data \mathbf{d} , and where the data

is taken to be the sum, $\mathbf{d} = \mathbf{s} + \mathbf{n}$, of the true signal, $\mathbf{s} = \mathbf{s}(\hat{\theta})$, and detector noise, \mathbf{n} , where $\hat{\theta}$ is the vector of signal parameters. We ignore the evidence here as we shall not be comparing different models; given data \mathbf{d} , $p(\mathbf{d})$ is a constant, and we shall require $p(\theta | \mathbf{d})$ up to proportionality only.

We may define our own prior based on expectations derived from other means; for the present study we will generally choose uniform priors with a suitably chosen ‘window’ centred on the true parameters. The remaining factor to be determined is the likelihood function. Assuming a Gaussian noise model, we can write the probability of some noise realisation \mathbf{n} , a vector of densities, occurring as

$$p(\mathbf{n}) = [(2\pi)^N \det(\boldsymbol{\Sigma})]^{-1/2} \exp\left[-\frac{1}{2}\langle \mathbf{n} | \mathbf{n} \rangle\right], \quad (2.7)$$

where we call $\boldsymbol{\Sigma}$ the (*noise*) *covariance matrix*, $N \equiv \dim(\mathbf{n})$, and where $\langle \cdot | \cdot \rangle$ is an inner product: the inner product and noise covariance matrix will be defined in both the time and frequency domains, and we will see how they are related in both continuous and discrete formalisms.

2.3.2 Inner products and transformations

Let us first consider the frequency domain, discrete inner product. We suppose that the individual frequency components of vector $\tilde{\mathbf{a}}$ are independent so that $\langle \tilde{\mathbf{a}}_i | \tilde{\mathbf{a}}_j \rangle \propto \delta_{ij}$. The standard expression of the multivariate Gaussian probability distribution is obtained if we define the inner product of a generally complex frequency domain noise vector, $\tilde{\mathbf{n}}$, as

$$\langle \tilde{\mathbf{n}} | \tilde{\mathbf{n}}^* \rangle \equiv \tilde{\mathbf{n}}^T \boldsymbol{\sigma}^{-2} \tilde{\mathbf{n}}^*, \quad (2.8)$$

where $\boldsymbol{\sigma}^2$ is the diagonal matrix (since we assume that the frequencies are uncorrelated) that encodes the variance of each frequency bin. Substituting this into equation (2.7), and making the identification $\boldsymbol{\Sigma} \equiv \boldsymbol{\sigma}^2$, we obtain the usual expression for the joint probability of finding each component of $\tilde{\mathbf{n}}$ in its respective frequency bin:

$$p(\tilde{\mathbf{n}}) = \left[\prod_{i=1}^N \sqrt{2\pi} \sigma_{ii} \right]^{-1} \exp\left[-\frac{1}{2}\tilde{\mathbf{n}}^T \boldsymbol{\sigma}^{-2} \tilde{\mathbf{n}}^*\right]. \quad (2.9)$$

Note that we define the (continuous) *two-sided PSD*, $S_n^{\text{two-sided}}(f)$, of detector noise, \tilde{n}

$$\mathbb{E}[\tilde{n}(f')^* \tilde{n}(f)] = \delta(f - f') S_n^{\text{two-sided}}(f), \quad (2.10)$$

where again \mathbb{E} is the expectation value operator and δ is the Dirac delta function. The discrete version of this can be written [69]

$$\mathbb{E}[\tilde{\mathbf{n}}_i \tilde{\mathbf{n}}_j^*] = \frac{N \delta_{ij}}{\Delta t} S_n^{\text{two-sided}}(f_i) = 2\sigma_{ij}^2, \quad (2.11)$$

where Δt is the sampling interval. We have taken the mean to be zero, and the factor of 2 on the right-hand-side comes from individual contributions from real and imaginary parts. Thus substituting for the variance matrix in the inner product, the discrete inner product of two vectors $\tilde{\mathbf{a}}$ & $\tilde{\mathbf{b}}$ may be written in terms of the PSD as

$$\langle \tilde{\mathbf{a}} | \tilde{\mathbf{b}} \rangle \equiv \Re \sum_{i>0} \frac{2\Delta t}{N} \frac{\tilde{\mathbf{a}}_i \tilde{\mathbf{b}}_i^*}{\frac{1}{2} S_n(f_i)} = \Re \sum_{i>0} 4\Delta f \frac{\tilde{a}(f_i) \tilde{b}^*(f_i)}{S_n(f_i)}, \quad (2.12)$$

since by taking the discrete Fourier Transform as an approximation of the continuous Fourier Transform, we make the identification $\tilde{\mathbf{a}}_i = \tilde{a}(f_i)/\Delta t$. The sum runs over the positive frequency bins present with size Δf , the tilde $\tilde{\cdot}$ denotes a Fourier Transform, and \Re returns the real part of its argument. One factor of 2 comes from the fact that we are summing over positive frequency bins only, and another comes from using the *one-sided PSD*, where $\frac{1}{2} S_n(f) \equiv S_n^{\text{two-sided}}(f)$ from equation (2.3).

Denoting the inner product of continuous functions with $\langle \cdot, \cdot \rangle$, we will now see how the continuous frequency domain inner product transforms to the continuous time domain inner product. The inner product of continuous functions clearly follows from the above discrete version by taking $\Delta f \rightarrow df$ and writing the sum as an integral. Thus

$$\langle \tilde{a}, \tilde{b} \rangle \equiv \Re \left(4 \int_0^\infty df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} \right).$$

However, it will be useful for the following to rewrite the continuous frequency domain inner product as

$$\langle \tilde{a}, \tilde{b} \rangle = \int_{-\infty}^\infty df \frac{\frac{1}{2} [\tilde{a}^*(f) \tilde{b}(f) + \tilde{a}(f) \tilde{b}^*(f)]}{\frac{1}{2} S_n(f)},$$

where it is implied that the real part of $\tilde{a}^* \tilde{b}$ is even. Denoting the Fourier Transform by \mathcal{F} , define a new function $\rho^{-1}(t)$ (the reasons for this will become clear) by the relation

$$\left(\frac{1}{2} S_n \right)^{-1} \equiv \mathcal{F} \{ \rho^{-1} \}. \quad (2.13)$$

Then writing the spectra as Fourier Transforms of their respective signals, the first inte-

gral term becomes

$$\begin{aligned}
 \int_{-\infty}^{\infty} df \frac{\frac{1}{2}\tilde{a}^*(f)\tilde{b}(f)}{\frac{1}{2}S_n(f)} &= \frac{1}{2} \int_{-\infty}^{\infty} df \mathcal{F}\{a(t)\}^* \cdot \mathcal{F}\{b(t)\} \cdot \mathcal{F}\{\rho^{-1}(t)\} \\
 &= \frac{1}{2} \int_{-\infty}^{\infty} df \mathcal{F}\{a(-t)\} \cdot \mathcal{F}\{b(t)\} \cdot \mathcal{F}\{\rho^{-1}(t)\} \\
 &= \frac{1}{2} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} dt' \int_{-\infty}^{\infty} dt'' a(-t) \cdot b(t') \cdot \rho^{-1}(t'') \cdot e^{-2\pi i f(t+t'+t'')} \\
 &= \frac{1}{2} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} dt' \int_{-\infty}^{\infty} dt'' a(-t) \cdot b(t') \cdot \rho^{-1}(t'') \cdot \delta(t+t'+t'') \\
 &= \frac{1}{2} \int_{-\infty}^{\infty} dt a(t) \int_{-\infty}^{\infty} dt' b(t') \cdot \rho^{-1}(t-t'),
 \end{aligned}$$

where we used the fact that $a(t)$ is real in the second line. We have only computed the first term of the $\langle \tilde{a}, \tilde{b} \rangle$ integral above, however, $a(t)$ and $b(t)$ are easily interchangeable throughout the preceding steps, so we find that

$$\begin{aligned}
 \langle \tilde{a}, \tilde{b} \rangle &\equiv \int_{-\infty}^{\infty} df \frac{\tilde{a}^*(f)\tilde{b}(f) + \tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} \\
 &= \int_{-\infty}^{\infty} dt a(t) \int_{-\infty}^{\infty} dt' b(t') \cdot \rho^{-1}(t-t') \\
 &\equiv \langle a, b \rangle.
 \end{aligned}$$

We now have both the frequency domain and time domain inner products of continuous functions. Since real data is never continuous however, the familiar form of the time domain inner product is the discrete version. Discretising the above

$$\langle a, b \rangle = \int_{-\infty}^{\infty} dt a(t) \int_{-\infty}^{\infty} dt' b(t') \cdot \rho^{-1}(t-t') \longrightarrow \sum_{i=-N/2}^{N/2-1} \Delta t \mathbf{a}_i \sum_{j=-N/2}^{N/2-1} \Delta t \mathbf{b}_j \rho^{-1}(t_i - t_j),$$

so that

$$\langle a, b \rangle \longrightarrow \langle \mathbf{a} | \mathbf{b} \rangle = \Delta t^2 \mathbf{a}_i \rho_{ij}^{-1} \mathbf{b}_j, \quad (2.14)$$

where the matrix $\rho_{ij}^{-1} \equiv \rho^{-1}(t_i - t_j)$ and summation over repeated indices is implied in the last step. Thus we have all desired conversions between frequency and time domain samples presented in the different bases.

The time domain version of the posterior in (2.9) can also be derived via a matrix basis transformation, specifically the (inverse) discrete Fourier transform (DFT) matrix, under which the determinant of the covariance matrix is unchanged. That is, there is

a DFT transformation matrix \mathbf{P} , given by $P_{jk} \equiv \Delta t \exp(-2\pi ijk/N)$, such that (introducing superscripts for the time domain, Σ^{td} , and frequency domain, Σ^{fd} , inner product matrices)

$$(\Sigma^{\text{fd}})^{-1} = \mathbf{P}(\Sigma^{\text{td}})^{-1}\mathbf{P}^{-1},$$

where we have that

$$\det(\Sigma^{\text{fd}})^{-1} = \det(\mathbf{P}) \det(\Sigma^{\text{td}})^{-1} \det(\mathbf{P}^{-1}) = \det(\Sigma^{\text{td}})^{-1}.$$

We can then replace $\prod_{i=1}^N \sigma_{ii} = \sqrt{\det(\Sigma^{\text{fd}})}$ in (2.9) with $\sqrt{\det(\Sigma^{\text{td}})}$. For the inner product in the exponent in (2.9), note that $\mathbf{P}\mathbf{P}^{-1} = \mathbb{1}$, so we can write

$$\tilde{\mathbf{n}}^{*\text{T}}(\Sigma^{\text{fd}})^{-1}\tilde{\mathbf{n}} = (\tilde{\mathbf{n}}^{*\text{T}}\mathbf{P})(\mathbf{P}^{-1}(\Sigma^{\text{fd}})^{-1}\mathbf{P})(\mathbf{P}^{-1}\tilde{\mathbf{n}}) = \mathbf{n}^{\text{T}}(\Sigma^{\text{td}})^{-1}\mathbf{n} = \langle \mathbf{n} | \mathbf{n} \rangle, \quad (2.15)$$

since $\mathbf{P}\mathbf{n} = \tilde{\mathbf{n}}$, and $\mathbf{P}^{-1} = \mathbf{P}^*$ so that $\tilde{\mathbf{n}}^{*\text{T}}\mathbf{P} = (\mathbf{P}^{-1}\tilde{\mathbf{n}})^{*T}$. Comparing this with (2.14) gives

$$(\Sigma^{\text{td}})^{-1}_{ij} = \Delta t^2 \rho_{ij}^{-1}. \quad (2.16)$$

Given this transformation between domains, the origin of the discrete time domain inner product is now clear. In the next sections we compute Σ^{td} using the PSD directly. To conclude the transformation of the likelihood in (2.9), we replace the inner product with the time domain version to see that the likelihood of obtaining noise \mathbf{n} is, naturally, independent of basis and may also be written

$$p(\tilde{\mathbf{n}}) = p(\mathbf{n}) = \left[(2\pi)^N \det(\Sigma^{\text{td}}) \right]^{-1/2} \exp \left[-\frac{1}{2} \mathbf{n}^{\text{T}}(\Sigma^{\text{td}})^{-1}\mathbf{n} \right]. \quad (2.17)$$

Since the data is ordinarily written as $\mathbf{d} = \mathbf{s} + \mathbf{n}$, we have that $\mathbf{n} = \mathbf{d} - \mathbf{s}$, and may think of $p(\mathbf{d} - \mathbf{h}(\theta))$ as the likelihood that the remainder of the data left over, after subtracting the model template \mathbf{h} with parameters θ , can be accounted for completely as detector noise. That is, when $\mathbf{h}(\theta) = \mathbf{s}$, we have $p(\mathbf{d} - \mathbf{h}) = p(\mathbf{n})$. The likelihood is thereby defined to be

$$\begin{aligned} p(\mathbf{d} | \theta) &\equiv p(\mathbf{d} - \mathbf{h}(\theta)) \\ &= \left[(2\pi)^N \det(\Sigma^{\text{td}}) \right]^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{h}(\theta))^{\text{T}} (\Sigma^{\text{td}})^{-1} (\mathbf{d} - \mathbf{h}(\theta)) \right], \end{aligned} \quad (2.18)$$

We will switch between the matrix product and standard inner product notation freely, when convenient. The vector $\mathbf{r}(\theta) \equiv \mathbf{d} - \mathbf{h}(\theta)$ will be referred to as the *residual*.

2.3.3 Noise covariance matrix

The noise power generated by the detector as a function of frequency (the PSD) determines the degree to which any two samples in the time domain representation of the noise are correlated with each other. In the current section and the next, we derive the time-domain inner product operator (the matrix $(\Sigma^{\text{td}})^{-1}$) from the PSD to account for this interdependence of the samples. It will be assumed that the noise is *wide sense stationary* (WSS), that is, that its mean and covariance are time invariant.

A *cross-covariance matrix* of two vectors describes the covariance between elements of the vectors. The i, j element of the cross-covariance matrix $\Sigma_{\mathbf{X}\mathbf{Y}}$ is the covariance between the i^{th} element of \mathbf{X} and the j^{th} element of \mathbf{Y} . A cross-covariance matrix is called an *auto-covariance*, or simply *covariance matrix* when describing the covariance between different elements of the same vector. Our noise covariance matrix will be denoted here as Σ , so that

$$\Sigma \equiv \Sigma_{\mathbf{nn}}, \quad (2.19)$$

where \mathbf{n} is a noise realisation. In terms of matrix elements, we define a cross-covariance matrix as

$$\Sigma_{ij} \equiv \text{cov}[X_i, Y_j] = \mathbb{E}[(X_i - \mathbb{E}[X_i])(Y_j - \mathbb{E}[Y_j])], \quad (2.20)$$

where \mathbb{E} is the expectation value operator. For vanishing mean of the noise, we have

$$\Sigma_{ij} = \mathbb{E}[n_i n_j] \equiv R_{ij}, \quad (2.21)$$

defining the *autocorrelation matrix*, R_{ij} . Now since the noise is assumed to be WSS, the expected value of the product of any two samples only depends on the time ‘lag’ between the samples, i.e., $\mathbb{E}[n_i n_j] = \mathbb{E}[n_{i+c} n_{j+c}]$ for any integer c . It will be useful to consider $c = -i$ where we have that

$$R_{ij} = \mathbb{E}[n_{i-i} n_{j-i}] = \mathbb{E}[n_0 n_{j-i}]. \quad (2.22)$$

In this form, we can see the matrix is both Toeplitz and symmetric. It is clear that the i, j element represents the expected value of noise correlations of time series samples separated by the absolute time lag between the i^{th} and j^{th} sample. The first row, $R_{0i} = \mathbb{E}[n_0 n_i]$, gives precisely the definition of the (discrete) *autocorrelation function* (ACF).

By the *Wiener-Khinchin theorem*, the ACF is the Fourier Transform of the PSD. However, we must be careful in the GW data analysis world since, again, we most often use

the one-sided PSD which introduces a factor of 2, such that the equation embodying the Wiener-Khinchin theorem reads

$$S_n(f) = 2 \int_{-\infty}^{\infty} dt R(t) e^{-2\pi i f t} = 2 \mathcal{F}\{R\}(f). \quad (2.23)$$

Thus given the PSD we can easily compute Σ , since it is defined by $R_{ij} = R(t_i - t_j)$. For the inner product however, we require Σ^{-1} . When $N = \dim(\mathbf{d})$ becomes large, numerical inversion of the resulting $N \times N$ noise covariance matrix becomes very unwieldy. We have devised a fast and accurate solution to this problem, which we now describe.

2.3.4 Inverting large autocorrelation matrices

We wish to find the matrix Σ^{-1} that satisfies $\Sigma \Sigma^{-1} = \mathbb{1}$. First, note that since Σ is symmetric Toeplitz, so then must be Σ^{-1} . In components, write

$$\mathbb{1}_{ik} = \Sigma_{ij} \Sigma_{jk}^{-1} = \Sigma_{ij} \Sigma_{kj}^{-1} = \Sigma_{0j} \Sigma_{(k-i)j}^{-1} = R_{0j} R_{(k-i)j}^{-1}, \quad (2.24)$$

where summation over j is implied by the repeated indices. It is clearer in the continuous regime, where the above can be written

$$\delta(t - t') = \int_{-\infty}^{\infty} d\tau R(t - \tau) R^{(-1)}(\tau - t') = \int_{-\infty}^{\infty} d\tau R(t - t' - \tau) R^{(-1)}(\tau), \quad (2.25)$$

where δ is the Dirac-delta function, R is the ACF, known by the Wiener-Khinchin theorem, and where $R^{(-1)}$ is the function we seek, which we call the *inverse autocorrelation function*. The second equation comes from setting $\tau - t' \rightarrow \tau$ in the integral. Now define $T \equiv t - t'$ to obtain

$$\delta(T) = \int_{-\infty}^{\infty} d\tau R(T - \tau) R^{(-1)}(\tau). \quad (2.26)$$

This integral is precisely the form of a convolution ($*$ operator) of R with $R^{(-1)}$, so we can write

$$\delta(t) = (R * R^{(-1)})(t). \quad (2.27)$$

If \mathcal{F} denotes the Fourier Transform and \mathcal{F}^{-1} its inverse, then the convolution theorem states that

$$\mathcal{F}\{R * R^{(-1)}\} = \mathcal{F}\{R\} \cdot \mathcal{F}\{R^{(-1)}\}, \quad (2.28)$$

so that after Fourier transforming (2.27), using the convolution theorem, and some rearranging, we obtain

$$R^{(-1)} = \mathcal{F}^{-1} \left\{ \left(\mathcal{F}\{R\} \right)^{-1} \right\} = \mathcal{F}^{-1} \left\{ \left(\frac{1}{2} S_n \right)^{-1} \right\}, \quad (2.29)$$

where we substituted for the PSD from (2.23). The inverse ACF coincides with the definition of the new function in (2.13), where $R^{(-1)} = \rho^{-1}$, from which it follows that the *inverse autocorrelation matrix* R_{ij}^{-1} is the time-domain inner product operator ρ_{ij}^{-1} . The discrete form of $R^{(-1)}$ is then related to the first row of $(\Sigma^{\text{td}})^{-1}$ by (2.16), which then sets the whole matrix (since we already know it is symmetric Toeplitz), and we can quickly compute $(\Sigma^{\text{td}})^{-1}$. A realistic realisation of the LISA PSD is shown in Fig 2.4, the inverse ACF of that PSD is shown in Fig 2.5.

Aside: An algorithm for inversion of general Toeplitz matrices

As a brief aside, we point out an interesting result that follows from the method of inverting symmetric Toeplitz matrices: that it is possible to extend this approach to determine the inverses of general Toeplitz matrices. Suppose then that we wish to determine the inverse of the general, non-symmetric Toeplitz matrix \mathbf{F} . We can embed \mathbf{F} into a larger block matrix

$$\mathbf{A} = \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix}, \quad (2.30)$$

which we define, by choice of \mathbf{E} , \mathbf{G} , \mathbf{H} , to be symmetric Toeplitz. Then we require

$$\begin{aligned} \mathbf{G} &= \mathbf{F}^T, \\ \mathbf{E} &= \mathbf{E}^T = \mathbf{H}, \end{aligned}$$

and where \mathbf{E} is chosen such that its entries are consistent with those of \mathbf{F} to make \mathbf{A} Toeplitz (note that the value on the leading diagonal can be chosen freely). I.e., for \mathbf{E} an $n \times n$ matrix: $E_{ni} = G_{0(i+1)}$, for $i < n$. Since \mathbf{A} is symmetric and Toeplitz, we can apply the DFT method to find its inverse, \mathbf{A}^{-1} . However, it is also known that the inverse of the block matrix \mathbf{A} defined by four ‘quadrant’ block matrices is equal to

$$\mathbf{A}^{-1} = \begin{pmatrix} \mathbf{E}^{-1} + \mathbf{E}^{-1}\mathbf{F}\mathbf{S}^{-1}\mathbf{G}\mathbf{E}^{-1} & -\mathbf{E}^{-1}\mathbf{F}\mathbf{S}^{-1} \\ -\mathbf{S}^{-1}\mathbf{G}\mathbf{E}^{-1} & \mathbf{S}^{-1} \end{pmatrix}, \quad (2.31)$$

where $\mathbf{S} \equiv \mathbf{H} - \mathbf{G}\mathbf{E}^{-1}\mathbf{F} = \mathbf{E} - \mathbf{F}^T\mathbf{E}^{-1}\mathbf{F}$ (see *Schur complementation*, [70]). Now, because \mathbf{A} is symmetric Toeplitz, then so is \mathbf{A}^{-1} , and so too are the upper left and lower right quadrants of \mathbf{A}^{-1} . Indeed, these two leading diagonal quadrants must be equal.

That is

$$\mathbf{S}^{-1} = \mathbf{E}^{-1} + \mathbf{E}^{-1}\mathbf{F}\mathbf{S}^{-1}\mathbf{F}^T\mathbf{E}^{-1},$$

which can be rearranged to give the inverse of our original matrix \mathbf{F} , as

$$\mathbf{F}^{-1} = \mathbf{S}^{-1}\mathbf{F}^T\mathbf{E}^{-1}(\mathbf{E}\mathbf{S}^{-1} - \mathbf{1})^{-1}, \quad (2.32)$$

where we also have that $\mathbf{E}\mathbf{S}^{-1} - \mathbf{1}$ is symmetric Toeplitz, and can also be inverted using the DFT method we previously described. Note that $\mathbf{S}^{-1}\mathbf{F}^T\mathbf{E}^{-1}$ is already given by the negative of the lower left quadrant of \mathbf{A}^{-1} , and \mathbf{S}^{-1} by the upper left/lower right of \mathbf{A}^{-1} . So we have found that the inverse of *any* Toeplitz matrix \mathbf{F} can be obtained simply by two DFTs and two matrix multiplications. Note that we have *assumed* invertibility of \mathbf{F} ; a test of invertibility of Toeplitz matrices is given in Ref. [71], where some special case solutions of Toeplitz matrix inverses are also given. An iterative procedure for determining the inverse of general Toeplitz matrices was proposed in Ref. [72], although the approach presented here appears considerably simpler.




Figure 2.4: A sample power spectral density (PSD) function for LISA. A detailed account of the origin of the structure of this function (and the source code defining it) can be found in [62].




Figure 2.5: The (one-sided) *inverse autocorrelation function* for the LISA PSD (shown in Fig 2.4). This defines the continuous inner product operator and shows that samples separated by around $\gtrsim 140$ seconds are not significantly correlated in comparison to the zero time lag value.

2.4 Sampling

With the likelihood function in equation (2.18) and hence the posterior probability distribution in equation (2.6) defined, we would like to be able to analyse the posterior in some way, perhaps finding the mean, or the variances/covariances of/between parameters, for example. Due to the often highly intricate form of the posterior function, it is generally exceedingly expensive to evaluate the posterior over a discrete lattice of points, and it is generally not possible to write down short, simple expressions for the desired quantities that are reasonably easy to compute.

The standard solution is to ‘draw a sample set’ (a discrete set of posterior samples) from the posterior distribution such that the sample density is proportional to the true (continuous) posterior distribution. This discrete set of points is then used as a representation of the true posterior. However, doing this is not at all straightforward, since a process devised to produce the random samples can only do so accurately if the posterior distribution is already known. Various techniques exist for approximating the posterior sample set however; algorithms that perform this task are known as *samplers*. There are many variations on the details of the operations of different samplers. These are not so

important to know for the particular aims of this thesis however, so we shall only briefly describe basic notions of two sampling methods: *MCMC sampling*, and *nested sampling*.

2.4.1 Markov chain Monte-Carlo sampling

Markov chain Monte Carlo (MCMC) samplers use *walkers* (the algorithm that proposes, tests and selects posterior samples) to generate *chains* (ordered sets) of posterior samples. An example of the procedure is as follows:

1. A walker will first choose a random starting point, θ_0 , in the parameter space, Θ , and evaluate the posterior there, $P(\theta_0)$: this point is set as the first element, $c_0 = \theta_0$, in the ‘chain’, $\mathcal{C} = \{c_0\}$. Set $i = 1$.
2. The walker proposes another random point, θ_i , and evaluates the posterior, $P(\theta_i)$, there. If $P(\theta_i) > P(\theta_{i-1})$, then $c_i = \theta_i$. Otherwise, if $P(\theta_i)/P(\theta_{i-1}) < U(0, 1)$, for $U(0, 1)$ a uniform distribution between 0 and 1, then $c_i = \theta_i$, else $c_i = \theta_{i-1}$. Add c_i to \mathcal{C} .
3. Increment i . If $i < n$, for n some chosen chain length, go back to step 2, otherwise, return chain \mathcal{C} .

There are finer details about how new points are proposed and tested by the walker. For example, points are usually proposed from a Gaussian centred on the previous point, making ‘distant jumps’ unlikely and small steps likely (hence the name ‘walker’). The design of these algorithms ensures that walkers tend towards higher probability regions of the parameter space, which are explored in more detail. More importantly however, since the probability of acceptance (after enough initial steps have been made) is roughly equal to the ratio of probabilities between steps, the accepted points correctly sample the posterior. Examples of MCMC samplers include: `emcee` [73], and `Bilby_MCMC` [74].

2.4.2 Nested sampling

Nested sampling is a method used for estimating the *evidence* (briefly mentioned in Section 2.3.1) from which one also obtains a set of posterior samples ‘for free’. As mentioned above, it is a difficult task to emulate drawing samples from a distribution when that distribution is not known. Nested sampling replaces this task with a series of simple tasks:

a series of successive, nested iso-likelihood contours (closed, $(n - 1)$ -dimensional surfaces in an n -dimensional parameter space, of equal likelihood) are created at increasing likelihood values. At each stage/level/contour, a new set of samples can be proposed relatively straightforwardly; the proposal distribution is just the prior distribution, constrained within the bounds of (i.e., greater than) the current iso-likelihood value. One then acquires consecutive sets of samples from within each iso-likelihood contour, which can be combined (using ‘weights’ associated with each region [75]) to reconstruct the posterior. Examples of nested sampling samplers include: `Nessai` [76], `CPNest` [77], and `dynesty` [78]. We use `Nessai` (with `Bilby` [79]) to provide all posteriors in Chapters 6, 7 & 8 (see Section 7.1).

The Large Dataset Problem

3.1 Introduction

In order to perform explorative analyses of very large, purely simulated data sets, with limited computing power and on manageable timescales, we investigate a means of reducing computational costs without jeopardising the quality of results. Our approach is to discard a number of data points and modify the properties of the remaining data, or the methods used to analyse it, in order to compensate for the original information content of interest. That is, we wish to ascertain the robustness, limitations and benefits of what we call *downsampling* of the given data set, followed by making necessary modifications to fulfil the requirement that the results of, in particular, parameter estimation (PE) analyses will be minimally affected. With this view, the task is an optimisation problem. Using fewer data points means fewer numerical operations are required to evaluate the likelihood function at a given point in the parameter space. We will see that the consequence of this is a drastic reduction in the amount of time required to compute the posterior probability distribution. The particular focus here is the study of ‘slowly evolving’ GW signals that we expect to obtain from the LISA detector [80] (with special emphasis on *compact binary inspirals*), but the methods employed here are general and may be used with other datasets of signals that can be said to be slowly evolving in time.

We can loosely define a waveform as being *slowly evolving* over a time interval \mathcal{S} if the average value of the *Fisher Information Matrix* (see Section 3.3.2) at time t over ‘intermediate duration’ time intervals \mathcal{J}_t (with a range of order ~ 100 times the reciprocal of the

Nyquist rate of the waveform in \mathcal{S}) centred at t is approximately constant for all $t \in \mathcal{S}$. A definition of this sort is required in order to be able to classify sinusoidal waveforms as being slowly evolving if their frequency and amplitude parameters are roughly constant, despite their oscillating nature; their Fisher information matrices also oscillate as functions of time, but taking the average value over intermediate time intervals ‘smooths out’ the oscillations.

To see why slow evolution is important, consider the simplest time-series signal that is clearly slowly evolving: a constant signal. One could take time series samples from any location; these would provide roughly equal information on that constant (depending on noise). Conversely, consider a signal which has an instantaneous step from one constant to another at a certain time; one should clearly classify this as not slowly evolving. To accurately determine the location of the step, the time samples just before and after the step would be preferred for optimal information on the location of the step. Downsampling in this case does not necessarily fail, but it would at least require a carefully constructed, signal-dependent setup to ensure the downsampled likelihood function is approximately equal to the original likelihood function across the whole parameter space. We aim to introduce a signal independent downsampling procedure, and thus we are restricted to signals of slow time evolution.

The inspiral stage of BHBs may last millions of years before merger, where the vast majority of the duration of GW emission constitutes a very slowly evolving signal of gradually increasing amplitude and frequency (a characteristic ‘chirp’ [81]). The LISA detector will be sensitive to a good portion of the lower frequency ($\sim 1\text{--}100$ mHz) part, and may be active for $\sim 5\text{--}10$ years [80]. For the most slowly evolving signals hovering near the upper frequency sensitivity of ~ 100 mHz throughout LISA’s lifetime, we may realistically expect datasets consisting of $\sim 10^9$ data points at the Nyquist rate.

Such a large dataset would require a great deal of computing resources for PE analyses, but since for preparatory research the entire dataset (just meaning both the signal and noise) is simulated, we have freedom to alter or redefine the signal and/or detector model at will to suit our investigative requirements. In particular, that liberty is made use of by the definition and application of the downsampling procedure that is the subject of this chapter. We shall therefore attempt to be comprehensive in the description and provision of such a scheme, and it shall be shown to be able to accurately reproduce

the results one can expect to obtain from a real-world analysis of a real-world dataset. For real-world scenarios, this procedure is not applicable since we have no control of either the signal or the noise. This procedure is only applicable in simulations and will therefore serve only to provide experimental insight into which sorts of analyses of large data sets could be the most fruitful, and the results we can expect to see when real-world analyses are performed on the real-world dataset.

It is instructive to place this method and its results in the context of related works that employ other approaches in either making approximations, or in some other way speed up the evaluation of PE analyses. Perhaps the most frequently employed formalism is the Fisher information approximation; see for example [82, 83, 84, 85, 86, 87]. This formalism provides a multivariate Gaussian approximation of the likelihood that gives a lower bound in the standard error of the signal’s true parameters. This can be of some use, but in many cases is a poor approximation of the posterior; for a detailed discussion on why this is not generally reliable (except in some particular regimes), see [88, 89] and Section 6.2.1 for further details. Other methods that are not approximations, but exploit common properties of the waveforms under study by finding ways to cheaply produce waveform templates include *reduced order quadrature (ROQ)* [90], *heterodyning/relative binning* [91], and so-called *adaptive frequency resolution* [92]. Finally, another approximation technique, perhaps most similar to that developed here, is *Template-Interpolation* [93]. This entails computing the frequency domain waveform, except in those parts which are known to be roughly linear, and interpolating over those gaps in the waveform between the points that are explicitly computed. We shall of course describe downsampling in detail, but one quick way to see how downsampling differs from this is to see that (after the original, complete dataset is used to inform the setup) the gaps in the data are fully discarded, rather than ‘filled in’; the definitions of signal/noise/inner product operations, or whatever else might be required, are modified such that the remaining datapoints constrain the signal parameters just as well as the fully sampled dataset would have.

This chapter is structured as follows. In Section 3.2, we introduce the framework in which to define the problem and useful results required for the analysis. We then define and develop the downsampling procedure in Section 3.3, discussing some of the concepts and problems that must be considered for a faithful reproduction of the LISA posterior. The procedure defined in Section 3.3 is tested for accuracy experimentally in Section

3.4, using two different measures of posterior convergence and lays out the convergence criteria we have adopted. Section 3.5 presents the projected improvements in the rate of evaluation of PE attained by downsampling.

3.2 Preliminaries

This section consists of a few details, results and definitions that are required to be clarified before moving on to defining and testing the downsampling procedure. The materials presented here are extensions of the foundational results presented in Sections 2.2 & 2.3.

3.2.1 Whitening transformation matrix

We will require an operator that transforms vectors of correlated samples to vectors of uncorrelated samples. The same mathematical techniques used to find the inverse auto-correlation function we used in Section 2.3.4 can be employed to do this.

Consider representing a noise weighted (time domain) inner product of signals in another basis in which each of the signal's samples are entirely uncorrelated from each other. The inner product matrix in this space is a diagonal matrix. Further suppose that the matrix is the identity matrix, so that each sample contributes to the inner product sum with equal weight. Since inner products remain invariant under basis transformations, we can define a transformed inner product of vectors \mathbf{a} and \mathbf{b} by

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \mathbf{b} = \bar{\mathbf{a}}^T \mathbb{1} \bar{\mathbf{b}}, \quad (3.1)$$

for general time domain inner product matrix $\boldsymbol{\Sigma}^{-1}$, and where $\bar{\mathbf{a}}$ and $\bar{\mathbf{b}}$ are the 'whitened' representations of vectors \mathbf{a} and \mathbf{b} . We shall solve this here by writing

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \mathbf{b} = \mathbf{a}^T \boldsymbol{\Sigma}^{(-1/2)} \boldsymbol{\Sigma}^{(-1/2)} \mathbf{b} = \mathbf{a}^T \left(\boldsymbol{\Sigma}^{(-1/2)} \right)^T \mathbb{1} \boldsymbol{\Sigma}^{(-1/2)} \mathbf{b}, \quad (3.2)$$

where $\boldsymbol{\Sigma}^{(-1/2)}$ is some matrix to be found, which we shall constrain to be symmetric Toeplitz. We can then read off the whitening transformation of the vector space, in which $\boldsymbol{\Sigma}^{-1} \rightarrow \mathbb{1}$ along with all vectors $\mathbf{v} \rightarrow \bar{\mathbf{v}} = \boldsymbol{\Sigma}^{(-1/2)} \mathbf{v}$. Notice, we have that

$$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}^{(-1/2)} \boldsymbol{\Sigma}^{(-1/2)}, \quad (3.3)$$

which in the continuous case can be written as

$$R^{(-1)}(t-t') = \int_{-\infty}^{\infty} d\tau R^{(-1/2)}(t-\tau)R^{(-1/2)}(\tau-t') = \int_{-\infty}^{\infty} d\tau R^{(-1/2)}(t-t'-\tau)R^{(-1/2)}(\tau), \quad (3.4)$$

where $R^{(-1)}$ is the inverse auto-correlation function we defined earlier in equation 2.29, and $R^{(-1/2)}$ shall be referred to either as the *whitening function* or *decorrelation function* and is the function sought. Define $T \equiv t - t'$ so that

$$R^{(-1)}(T) = \int_{-\infty}^{\infty} d\tau R^{(-1/2)}(T - \tau)R^{(-1/2)}(\tau) = R^{(-1/2)} * R^{(-1/2)}(T). \quad (3.5)$$

The Fourier transform gives

$$\mathcal{F}\{R^{(-1)}\} = \left[\mathcal{F}\{R^{(-1/2)}\} \right]^2, \quad (3.6)$$

so that

$$R^{(-1/2)} = \mathcal{F}^{-1} \left\{ \sqrt{\mathcal{F}\{R^{(-1)}\}} \right\}. \quad (3.7)$$

However, we know that $\mathcal{F}\{R^{(-1)}\} = 2/S_n$, as in equation (2.29) for example, so we finally write the whitening function as

$$R^{(-1/2)} = \mathcal{F}^{-1} \left\{ \sqrt{2/S_n} \right\}, \quad (3.8)$$

where the matrix version $R_{ij}^{(-1/2)} \equiv R^{(-1/2)}(t_i - t_j)$. The whitening function for the LISA PSD is shown in Figure 3.1. From (3.4) we can then write

$$R_{ik}^{-1} = R_{ij}^{(-1/2)} R_{jk}^{(-1/2)} \quad (3.9)$$

with which, using (2.16) and (3.3) we can find the *whitening matrix*, or *decorrelation matrix*,

$$\Sigma_{ij}^{(-1/2)} = \Delta t R_{ij}^{(-1/2)}. \quad (3.10)$$

3.2.2 Data vector spaces

The inverse noise covariance matrix can be thought of as a metric in the space of data vectors, \mathcal{D} , giving rise to the notion of distance. The so-called *Mahalanobis distance* between points in \mathcal{D} represented by the position vectors \mathbf{X}, \mathbf{Y} (this interpretation is valid when Σ is constant) is given by

$$d_M(\mathbf{X}, \mathbf{Y}) = \sqrt{(\mathbf{X} - \mathbf{Y})^\top \Sigma^{-1} (\mathbf{X} - \mathbf{Y})}. \quad (3.11)$$




Figure 3.1: The (one-sided) *whitening* or *decorrelation* function, derived from the LISA PSD given in Figure 2.4. This is defined in equation (3.8) and whitens (continuous) LISA noise via convolution.

The likelihood in (2.18) can thus be written

$$p(\mathbf{d} | \theta) = [(2\pi)^N \det(\boldsymbol{\Sigma})]^{-1/2} \exp \left[-\frac{1}{2} d_{\mathbf{M}}(\mathbf{d}, \mathbf{h}(\theta))^2 \right]. \quad (3.12)$$

Recall that what we are trying to achieve is retaining the form of the likelihood function after discarding samples. This amounts to demanding that the Mahalanobis distance is unchanged after the downsampling procedure (yet to be defined). Suppose we have some initial (‘full’) dataset \mathbf{d}_f consisting of $N_f = \dim(\mathbf{d}_f)$ samples, and a downsampled dataset (a subset of the provided initial dataset) \mathbf{d}_s with $N_s = \dim(\mathbf{d}_s)$ samples. What we would like, ideally, is to have

$$p(\mathbf{d}_s | \theta) \stackrel{?}{=} p(\mathbf{d}_f | \theta),$$

for all θ , where $\stackrel{?}{=}$ indicates that we seek a definition $p(\mathbf{d}_s | \theta)$ that satisfies the above equation. Aside from an inconsequential proportionality constant which can be ignored, this means we require

$$d_{\mathbf{M},s}(\mathbf{d}_s, \mathbf{h}_s(\theta)) \stackrel{?}{=} d_{\mathbf{M},f}(\mathbf{d}_f, \mathbf{h}_f(\theta)), \quad (3.13)$$

where the subscripts describe whether the object lives in the full data space, or the subset data space.

Seemingly, this is a poorly defined and analytically impossible problem; it constitutes an underdetermined system of equations. It is only possible to ensure the distances are

equal by either redefining the signal model $\mathbf{h}_s \rightarrow \mathbf{h}_s' = \mathbf{h}_s'(\mathbf{d}_f, \mathbf{h}_f(\boldsymbol{\theta}))$, or defining $d_{M,s}$ with a covariance matrix $\boldsymbol{\Sigma}_s^{-1} = \boldsymbol{\Sigma}_s^{-1}(\mathbf{d}_f, \mathbf{h}_f(\boldsymbol{\theta}))$, as a function of all of the original data. This almost certainly *increases* the difficulty of the problem, and we must instead consider approximation methods. These will be discussed in subsequent sections however, the main point of this section is simply to introduce some of the geometric intuition behind the data vector spaces that will be helpful for our definition of the downsampling procedure.

Finally, consider the N_f dimensional ('full') data space, \mathcal{D}_f , with its metric $\boldsymbol{\Sigma}_f^{-1}$. This is provided by the detector model and data \mathbf{d}_f , and is considered fixed. The N_s dimensional subset space, \mathcal{D}_s , and metric, $\boldsymbol{\Sigma}_s^{-1}$, can be considered 'invented', (perhaps the subset signal model, \mathbf{h}_s , too) with the demand that (3.13) is satisfied, at least to some acceptable degree of accuracy to be determined. As a first attempt for finding the form of $\boldsymbol{\Sigma}_s^{-1}$, one may envisage simply deleting those rows and columns from $\boldsymbol{\Sigma}_f^{-1}$ corresponding to the samples dropped in going from $\mathbf{d}_f \rightarrow \mathbf{d}_s$, so that correlations between the remaining samples are still accounted for in the resulting $N_s \times N_s$ matrix $\boldsymbol{\Sigma}_s^{-1}$. It is not strictly necessary to keep account of sample correlations; our only aim is to satisfy (3.13) as closely as possible in some way. Unfortunately however, we will see that it is not so simple as naively choosing a few time-series samples, since it turns out that to accurately reproduce a LISA likelihood, with its implications of correlations between neighbouring time-domain samples, one must first *whiten* (decorrelate) the residual before downsampling, or else equivalently account for the correlations directly when computing the downsampled likelihood (see Section 3.3.1.3 for more detail).

3.2.3 Note on time vs. frequency domain

It is appropriate to comment briefly on the choice of performing this study on time-domain signals, rather than frequency domain. This is a significant point, as a great deal of signal analysis, and particularly in the field of GW data analysis, is performed in the frequency domain. A primary reason for this is that the inner product of vectors is especially easy to calculate in the frequency domain due to the fact that samples are uncorrelated there.

A guiding initial assumption was that samples remaining after downsampling would be separated in time to such a high extent that any sample correlations could be com-

pletely neglected. The initial guess of the form of Σ_s^{-1} suggested above would thus be proportional to the identity matrix, to a very close approximation. We could then choose $\Sigma_s^{-1} \propto \mathbf{1}$, and consider later revisions if required. Consequently, we expected no detrimental effects to the likelihood evaluation speed by working in the time domain. However we found (as described in Section 3.3.1.3) that first performing the whitening transformation is required, increasing the number of data points at which the residual must be computed, and so the number of operations required to compute the inner product. This hampers the evaluation time somewhat, but not so much as to render the time domain approach ineffective.

Another advantage was expected from a time domain analysis over the frequency domain: fewer samples were expected to be required to reproduce the likelihood function accurately. This is because, for the slowly evolving signals we are interested in, the (Fisher) information in the time domain is far more uniformly distributed than it is in the frequency domain. To retain a precise ratio of information on each parameter then, it is expected that, while samples can be selected at random from the time series, either the samples should have to be more carefully selected from the frequency series, so as not to constrain some parameters too much, and others too little, or more samples would have to be chosen.

In hindsight, this may not have made a great deal of difference since we are forced to use more samples for decorrelation in the time domain; the number of frequency domain samples required may well be of a similar number. A frequency domain version of the downsampling procedure would also be useful and will likely be developed. However, there are still very useful benefits of the time domain version, especially, for example, when a signal spans a highly non-uniform region of a PSD or if parameter information is highly non-uniform in the frequency domain, and, perhaps more practically, from the perspective of ease of producing (or modifying) time series models.

3.2.4 Divergences

We will require a measure of the amount by which two distributions differ. There are a number of ways one can go about this (a *P-P plot* is one example) but perhaps the most robust and meaningful are the measures known as *divergences*. There are different interpretations of divergences between distributions. Consider, for instance, the *Kullback-*

Liebler (KL) divergence, $D_{\text{KL}}(P \parallel Q)$, between (discrete) distributions P and Q . The KL divergence is often thought of as being a measure of the (mean) extra *information*, measured in bits, required to describe the distribution P using a code which is optimal for encoding distribution Q . Equivalently, more technically, it is the difference between the self-information or *entropy* of P , given by

$$H(P) \equiv - \sum_{i \in \mathcal{J}} P_i \log_2 P_i,$$

where \mathcal{J} is the set of possible outcomes and P_i is the probability of outcome i , which measures the average information conveyed about P and by P , and the *cross-entropy* of P and Q , given by

$$H(P, Q) \equiv - \sum_{i \in \mathcal{J}} P_i \log_2 Q_i,$$

which measures the average information conveyed about P , but carried by Q . Thus if $Q \neq P$, then one requires extra information to convey P using Q , but as $Q \rightarrow P$, the amount of extra information (i.e., the KL divergence) goes to zero. We can write the KL divergence as

$$D_{\text{KL}}(P \parallel Q) = H(P, Q) - H(P).$$

Another useful interpretation is that, if we suppose Q is a sample distribution drawn from P , then the KL divergence can be understood as the amount of *surprise* that Q was drawn from P . The divergences also have formulations for continuous variables, the KL divergence being written:

$$D_{\text{KL}}(P(\theta) \parallel Q(\theta)) \equiv \int_{\Theta} d^k \theta P(\theta) \ln \left(\frac{P(\theta)}{Q(\theta)} \right), \quad (3.14)$$

where we now write the distributions P and Q as functions of the continuous variable θ , and where Θ is the space of possible outcomes (parameter space).

However, yet another useful interpretation, from the field of information geometry, is that the KL divergence is a (generalised) *distance* in the space of distributions. The KL divergence is not a metric distance, since it is not symmetric in its arguments and does not satisfy the axioms of a metric. This property reflects the notion that the first argument is understood to be the ‘true’ distribution, and the second argument is, for example, a test model, an approximation or sample set that diverges from the true distribution to some degree, as implied by the discussion above. The *Jensen-Shannon* divergence remedies this. It is a divergence that is symmetric in its arguments and which satisfies the metric

axioms, and can thus be used to specify distance in the usual way, between two distributions. The Jensen-Shannon (JS) divergence is defined in terms of the KL divergence as:

$$\begin{aligned}
 D_{\text{JS}}(P(\theta) \parallel Q(\theta)) &\equiv \frac{1}{2} D_{\text{KL}}\left(P \parallel \frac{1}{2}[P + Q]\right) + (Q \leftrightarrow P) & (3.15) \\
 &= \frac{1}{2} \int_{\Theta} d^k \theta P \ln \left(\frac{P}{\frac{1}{2}[P + Q]} \right) + (Q \leftrightarrow P) \\
 &= \frac{1}{2} \int_{\Theta} d^k \theta P \ln P + Q \ln Q - \int_{\Theta} d^k \theta \frac{1}{2}(P + Q) \ln \frac{1}{2}(P + Q),
 \end{aligned}$$

where the $(Q \leftrightarrow P)$ indicates a copy of the previous terms with P and Q switching places. This divergence will be particularly useful since for the majority of our investigations we shall not be in possession of ‘true’ posteriors (those defined using the complete dataset).

3.3 Defining a new likelihood function

3.3.1 Downsampling

Downsampling is simply the process by which a dataset, \mathbf{d}_s , say, consisting of N_s samples is defined, by some means or another, as a subset of an original dataset \mathbf{d}_f which consists of $N_f > N_s$ samples; thus $\mathbf{d}_s \subseteq \mathbf{d}_f$. For example, *decimation* is a well-known and common method whereby every tenth sample is selected from some original ordered dataset, to form a new dataset one tenth of the original’s size.

3.3.1.1 Downsampling schemes

After some preliminary investigation, there appears to be rather limited advantages (and some disadvantages) to using specific downsampling (data sample selection) schemes. This is, however, limited by our simple model which has the important property of being slowly evolving. Some examples of possible schemes are:

- Uniform downsampling/‘decimation’ - taking every n^{th} sample ($n > 1$),
- Random downsampling - selecting N_s samples at random from the full dataset,
- Block downsampling - selecting blocks/chunks of given sizes from various regions of the original dataset,

- Non-uniform probability density random sampling - selecting samples at random with higher probability from certain regions of the original dataset. The probability density could be determined by considering average signal-to-noise ratio or Fisher Information contributions over short regions (for example),
- Combinations of the above methods.

One sample selection scheme may allow recovery of the fully sampled posterior in fewer samples than some other sampling scheme. However, attempts to optimise the down-sampling procedure using specific schemes should likely need to be studied for each type of model (even for each dataset considered); the efforts required for this will almost certainly outweigh any benefits and shall not be considered here.

The scheme we opted for, which gives a good representation of the signal, is simply to choose data samples at random from the available samples with uniform probability density. Uniform (periodic/regular) downsampling will lead to aliasing, causing the posterior to be a poor representation of the truth. By randomly sampling, this problem can be mostly eliminated [94], so long as the number of samples used is not catastrophically low, causing the structure of the posterior to break down.

3.3.1.2 Residuals and metrics

There are a number of ways to go about expressing the downsampled residual in terms of the fully sampled residual. This depends on the desired resulting dimensionality of the subset data space, \mathcal{D}_s . First, consider the $N_s \times N_f$ matrix

$$\mathbf{D} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \end{pmatrix}, \quad (3.16)$$

which, acting on the left of the N_f dimensional vector $\mathbf{v} = (v_1, v_2, v_3, v_4, v_5, \dots, v_{N_f})^T$ produces the N_s dimensional vector $\mathbf{v}' = \mathbf{D}\mathbf{v} = (v_2, v_3, v_6, v_7, \dots)^T$. Setting the Mahalanobis

distances equal, as in (3.13), at least approximately, is equivalent to setting

$$\mathbf{v}'^T \mathbf{\Gamma}_s^{-1} \mathbf{v}' \approx \mathbf{v}^T \mathbf{\Gamma}_f^{-1} \mathbf{v}, \quad (3.17)$$

for all \mathbf{v} , where $\mathbf{\Gamma}_s^{-1}$ and $\mathbf{\Gamma}_f^{-1}$ are the metrics on the subspace and full data space respectively. An exact solution for $\mathbf{\Gamma}_s^{-1}$ would necessarily be in terms of \mathbf{v} , and using the entire vector defeats the purpose. Our method will be to make an initial guess of a metric on the subspace $\mathbf{\Gamma}'_s$, and define the final metric $\mathbf{\Gamma}_s^{-1} = \mathbf{\Gamma}_s^{-1}(\mathbf{\Gamma}'_s)$ in terms of the initial guess in some appropriate fashion, as shall be described in the following sections. Of course $\mathbf{\Gamma}'_s$ must be an $N_s \times N_s$ matrix, so we might suppose, given that we are already downsampling \mathbf{v} using \mathbf{D} , that

$$\mathbf{\Gamma}'_s = \mathbf{D} \mathbf{\Gamma}_f^{-1} \mathbf{D}^T. \quad (3.18)$$

As an example, we will later try to find some constant m to set $\mathbf{\Gamma}_s^{-1} = m \mathbf{\Gamma}'_s$, in which case one may write

$$\mathbf{v}^T \mathbf{\Gamma}_f^{-1} \mathbf{v} \approx \mathbf{v}'^T \mathbf{\Gamma}_s^{-1} \mathbf{v}' = m \mathbf{v}^T \mathbf{D}^T \mathbf{D} \mathbf{\Gamma}_f^{-1} \mathbf{D}^T \mathbf{D} \mathbf{v}. \quad (3.19)$$

Note other equivalent ways of writing downsampling operations, in particular using a square matrix \mathbf{S} to select samples from \mathbf{v} , such that \mathbf{v}' remains N_f -dimensional, i.e, the identity matrix but with some diagonals set to zero, e.g.: $\mathbf{S} = \text{diag}(0, 0, 1, 0, \dots, 0)$. This is equivalent to the above since such a matrix may be defined as $\mathbf{S} = \mathbf{D}^T \mathbf{D}$.

Another useful way to downsample, rather than thinking of downsampling as reducing the number of dimensions by projecting the data into a smaller space using \mathbf{D} , is to define an N_f -dimensional vector, $\mathbf{k} = (0, 1, 1, 0, 0, 1, 1, 0, \dots, 0)$, a *sample selection vector*, and write $\mathbf{v}' = \mathbf{v} \odot \mathbf{k}$, where \odot is the Hadamard (element-wise) product, then using $\mathbf{\Gamma}'_s = \mathbf{\Gamma}_f$ as the inner product matrix. This retains the dimensionality and simply ‘deletes’ the data we do not intend to use; the usefulness of this approach will become apparent in the next part.

3.3.1.3 Time domain sample correlations

The acceptable downsampling rate we expect is very high, meaning that remaining samples would be greatly separated in time. The correlations effectively vanish between the remaining samples and can be ignored. One might therefore suppose one could treat the remaining samples as a small dataset effectively produced by a detector with a white

noise profile, and compute inner products quickly since the noise covariance matrix would be diagonal. Whilst it is true that these samples are effectively not correlated with each other, they will, in the full LISA dataset, be highly correlated with their immediately neighbouring samples. In the full LISA dataset analysis, these correlated neighbours are of course *not* neglected, and inform the likelihood function.

Importantly, we must be aware that we do not want to produce the posterior that some detector recording few, widely spaced, uncorrelated samples with very little noise would define (although this in itself is a rather interesting concept); rather, we would like to know and reproduce the posteriors that LISA in particular will define: we are of course trying to recreate in particular the LISA PE environment. To faithfully reproduce a LISA posterior by downsampling, we must be mindful of correlations between samples implied by the definition of the likelihood function, since they help to determine the degree to which LISA is sensitive to deviations from a pure noise dataset. In other words, since the likelihood function essentially tests the goodness of fit of the residual as detector noise, the samples in the residual are always assumed to be noise only and as such are assumed correlated as per the definition of the detector noise correlation function.

It is not obvious that neglecting correlations will cause a significant deviation from the true posterior, but our posterior convergence validity tests showed that it is generally unsafe to do so. Upon reflection, the effect is quite reasonable and may be explained with a simple example. Consider estimating the phase of a dirac-comb signal of known amplitude and frequency from a detector with a white noise profile (where the autocorrelation function is a delta function). Now, compare this to estimating the phase of the same signal measured by another detector with a noise profile such that immediately neighbouring samples are highly anti-correlated; in this case, one has more freedom to shift the phase from the truth by one sample without changing the Mahalanobis distance (and so the posterior) as much, since the residual in that case is more likely to be attributable to noise only. Hence in the downsampled case, if one selects a sample without considering how its value is affected by its neighbours, one will likely arrive at the wrong conclusions about how it constrains the model parameters.

The above example shows how correlations cause the information in the time-domain signal to which the detector is sensitive to become mixed between neighbouring samples. In order to prevent this from impinging upon the Bayesian inference after downsam-

pling, we must whiten, or in some other way decorrelate the data samples (for example, frequency representations from fast Fourier transform (FFT)s often consist of uncorrelated samples).

Thus in fact the preceding prescription for downsampling given in Section 3.3.1.2 *should only be carried out on uncorrelated data*: since frequency domain data is naturally uncorrelated, the task would be straightforward and safe (at least from the point of view of sample correlations) to downsample frequency samples. In the time domain, we must first whiten residuals before downsampling. This simply means setting

$$\mathbf{v} = \boldsymbol{\Sigma}^{(-1/2)} \mathbf{r}, \quad (3.20)$$

where \mathbf{v} is the whitened residual, so that the inner product of time domain residuals can be written

$$\langle \mathbf{r} | \mathbf{r} \rangle = \mathbf{r}^T \boldsymbol{\Sigma}^{-1} \mathbf{r} = (\boldsymbol{\Sigma}^{(-1/2)} \mathbf{r})^T \boldsymbol{\Sigma}^{(-1/2)} \mathbf{r} = \mathbf{v}^T \boldsymbol{\Gamma}^{-1} \mathbf{v}, \quad (3.21)$$

where the metric in the ‘whitened basis’ $\boldsymbol{\Gamma}^{-1} = \mathbb{1}$.

Whitening time-domain signals via the matrix operation in (3.20) is generally extremely costly. However, due to being able to downsample our data to a high degree, a fairly efficient algorithm is possible for computing only the required template samples (to within some chosen degree of accuracy regarding the number of ‘significant’ neighbouring samples, see Section 3.3.1.3.1 below) to produce the downsampled residual

$$\mathbf{v} \odot \mathbf{k} = \boldsymbol{\Sigma}^{(-1/2)} \mathbf{r} \odot \mathbf{k}, \quad (3.22)$$

where as before \mathbf{k} is a sample selection vector. Then, for example, if it can be shown that the approximation method where we find the optimal noise reduction factor m given in (3.19) is accurate, then we can write the approximate inner product as

$$\langle \mathbf{r} | \mathbf{r} \rangle \approx m (\boldsymbol{\Sigma}^{(-1/2)} \mathbf{r} \odot \mathbf{k})^T (\boldsymbol{\Sigma}^{(-1/2)} \mathbf{r} \odot \mathbf{k}). \quad (3.23)$$

3.3.1.3.1 Maximum correlated samples

The downsampled residual defined in (3.22) requires an expensive matrix multiplication to evaluate. As we will see later, we are generally allowed a definition of \mathbf{k} that vastly reduces the number of signal samples that are required to be calculated to compute $\mathbf{v} \odot \mathbf{k}$. However, for each ‘1’ in \mathbf{k} corresponding to a selected *decorrelated* sample, we are still required to know a few more samples from \mathbf{r} (the ‘significantly correlated’ samples) to

be able to perform the decorrelation and compute the element of $\Sigma^{(-1/2)} \mathbf{r}$ corresponding to the relevant ‘1’ in \mathbf{k} .

The number of elements of \mathbf{r} required to be computed for each ‘1’ in \mathbf{k} depends on the LISA ACF. The ACF always tends to zero for large values of the time lag. Generally however, an ACF tends to zero even for relatively small values of the time lag, so if we set some level of precision beyond which correlations are ignored, the number of elements of \mathbf{r} required to compute an element of $\mathbf{v} \odot \mathbf{k}$ can be reduced significantly. For our analysis, we used the level of accuracy given by:

Definition. The index of last element of the ACF with magnitude greater than 0.5% of the value of the zeroth element of the ACF is an integer known as the (approximate number of) *maximum correlated samples (MCS)*, denoted \mathcal{M} .

Any given signal has its own Nyquist rate. This determines the frequency bin/time sample size, and so the resolution of the ACF, thus the MCS will be different for each signal. If some system has an MCS of 2, then 5 elements of \mathbf{r} will be required to be known for each ‘1’ in \mathbf{k} : one corresponding to the ‘1’ in \mathbf{k} , the two preceding samples and the two succeeding samples. For example, if $k^{76} = 1$, then we are required to know r^{74}, \dots, r^{78} in order to accurately compute $v_{N_t}^{76}$. In general, $2\mathcal{M} + 1$ elements of \mathbf{r} must be computed for each ‘1’ in \mathbf{k} . This is approximate however: suppose also that $k^{77} = 1$, we then need r^{75}, \dots, r^{79} , but r^{75}, \dots, r^{78} will already have been found from the requirements of having $k^{76} = 1$. We shall write the *approximate* total number of samples evaluated as $N_{\mathcal{M}}$, which is given by

$$N_{\mathcal{M}} \approx (2\mathcal{M} + 1)N_s. \quad (3.24)$$

3.3.2 Likelihood approximation methods

As mentioned, a number of approaches have been considered for evaluating likelihood functions of LISA-like signals quickly, either exactly or by approximation. These include using ROQ waveform synthesis, using the *Fisher information* formalism, *relative binning* or simply brute force methods using high performance computing clusters [95]. Another approach to approximating the likelihood function we considered briefly, straddling the realm of posterior sampling, is by Taylor expanding the likelihood at various points. The original function can be reproduced, in principle, to an arbitrary degree of accuracy by

including the appropriate number of terms, or by using few terms at multiple points, covering small patches which can be glued together at appropriately chosen boundaries. There are good reasons, however, as will be discussed, to expect downsampling to produce accurate results very quickly, as a relatively straightforward procedure, at the same time as fitting into a well tested and understood data analysis pipelines, which it is sensible to make use of and contribute to. To our knowledge, the downsampling techniques we introduce here for PE have not previously been considered.

As mentioned in Section 3.2.2, for reproducing the likelihood function via downsampling, it is necessary to find methods to approximate the likelihood function: an exact replica of some likelihood given a detector, data and a model would necessarily require using all of the original data at each point in the parameter space, which is what we are aiming to avoid. However, we *do* have the data, detector model and signal model. Guided by some of the approximation ideas discussed above, the direction we will take is to essentially try to reshape the downsampled likelihood into the same form, correcting, or compensating, for the variation in the distribution occurring from discarding samples by matching the standardised moments (mean, variance, skewness, and so on) in part at least, to those of the original likelihood.

For simplicity, let us consider posterior defined by data consisting of a signal and a vanishing noise realisation. The first moment of the distribution is the location of the maximum likelihood value, which is automatically conserved under downsampling; the log likelihood function will always vanish at the true parameters given any selection of data samples. The second moment of the distribution is then what one might consider as the first-order contribution to the variation of the posterior from downsampling. This is known as the *parameter covariance matrix (PCM)*. Its inverse is well-known as the *Fisher information matrix (FIM)*, given by [96, 97]:

$$F_{ij} \equiv \left[\frac{\partial h^T}{\partial \theta_i} \Sigma^{-1} \frac{\partial h}{\partial \theta_j} + \frac{1}{2} \text{tr} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_i} \Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_j} \right) \right]_{\theta=\hat{\theta}}, \quad (3.25)$$

which is evaluated here at $\hat{\theta}$, the maximum likelihood estimate. For stationary noise (where Σ is constant) this reduces to

$$F_{ij} \equiv \left[\frac{\partial h^T}{\partial \theta_i} \Sigma^{-1} \frac{\partial h}{\partial \theta_j} \right]_{\theta=\hat{\theta}}. \quad (3.26)$$

It is useful to abbreviate this, dropping the evaluation point instruction and taking this

to be implied, writing

$$F_{ij} = h_{k,i} \Sigma_{kl}^{-1} h_{l,j}, \quad (3.27)$$

where summing over repeated indices is assumed, and where we defined $h_{k,i} \equiv \frac{\partial}{\partial \theta^i} h_k$.

3.3.3 Exact preservation of Fisher information

We show here that the FIM (at the true parameters) can be preserved exactly after down-sampling. This can be achieved by either modifying the signal model or the sample noise covariance matrix (of remaining, pre-whitened samples). This technique should prove particularly useful for likelihood functions that are exceedingly expensive to compute and require a very large degree of downsampling, or for waveforms that evolve sufficiently quickly. In such cases it is likely that the downsampling degrades the relative abundance of Fisher information for each parameter pair. For the slowly evolving signals we will examine later however, we predict that this will not occur and a more basic solution, with an accuracy similar to using this exact FIM preservation method, could be adequate for our Bayesian inference goals: namely, computing a single noise reduction factor to use as the ‘new detector’ model noise.

It may not always be possible to find solutions that exactly preserve the FIM; certain pathological cases can occur in which a sampling choice renders the precise recovery of Fisher information impossible. For example, if a certain signal contains more information on some parameter θ_1 overall, and the subset choice contains only samples that have more information on parameter θ_2 , then no amount of signal/noise reweighting of those samples will be able to compensate for lost Fisher information on θ_1 without overcompensating for Fisher information on θ_2 . This is unlikely, but if necessary can be avoided in the initial sample selection stage.

Given an N_f -dimensional data space and metric as the pair $(\mathcal{D}_f, \Sigma_f^{-1})$ with FIM F_{ij} , we suppose there exists an N_s -dimensional data space and metric as the pair $(\mathcal{D}_s, \Sigma_s^{-1})$ with FIM F'_{ij} , such that

$$F'_{ij} = F_{ij}. \quad (3.28)$$

For the FIM of the subset, where the noise covariance matrix is written Σ_s^{-1} , we have

$$\begin{aligned} F'_{ij} &= h_{k,i} (\Sigma_f^{-1})_{kl} h_{l,j} \\ &= \bar{h}'_{k',i} (\Gamma_s^{-1})_{k'l'} \bar{h}'_{l',j}, \end{aligned} \quad (3.29)$$

where primes indicate objects on the N_s -dimensional data subspace, and where the $\bar{h}'_{k',i}$ are the prewhitened model derivative vectors on \mathcal{D}_s . That is, given a downsampling matrix \mathbf{D} such as the example in equation (3.16), we have

$$\bar{h}'_{k',i} = D_{k'm}(\Sigma_f^{-1/2})_{mk} h_{k,i}. \quad (3.30)$$

The matrix Γ_s is unknown at this point, the form of this is what we wish to ascertain to satisfy equation (3.28). Thus let us further suppose that

$$F'_{ij} = (\omega \odot \bar{h}')_{k',i} \mathbf{Q}_{k'l'} (\omega \odot \bar{h}')_{l',j}, \quad (3.31)$$

where ω is a vector to be determined, and \mathbf{Q} is some guess of sample covariance matrix. Since the samples are prewhitened, we can choose \mathbf{Q} to be diagonal, and our Ansatz shall be $\mathbf{Q} = \mathbb{1}$. Note, however, that there may be some other \mathbf{Q} that minimises the distance between the original and downsampled posterior over the entire parameter space. With our Ansatz, we can still set the FIMs equal, but the likelihood functions may diverge faster (than the ‘optimal’ choice of \mathbf{Q} , which is not known) as one moves away from the true parameters. Proceeding with the simple, $\mathbf{Q} = \mathbb{1}$ case, we have that

$$(\omega \odot \bar{h}')_{k',i} \mathbf{Q}_{k'l'} (\omega \odot \bar{h}')_{l',j} = \sum_{k'=0}^{N_s} \omega_{k'}^2 \bar{h}'_{k',i} \bar{h}'_{k',j} = (\bar{h}')_i^T \cdot \text{diag}(\omega_1^2, \omega_2^2, \dots, \omega_{N_s}^2) \cdot (\bar{h}')_j, \quad (3.32)$$

where we can easily read off, from equation (3.29), that

$$\Gamma_s^{-1} = \text{diag}(\omega_1^2, \omega_2^2, \dots, \omega_{N_s}^2). \quad (3.33)$$

If the FIM is an $n \times n$ matrix, i.e. there are n parameters in our model, then since the FIM is symmetric, there will be $n_b = (n^2 + n)/2$ unique equations in (3.28) given by combinations of the i, j indices. We may define the weighting ω to be constrained to have the form of a sum of n_b basis functions, for example:

$$\omega_{k'}^2 = \sum_{p=0}^{n_b-1} a_p t_{k'}^p, \quad (3.34)$$

where t is the time, $t_{k'}^p$ is the k' th time value raised to the p th power, and the a_p are constants. Thus we will have n_b equations and n_b unknowns (the a_p 's), and therefore we have the set of n_b simultaneous equations

$$F_{ij} = \sum_{k'=0}^{N_s} \left(\sum_{p=0}^{n_b-1} a_p t_{k'}^p \bar{h}'_{k',i} \bar{h}'_{k',j} \right) = \sum_{p=0}^{n_b-1} a_p \sum_{k'=0}^{N_s} t_{k'}^p \bar{h}'_{k',i} \bar{h}'_{k',j}. \quad (3.35)$$

which can be written in matrix form

$$\begin{pmatrix} F_{00} \\ \vdots \\ F_{ij} \\ \vdots \\ F_{nn} \end{pmatrix} = \begin{pmatrix} (t^0 \odot \bar{h}')_{,0}^T \bar{h}'_{,0} & \dots & (t^p \odot \bar{h}')_{,0}^T \bar{h}'_{,0} & \dots & (t^n \odot \bar{h}')_{,0}^T \bar{h}'_{,0} \\ \vdots & & \vdots & & \vdots \\ (t^0 \odot \bar{h}')_{,i}^T \bar{h}'_{,j} & \dots & (t^p \odot \bar{h}')_{,i}^T \bar{h}'_{,j} & \dots & (t^n \odot \bar{h}')_{,i}^T \bar{h}'_{,j} \\ \vdots & & \vdots & & \vdots \\ (t^0 \odot \bar{h}')_{,n}^T \bar{h}'_{,n} & \dots & (t^p \odot \bar{h}')_{,n}^T \bar{h}'_{,n} & \dots & (t^n \odot \bar{h}')_{,n}^T \bar{h}'_{,n} \end{pmatrix} \begin{pmatrix} a_0 \\ \vdots \\ a_p \\ \vdots \\ a_n \end{pmatrix}, \quad (3.36)$$

for $i \leq j \leq n$, and thereby, since we know the F_{ij} and we can compute all the matrix elements, we can solve the above for the coefficients a_p , giving us the form of Γ_s^{-1} using (3.33) and (3.34). One might expect that if samples are chosen such that the Fisher information is not precisely recoverable, as per our pathological example case discussed above, then this matrix will be singular (although we shall not attempt to prove this here).

If $\mathbf{Q} \neq \mathbf{1}$ however, one finds that upon expanding, the resulting equations in the a_p are quadric equations. These are difficult to solve analytically, and the difficulty of finding solutions quickly increases as the number of parameters of the model increases. Numerical methods could be employed in this case, to find solutions for the coefficients of the chosen basis function set in (3.34).

The ‘single noise reduction factor’ approach discussed in the next section is equivalent to setting the FIMs approximately equal, with the assumption that $a_p \approx 0$ for $p > 0$ and finding the a_0 that minimises the variation in the FIMs. If a situation occurs in which this single factor is not precise enough to recover lost information, which could be expected when the Fisher information transmission rate changes significantly on timescales similar to the average duration between time samples¹, it is possible that the solution described here will be more robust (an interesting case to consider, for example, is extreme mass ratio inspirals in LISA; the rate of transmission of Fisher information is far less uniform than slow evolution BHBS, and approximate FIM recovery using a single factor may not be feasible).

3.3.4 An optimal noise reduction factor

As we have seen, setting even the first moments (the PCMs) of the distributions to be equal turned out not to be a trivial task. Seeking to simplify this for the limited class of

¹In other words, the signal cannot be said to be *slowly evolving*. We can define *Fisher information transmission rate* as the time derivative of the expectation value of the Fisher information: $\partial_t \langle F_{ij}(t) \rangle$.

signals of our interest, we proposed the conjecture that a single noise reweighting factor applied to all remaining samples would be sufficient to recover the general structure of the distribution, given enough samples remain to capture a fair representation of the information content.

To introduce the motivation of this conjecture, consider the simplest time series model, $\mathbf{h}(\theta) = \text{const}$. Suppose the detector produces time series $\mathbf{d} = \mathbf{s} + \mathbf{n}$, where \mathbf{n} is a realisation of the detector noise. Each sample contains a certain amount of information about (in the form of a constraint on) the signal's parameter. The amount of information present is thus strongly dependent on inherent detector noise (i.e. greater noise \Rightarrow less information), and the parameter constraints are manifest in the *likelihood function*.

Assuming Gaussian noise, the log likelihood function (discarding a constant that is of no consequence) of an 'original' (full) data set can be written as

$$\ell_f(\mathbf{d}_f, \theta) \equiv -\frac{1}{2} \mathbf{r}_f^T \mathbf{C}_f^{-1} \mathbf{r}_f - \frac{1}{2} \ln \left[(2\pi)^{N_f} \det(\mathbf{C}_f) \right] \quad (3.37)$$

where $\mathbf{r}_f = \mathbf{r}_f(\theta) \equiv \mathbf{d}_f - \boldsymbol{\mu}_f(\theta)$, and $\boldsymbol{\mu}_f(\theta)$ is the signal model function at θ . Here $\dim(\mathbf{d}_f) = \dim(\boldsymbol{\mu}_f) = N_f$. The noise covariance matrix will be denoted \mathbf{C}_f . For white noise, $\mathbf{C}_f = \mathbb{1} \sigma_f^2$, where σ_f is the noise standard deviation. The data vector $\mathbf{d}_f = \mathbf{s}_f + \mathbf{n}_f$, and suppose that $\mathbf{s}_f = \mathbf{1}_f \cdot \hat{\theta}$, where $\hat{\theta}$ is a constant. The model

$$\boldsymbol{\mu}_f(\theta) = \mathbf{1}_f \cdot \theta_A, \quad \theta = (\theta_A), \quad \theta_A \in \mathbb{R},$$

(i.e. θ is just a 1-dimensional parameter vector) then the log-likelihood function can be written as

$$\begin{aligned} \ell_f(\mathbf{h}_f, \theta) &= -\frac{1}{2\sigma_f^2} \left(\hat{\theta}^2 \mathbf{1}_f^T \mathbf{1}_f + \mathbf{n}_f^T \mathbf{n}_f + 2\hat{\theta} \mathbf{1}_f^T \mathbf{n}_f - 2\hat{\theta} \mathbf{n}_f^T \boldsymbol{\mu}_f - 2\hat{\theta} \mathbf{1}_f^T \boldsymbol{\mu}_f + \boldsymbol{\mu}_f^T \boldsymbol{\mu}_f \right) \\ &\approx -\frac{1}{2\sigma_f^2} \left(\underbrace{N_f \hat{\theta}^2 + N_f \langle \mathbf{n}_f^2 \rangle}_{\text{discard}} + \underbrace{2\hat{\theta} \mathbf{n}_f^T \mathbf{1}_f [1 - \theta_A]}_{\approx 0} - 2\hat{\theta} \mathbf{1}_f^T \mathbf{1}_f \cdot \theta_A + N_f \theta_A^2 \right) \\ &\approx -\frac{N_f}{2\sigma_f^2} \theta_A (\theta_A - 2\hat{\theta}), \end{aligned}$$

where we have discarded another constant and assumed that $\langle \mathbf{n} \rangle \approx 0$. Following the same procedure for a downsampled dataset, we find that

$$\ell_s(\mathbf{h}_s, \theta) \approx -\frac{N_s}{2\sigma_s^2} \theta_A (\theta_A - 2\hat{\theta}),$$

and thus, as we initially set out to do, we can set $\ell_f(\mathbf{h}_f, \theta) \approx \ell_s(\mathbf{h}_s, \theta)$ if we *choose*:

$$\sigma_s = \sigma_f \sqrt{\frac{N_s}{N_f}},$$

to be the noise standard deviation on the downsampled dataset.

We can try to extend this sort of analysis to the signals and detector noise profile we are interested in. However, it is particularly useful to have a more general result that is model independent, to widen the scope of applicability. We therefore opt to simply raise as a conjecture that in general cases of more realistic systems, the same principle will still apply in practice; that one may downsample and modify the detector noise with a single factor to compensate for lost information. We seek to verify this experimentally using a simple GW model. The (model independent) noise reduction factor is required, which is derived in the following section.

3.3.4.1 Deriving the optimal noise reduction factor

By approximating posteriors as Gaussian distributions and minimising *their* divergence (of some type), we hope to find the approximate optimal noise reduction factor that retains the shape of the likelihood function. In the high signal to noise ratio (SNR) limit, where posteriors closely approximate Gaussians, this is a good approximation (if the semi-axes of the Gaussians are well-aligned and proportionately similar). In other words, it is expected that only the overall *precision* of the likelihood will be significantly affected by downsampling. We suppose that, for our restricted class of signals of interest (LISA inspiral only signals) this assumption will generally be highly accurate, details of the results of testing this claim are given in Section 3.4.4. Since the parameter covariance matrices (PCMs) can be used to describe the aspect ratio of a k -dimensional ellipsoid representing the 1-sigma contour of the probability distribution [88], one might already guess that if the two ellipsoids are similar in orientation and aspect ratio, then the noise reduction factor required is the one that scales the ellipsoids to be equal in size. This turns out to be trivially true; for completeness we work through some details in Section 3.3.4.1.1. Then for the more robust treatment in Section 3.3.4.1.2, we will also compute the factor which accounts for the Gaussians being misaligned with respect to each other and/or having unequal aspect ratios.

3.3.4.1.1 Primitive method

As mentioned, we might consider performing an analysis under the assumption that the PCMs/posteriors have the same orientation and aspect ratio. The FIMs would thus be similar, too, differing only by a scaling factor. Then we simply note that the $m = m_{\text{det}}$ which minimises the KL divergence is the one that sets $p_s(\boldsymbol{\theta}) = p_f(\boldsymbol{\theta})$, where p_s is the posterior defined with the subset of the data and p_f is the original posterior defined with the full dataset, and thus $m_{\text{det}}\mathbf{F}'_s = \mathbf{F}_s = \mathbf{F}_f$, since in that case, both $m_{\text{det}}\mathbf{F}'_s - \mathbf{F}_f = 0$ and

$$\ln \left(m_{\text{det}}^{-k/2} \sqrt{\frac{|\mathbf{F}'_s{}^{-1}|}{|\mathbf{F}_f{}^{-1}|}} \right) = \ln 1 = 0 \quad (3.38)$$

by which (3.14) is easily shown to be

$$D_{\text{KL}}(p_f(\boldsymbol{\theta}) || p_s(\boldsymbol{\theta})) = 0. \quad (3.39)$$

This is a minimum since $D_{\text{KL}}(\cdot || \cdot) \geq 0$ always. We conjectured that only the precisions differ between posteriors, such that the factor which sets the FIMs (almost) equal automatically extends to the entire posteriors where the KL divergence between them vanishes. By (3.38), or since $m_{\text{det}}\mathbf{F}'_s = \mathbf{F}_f$, we have

$$m_{\text{det}} = \sqrt[k]{\frac{\det \mathbf{F}_f}{\det \mathbf{F}'_s}}. \quad (3.40)$$

In this case, the optimal noise reduction factor is that factor which scales the k -dimensional ellipsoids to be of the same volume; the factor m_{det} derived sets the determinants of \mathbf{F}_s and \mathbf{F}_f equal. This is certainly reasonable. Of course if the PCMs were truly aligned and similar, then rescaling such that any one particular component becomes equal will scale the others similarly too.

3.3.4.1.2 General Gaussian approximations

Let Θ be the space of parameters $\boldsymbol{\theta}$, and take k to be the dimension of the parameter space, $k = \dim(\Theta)$. Let $(\mathcal{D}_f, \mathbb{C}_f^{-1})$ be the data space and inner product operator pair for the full dataset (subscript f). If the posterior defined on the parameter space, $p_f(\boldsymbol{\theta})$, is Gaussian, it can be written

$$p_f(\boldsymbol{\theta}) = [(2\pi)^k |\mathbf{F}_f^{-1}|]^{-1/2} \exp \left[-\frac{1}{2} \boldsymbol{\Delta}^T \mathbf{F}_f \boldsymbol{\Delta} \right], \quad (3.41)$$

where \mathbf{F}_f is the FIM of the full dataset, $|\mathbf{F}_f^{-1}|$ is the determinant of its inverse, $\mathbf{\Delta} \equiv \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\theta}}$ is the location of the peak of the posterior. Suppose there also exists the space $(\mathcal{D}_s, \mathbf{M})$, with the posterior $p'_s(\boldsymbol{\theta})$, and where \mathbf{F}'_s is the FIM in $(\mathcal{D}_s, \mathbf{M})$, where \mathbf{M} is some Ansatz matrix (we will set this equal to $\mathbb{1}$ so that each sample is weighted equally). This is the (unmodified) downsampled subspace, that is: $\mathcal{D}_s \subset \mathcal{D}_f$.

Finally, let $(\mathcal{D}_s, \mathbb{C}_s^{-1} = m\mathbf{M})$ be a data space and inner product operator pair for a downsampled dataset, where m is the constant we seek. Write the Gaussian posterior similarly on this pair, $p_s(\boldsymbol{\theta})$, as

$$p_s(\boldsymbol{\theta}) = [(2\pi)^k |\mathbf{F}'_s{}^{-1}|]^{-1/2} \exp \left[-\frac{1}{2} \mathbf{\Delta}^\top \mathbf{F}'_s \mathbf{\Delta} \right], \quad (3.42)$$

where \mathbf{F}'_s is the FIM of the downsampled dataset.

Ideally, we would minimise the JS divergence, since this is a well-known, symmetric measure of distance between distributions and is impartial to either distribution, unlike the KL divergence, which implies its first argument is the true distribution and its second is a sample distribution. However, we are choosing our posterior to be a *replacement* for the truth, not some sample distribution with minimal divergence from it, so we require a symmetric divergence measure. The JS divergence is defined in equation (3.15). It contains the integral of the log of a sum and only special cases are able to be integrated analytically. In order to avoid prioritisation of any given distribution, we shall attempt to minimise another symmetric divergence known as the *Jeffreys'-divergence* instead, simply defined by

$$D_J(p_f \parallel p_s) \equiv D_{\text{KL}}(p_f \parallel p_s) + (p_s \leftrightarrow p_f). \quad (3.43)$$

First however, we warm up by minimising the KL divergence, before moving on to the Jeffreys' divergence.

We have the unknown m in $(\mathcal{D}_s, \mathbb{C}_s^{-1})$, but $(\mathcal{D}_s, \mathbf{M})$ and $(\mathcal{D}_f, \mathbb{C}_f^{-1})$ are known completely. Using the Gaussian approximations, we write down the KL divergence to find

the m that minimises the KL divergence between $p_s(\boldsymbol{\theta})$ and $p_f(\boldsymbol{\theta})$:

$$\begin{aligned}
 D_{\text{KL}}(p_f(\boldsymbol{\theta}) \parallel p_s(\boldsymbol{\theta})) &\equiv \int_{\Theta} d^k \boldsymbol{\theta} p_f(\boldsymbol{\theta}) \ln \left(\frac{p_f(\boldsymbol{\theta})}{p_s(\boldsymbol{\theta})} \right) \\
 &= \int_{\Theta} d^k \boldsymbol{\theta} p_f(\boldsymbol{\theta}) \left[\ln \left(\sqrt{\frac{|\mathbf{F}_s^{-1}|}{|\mathbf{F}_f^{-1}|}} \right) + \frac{1}{2} \boldsymbol{\Delta}^T (\mathbf{F}_s - \mathbf{F}_f) \boldsymbol{\Delta} \right] \\
 &= \ln \left(m^{-k/2} \sqrt{\frac{|\mathbf{F}'_s{}^{-1}|}{|\mathbf{F}'_f{}^{-1}|}} \right) + \frac{1}{2} \int_{\Theta} d^k \boldsymbol{\theta} p_f(\boldsymbol{\theta}) \cdot \boldsymbol{\Delta}^T (m\mathbf{F}'_s - \mathbf{F}_f) \boldsymbol{\Delta}, \quad (3.44)
 \end{aligned}$$

since the integral of the posterior (which we take to be normalised) over parameter space is equal to 1. In the last line, we also used the definition of the FIM in equation (3.27) and $\mathbb{C}_s^{-1} = m\mathbf{M}$ to find that

$$F_{s,ij} = h_{k,i}(\mathbb{C}_s^{-1})_{kl} h_{l,j} = m h_{k,i} M_{kl} h_{l,j} = m F'_{s,ij}, \quad (3.45)$$

or $\mathbf{F}_s = m\mathbf{F}'_s$, from which we obtain

$$|\mathbf{F}_s| = m^k |\mathbf{F}'_s|. \quad (3.46)$$

The KL divergence is a function of m , with the minimum being the m at which the derivative with respect to m vanishes (it is clear, given our restriction to considering Gaussians, that there will always be one extremum only, which is a minimum). The minimum is given by the $m = m_{\text{KL}}$ solving $\frac{d}{dm} D_{\text{KL}}(p_f(\boldsymbol{\theta}) \parallel p_s(\boldsymbol{\theta})) = 0$, which is

$$m_{\text{KL}} = k \left(\int_{\Theta} d^k \boldsymbol{\theta} p_f(\boldsymbol{\theta}) \cdot \boldsymbol{\Delta}^T \mathbf{F}'_s \boldsymbol{\Delta} \right)^{-1}. \quad (3.47)$$

To compute the integral, let us change coordinates first by a translation $\theta_\mu \rightarrow \theta_\mu^*$ that takes $\hat{\boldsymbol{\theta}} \rightarrow \hat{\boldsymbol{\theta}}^* = \mathbf{0}$, centering the peak of the posterior at the origin, so that $\boldsymbol{\Delta} \rightarrow \boldsymbol{\Delta}^* = \boldsymbol{\theta}^*$. Then perform a rotation $\theta_\mu^* \rightarrow \tilde{\theta}_\mu$ such that $\mathbf{F}'_s \rightarrow \tilde{\mathbf{F}}_s$, where the matrix $\tilde{\mathbf{F}}_s$ is diagonal, i.e., aligning the coordinate axes with the ellipsoid axes. We denote all objects in the final coordinates with a tilde accent. Both transformations are unitary resulting in a Jacobian determinant equal to 1, and if we further suppose that $\Theta \sim \mathbb{R}^k$ then also $\tilde{\Theta} \sim \mathbb{R}^k$, so we

can write the integral in (3.47) as

$$\begin{aligned}
 \int_{\Theta} d^k \theta p_f(\theta) \cdot \Delta^T \mathbf{F}'_s \Delta &= \int_{\tilde{\Theta}} d^k \tilde{\theta} p_f(\tilde{\theta}) \cdot \tilde{\theta}^T \tilde{\mathbf{F}}'_s \tilde{\theta} \\
 &= [(2\pi)^k |\tilde{\mathbf{F}}_f^{-1}|]^{-1/2} \int_{\tilde{\Theta}} d^k \tilde{\theta} \exp \left[-\frac{1}{2} \tilde{\theta}^T \tilde{\mathbf{F}}_f \tilde{\theta} \right] \cdot \tilde{\theta}^T \tilde{\mathbf{F}}'_s \tilde{\theta} \\
 &= [(2\pi)^k |\tilde{\mathbf{F}}_f^{-1}|]^{-1/2} \int_{\tilde{\Theta}} d^k \tilde{\theta} \exp \left[-\frac{1}{2} \sum_{l=1}^k \tilde{F}_{f,ll} \tilde{\theta}_l^2 \right] \sum_{i,j=1}^k \tilde{\theta}_i \tilde{F}'_{s,ij} \tilde{\theta}_j \\
 &= [(2\pi)^k |\tilde{\mathbf{F}}_f^{-1}|]^{-1/2} \sum_{i,j=1}^k \tilde{F}'_{s,ij} \int_{\tilde{\Theta}} d^k \tilde{\theta} \prod_{l=1}^k \exp \left[-\frac{1}{2} \tilde{F}_{f,ll} \tilde{\theta}_l^2 \right] \tilde{\theta}_i \tilde{\theta}_j.
 \end{aligned}$$

Noting the well-known integral formulas:

$$\int_{-\infty}^{\infty} dx e^{-ax^2} = \sqrt{\pi/a}, \quad \int_{-\infty}^{\infty} dx x e^{-ax^2} = 0, \quad \int_{-\infty}^{\infty} dx x^2 e^{-ax^2} = \frac{\sqrt{2\pi}}{(2a)^{3/2}}, \quad (3.48)$$

we find that the only surviving terms in the i, j sum are the $i = j$ terms. The integral becomes

$$\begin{aligned}
 \int_{\Theta} d^k \theta p_f(\theta) \cdot \Delta^T \mathbf{F}'_s \Delta &= [(2\pi)^k |\tilde{\mathbf{F}}_f^{-1}|]^{-1/2} \sum_{i=1}^k \tilde{F}'_{s,ii} \int_{-\infty}^{\infty} d\tilde{\theta}_i \tilde{\theta}_i^2 e^{-\frac{1}{2} \tilde{F}_{f,ii} \tilde{\theta}_i^2} \prod_{j \neq i} \int_{-\infty}^{\infty} d\tilde{\theta}_j e^{-\frac{1}{2} \tilde{F}_{f,jj} \tilde{\theta}_j^2} \\
 &= [(2\pi)^k |\tilde{\mathbf{F}}_f^{-1}|]^{-1/2} \sum_{i=1}^k \tilde{F}'_{s,ii} \sqrt{\frac{2\pi}{\tilde{F}_{f,ii}^3}} \prod_{j \neq i} \sqrt{\frac{2\pi}{\tilde{F}_{f,jj}}} \\
 &= \sum_{i=1}^k \frac{\tilde{F}'_{s,ii}}{\tilde{F}_{f,ii}}. \quad (3.49)
 \end{aligned}$$

If the matrix \mathbf{G} diagonalises \mathbf{F}_f , that is, if $\mathbf{F}_f \rightarrow \tilde{\mathbf{F}}_f = \mathbf{G}^{-1} \mathbf{F}_f \mathbf{G}$ for diagonal $\tilde{\mathbf{F}}_f$, then the integral may be written

$$\int_{\Theta} d^k \theta p_f(\theta) \cdot \Delta^T \mathbf{F}'_s \Delta = \sum_{i=1}^k \frac{G_{ij}^{-1} F'_{s,jl} G_{li}}{G_{ij}^{-1} F_{f,jl} G_{li}}, \quad (3.50)$$

or for general symmetric matrices \mathbf{A} and \mathbf{B} , and with \mathbf{P} diagonalising \mathbf{A} , we have

$$\int_{\Theta} d^k \theta e^{-\frac{1}{2} \Delta^T \mathbf{A} \Delta} \cdot \Delta^T \mathbf{B} \Delta = \sqrt{\frac{(2\pi)^k}{|\mathbf{A}|}} \sum_{i=1}^k \frac{P_{ij}^{-1} B_{jl} P_{li}}{P_{ij}^{-1} A_{jl} P_{li}}. \quad (3.51)$$

Therefore our factor m_{KL} , using (3.47) is

$$m_{\text{KL}} = k \left(\sum_{i=1}^k \frac{G_{ij}^{-1} F'_{s,jl} G_{li}}{G_{ij}^{-1} F_{f,jl} G_{li}} \right)^{-1}. \quad (3.52)$$

This in particular minimises $D_{\text{KL}}(p_f || p_s)$, which should be noted is not equal to $D_{\text{KL}}(p_s || p_f)$: the KL divergence not symmetric. Let us now minimise the symmetric Jeffreys' divergence. With equation (3.43), this can be written as

$$\begin{aligned} D_J(p_f || p_s) &\equiv \int_{\Theta} d^k \theta [p_f - p_s] \cdot \ln \left(\frac{p_f}{p_s} \right) \\ &= \int_{\Theta} d^k \theta [p_f - p_s] \cdot \left[\ln \left(\sqrt{\frac{|\mathbf{F}_s^{-1}|}{|\mathbf{F}_f^{-1}|}} \right) + \frac{1}{2} \Delta^T (\mathbf{F}_s - \mathbf{F}_f) \Delta \right] \\ &= \frac{1}{2} \int_{\Theta} d^k \theta [p_f - p_s] \cdot \Delta^T (m \mathbf{F}'_s - \mathbf{F}_f) \Delta, \end{aligned} \quad (3.53)$$

again since the integral of both p_f and p_s over the whole space must be equal to 1. Writing out p_f and p_s in full we have

$$\begin{aligned} D_J(p_f || p_s) &= \frac{1}{2} \int_{\Theta} d^k \theta \left[\sqrt{\frac{|\mathbf{F}_f|}{(2\pi)^k}} e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} - \sqrt{\frac{m^k |\mathbf{F}'_s|}{(2\pi)^k}} e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \right] \cdot \Delta^T (m \mathbf{F}'_s - \mathbf{F}_f) \Delta \\ &= C \int_{\Theta} d^k \theta \left[C' e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} - m^{k/2} e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \right] \cdot \Delta^T (m \mathbf{F}'_s - \mathbf{F}_f) \Delta, \end{aligned}$$

where $C = \frac{1}{2} \sqrt{\frac{|\mathbf{F}'_s|}{(2\pi)^k}}$ and $C' = \sqrt{\frac{|\mathbf{F}_f|}{|\mathbf{F}'_s|}}$. Now differentiate with respect to m :

$$\begin{aligned} \frac{d}{dm} D_J(p_f || p_s) &= C \int_{\Theta} d^k \theta \left[C' e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} - m^{k/2} e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \right] \cdot \Delta^T \mathbf{F}'_s \Delta \\ &\quad - \left[\frac{k}{2} m^{k/2-1} e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} - \frac{1}{2} m^{k/2} \Delta^T \mathbf{F}'_s \Delta e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \right] \cdot \Delta^T (m \mathbf{F}'_s - \mathbf{F}_f) \Delta \\ &= C \int_{\Theta} d^k \theta C' e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} \cdot \Delta^T \mathbf{F}'_s \Delta \\ &\quad - m^{k/2} e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \left\{ \left[1 + \frac{k}{2} - \frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta \right] \cdot \Delta^T \mathbf{F}'_s \Delta \right. \\ &\quad \left. - \left[\frac{k}{2m} - \frac{1}{2} \Delta^T \mathbf{F}'_s \Delta \right] \cdot \Delta^T \mathbf{F}_f \Delta \right\} \\ &= CC' \int_{\Theta} d^k \theta e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} \cdot \Delta^T \mathbf{F}'_s \Delta \\ &\quad - C m^{k/2} \int_{\Theta} d^k \theta e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \cdot \Delta^T \left(\left[1 + \frac{k}{2} \right] \mathbf{F}'_s - \frac{k}{2m} \mathbf{F}_f \right) \Delta \\ &\quad + \frac{1}{2} C m^{k/2} \int_{\Theta} d^k \theta e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \{ m (\Delta^T \mathbf{F}'_s \Delta)^2 - \Delta^T \mathbf{F}'_s \Delta \cdot \Delta^T \mathbf{F}_f \Delta \}. \end{aligned} \quad (3.54)$$

The first two integrals of the last line are computed easily using (3.51) by making the relevant matrix substitutions;

$$\int_{\Theta} d^k \theta e^{-\frac{1}{2} \Delta^T \mathbf{F}_f \Delta} \cdot \Delta^T \mathbf{F}'_s \Delta = \sqrt{\frac{(2\pi)^k}{|\mathbf{F}_f|}} \sum_{i=1}^k \frac{D_{f,ij}^{-1} F'_{s,jl} D_{f,li}}{D_{f,ij}^{-1} F_{f,jl} D_{f,li}} \quad (3.55)$$

for diagonal $\mathbf{D}_f^{-1} \mathbf{F}_f \mathbf{D}_f$, and

$$\begin{aligned}
 & \int_{\Theta} d^k \theta e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \cdot \Delta^T \left(\left[1 + \frac{k}{2} \right] \mathbf{F}'_s - \frac{k}{2m} \mathbf{F}_f \right) \Delta \\
 &= \sqrt{\frac{(2\pi)^k}{|\mathbf{mF}'_s|}} \sum_{i=1}^k \frac{D_{s,ij}^{-1} \left(\left[1 + \frac{k}{2} \right] F'_s - \frac{k}{2m} F_f \right)_{jl} D_{s,li}}{D_{s,ij}^{-1} m F'_{s,jl} D_{s,li}} \\
 &= km^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\mathbf{F}'_s|}} \left(\left[1 + \frac{k}{2} \right] m - \frac{1}{2} \sum_{i=1}^k \frac{D_{s,ij}^{-1} F_{f,jl} D_{s,li}}{D_{s,ij}^{-1} F'_{s,jl} D_{s,li}} \right) \quad (3.56)
 \end{aligned}$$

for diagonal $\mathbf{D}_s^{-1} \mathbf{F}'_s \mathbf{D}_s$.

For the final integral in (3.54) we again diagonalise the matrix in the exponent to make the integration simple. Similarly to how we transformed earlier, with a translation and a rotation, let us take $\theta_\mu \rightarrow \check{\theta}_\mu$ so that $\mathbf{F}'_s \rightarrow \check{\mathbf{F}}'_s$ where $\check{\mathbf{F}}'_s$ is diagonal, and such that $\check{\Delta} = \check{\theta}$. Then we can rewrite the integral as

$$\begin{aligned}
 I &= \int_{\Theta} d^k \theta e^{-\frac{1}{2} m \Delta^T \mathbf{F}'_s \Delta} \{ m (\Delta^T \mathbf{F}'_s \Delta)^2 - \Delta^T \mathbf{F}'_s \Delta \cdot \Delta^T \mathbf{F}_f \Delta \} \\
 &= \int_{\check{\Theta}} d^k \check{\theta} e^{-\frac{1}{2} m \check{\theta}^T \check{\mathbf{F}}'_s \check{\theta}} \{ m (\check{\theta}^T \check{\mathbf{F}}'_s \check{\theta})^2 - \check{\theta}^T \check{\mathbf{F}}'_s \check{\theta} \cdot \check{\theta}^T \check{\mathbf{F}}_f \check{\theta} \} \\
 &= \int_{\check{\Theta}} d^k \check{\theta} e^{-\frac{1}{2} m \check{\theta}^T \check{\mathbf{F}}'_s \check{\theta}} \left\{ m \sum_{i=1}^k \check{F}'_{s,ii} \check{\theta}_i^2 \sum_{j=1}^k \check{F}'_{s,jj} \check{\theta}_j^2 - \sum_{i=1}^k \check{F}'_{s,ii} \check{\theta}_i^2 \sum_{j=1}^k \check{F}_{f,jj} \check{\theta}_j^2 \right\} \\
 &= \int_{\check{\Theta}} d^k \check{\theta} e^{-\frac{1}{2} m \check{\theta}^T \check{\mathbf{F}}'_s \check{\theta}} \sum_{i=1}^k \check{F}'_{s,ii} \check{\theta}_i^2 \sum_{j=1}^k (m \check{F}'_{s,jj} - \check{F}_{f,jj}) \check{\theta}_j^2 \\
 &= \sum_{i=1}^k \check{F}'_{s,ii} \sum_{j=1}^k (m \check{F}'_{s,jj} - \check{F}_{f,jj}) \int_{\check{\Theta}} d^k \check{\theta} \prod_{l=1}^k e^{-\frac{1}{2} m \check{F}'_{s,ll} \check{\theta}_l^2} \check{\theta}_i^2 \check{\theta}_j^2 \\
 &= \sum_{i=1}^k (m (\check{F}'_{s,ii})^2 - \check{F}'_{s,ii} \check{F}_{f,ii}) \int_{-\infty}^{\infty} d\check{\theta}_i e^{-\frac{1}{2} m \check{F}'_{s,ii} \check{\theta}_i^2} \cdot \check{\theta}_i^4 \prod_{j \neq i} \int_{-\infty}^{\infty} d\check{\theta}_j e^{-\frac{1}{2} m \check{F}'_{s,jj} \check{\theta}_j^2} \\
 &\quad + \sum_{i=1}^k \check{F}'_{s,ii} \sum_{j \neq i} (m \check{F}'_{s,jj} - \check{F}_{f,jj}) \prod_{l \in \{i,j\}} \int_{-\infty}^{\infty} d\check{\theta}_l e^{-\frac{1}{2} m \check{F}'_{s,ll} \check{\theta}_l^2} \cdot \check{\theta}_l^2 \prod_{h \notin \{i,j\}} \int_{-\infty}^{\infty} d\check{\theta}_h e^{-\frac{1}{2} m \check{F}'_{s,hh} \check{\theta}_h^2},
 \end{aligned}$$

where in the third line terms relating to the off-diagonal elements of $\check{\mathbf{F}}_f$ drop out since integrands of those terms are odd functions. In the last step, we split the sum into an $i = j$ part (first line) and a $i \neq j$ part (second line) for the $\check{\theta}_i^2 \check{\theta}_j^2$ factor in the integrand. Along with (3.48) we require the following well-known identity

$$\int_{-\infty}^{\infty} dx x^4 e^{-ax^2} = \frac{3}{4} \sqrt{\frac{\pi}{a^5}}. \quad (3.57)$$

Then we have that

$$\begin{aligned}
 I &= \sum_{i=1}^k (m(\check{F}'_{s,ii})^2 - \check{F}'_{s,ii}\check{F}'_{f,ii}) \frac{3}{4} \sqrt{\frac{\pi}{(\frac{1}{2}m\check{F}'_{s,ii})^5}} \prod_{j \neq i} \sqrt{\frac{\pi}{\frac{1}{2}m\check{F}'_{s,jj}}} \\
 &\quad + \sum_{i=1}^k \check{F}'_{s,ii} \sum_{j \neq i} (m\check{F}'_{s,jj} - \check{F}'_{f,jj}) \prod_{l \in \{i,j\}} \sqrt{\frac{2\pi}{(m\check{F}'_{s,ll})^3}} \prod_{h \notin \{i,j\}} \sqrt{\frac{\pi}{\frac{1}{2}m\check{F}'_{s,hh}}} \\
 &= 3m^{-\frac{k+4}{2}} \sum_{i=1}^k (m(\check{F}'_{s,ii})^2 - \check{F}'_{s,ii}\check{F}'_{f,ii})(\check{F}'_{s,ii})^{-2} \prod_{j=1}^k \sqrt{\frac{2\pi}{\check{F}'_{s,jj}}} \\
 &\quad + m^{-\frac{k+4}{2}} \sum_{i=1}^k \check{F}'_{s,ii} \sum_{j \neq i} (m\check{F}'_{s,jj} - \check{F}'_{f,jj}) \prod_{l \in \{i,j\}} (\check{F}'_{s,ll})^{-1} \prod_{h=1}^k \sqrt{\frac{2\pi}{\check{F}'_{s,hh}}} \\
 &= 3m^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \sum_{i=1}^k \left(m - \frac{\check{F}'_{f,ii}}{\check{F}'_{s,ii}} \right) + m^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \sum_{i=1}^k \sum_{j \neq i} \left(m - \frac{\check{F}'_{f,jj}}{\check{F}'_{s,jj}} \right) \\
 &= m^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \left[2 \sum_{i=1}^k \left(m - \frac{\check{F}'_{f,ii}}{\check{F}'_{s,ii}} \right) + k \sum_{j=1}^k \left(m - \frac{\check{F}'_{f,jj}}{\check{F}'_{s,jj}} \right) \right] \\
 &= m^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \left[(2+k) \left(km - \sum_{i=1}^k \frac{\check{F}'_{f,ii}}{\check{F}'_{s,ii}} \right) \right]. \tag{3.58}
 \end{aligned}$$

Substituting (3.55), (3.56), and (3.58) into (3.54), we have

$$\begin{aligned}
 \frac{d}{dm} D_J(p_f || p_s) &= CC' \sqrt{\frac{(2\pi)^k}{|\mathbf{F}_f|}} \sum_{i=1}^k \frac{D_{f,ij}^{-1} F'_{s,jl} D_{f,li}}{D_{f,ij}^{-1} F_{f,jl} D_{f,li}} \\
 &\quad - Cm^{k/2} km^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \left(\left[1 + \frac{k}{2} \right] m - \frac{1}{2} \sum_{i=1}^k \frac{D_{s,ij}^{-1} F_{f,jl} D_{s,li}}{D_{s,ij}^{-1} F'_{s,jl} D_{s,li}} \right) \\
 &\quad + \frac{1}{2} Cm^{k/2} m^{-\frac{k+4}{2}} \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \left[(2+k) \left(km - \sum_{i=1}^k \frac{\check{F}'_{f,ii}}{\check{F}'_{s,ii}} \right) \right] \\
 &= C \sqrt{\frac{(2\pi)^k}{|\check{\mathbf{F}}'_s|}} \left\{ \sum_{i=1}^k \frac{D_{f,ij}^{-1} F'_{s,jl} D_{f,li}}{D_{f,ij}^{-1} F_{f,jl} D_{f,li}} - km^{-2} \left(\left[1 + \frac{k}{2} \right] m - \frac{1}{2} \sum_{i=1}^k \frac{D_{s,ij}^{-1} F_{f,jl} D_{s,li}}{D_{s,ij}^{-1} F'_{s,jl} D_{s,li}} \right) \right. \\
 &\quad \left. + m^{-2} \left[\left(1 + \frac{k}{2} \right) \left(km - \sum_{i=1}^k \frac{\check{F}'_{f,ii}}{\check{F}'_{s,ii}} \right) \right] \right\} \\
 &= \frac{1}{2} \left[\sum_{i=1}^k \frac{D_{f,ij}^{-1} F'_{s,jl} D_{f,li}}{D_{f,ij}^{-1} F_{f,jl} D_{f,li}} - m^{-2} \sum_{i=1}^k \frac{D_{s,ij}^{-1} F_{f,jl} D_{s,li}}{D_{s,ij}^{-1} F'_{s,jl} D_{s,li}} \right].
 \end{aligned}$$

We want the $m = m_J$ such that this derivative vanishes, so finally we see that the factor

minimising the Jeffreys' divergence is given by:

$$m_J = \sqrt{\frac{\sum_{i=1}^k \frac{D_{s,ij}^{-1} F_{f,jl} D_{s,li}}{D_{s,ij}^{-1} F'_{s,jl} D_{s,li}}}{\sum_{i=1}^k \frac{D_{f,ij}^{-1} F'_{s,jl} D_{f,li}}{D_{f,ij}^{-1} F_{f,jl} D_{f,li}}}}, \quad (3.59)$$

where recall that \mathbf{D}_s diagonalises \mathbf{F}'_s and \mathbf{D}_f diagonalises \mathbf{F}_f . This is a sensible result; it is similar to m_{KL} in form but clearly balanced between the distributions.

3.3.5 Summary

The downsampling procedure we have devised, though not necessarily programmed precisely as described here for reasons of computational efficiency and/or algorithmic practicality, will consist of the following steps:

1. Compute the FIM of the full LISA posterior and $R^{(-1/2)} \sim \Sigma^{(-1/2)}$ from the PSD
2. Generate an N_f -dimensional vector, $\mathbf{k} = (0, 0, 0, 1, 0, 1, 0, 0, 1, 0, 0, \dots)$, containing N_s 1's and $(N_f - N_s)$ 0's, randomly positioned (or equivalently, a matrix \mathbf{D} as in (3.16))
3. Use \mathbf{k} to define the (whitened) N_f -dimensional, downsampled residual; $\hat{\mathbf{r}}_{N_f} = \Sigma^{(-1/2)} \mathbf{r} \odot \mathbf{k}$. This seems somewhat more natural to implement for efficient computation, given the whitening required. However, an N_s -dimensional version, $\hat{\mathbf{r}}_{N_s} = \mathbf{D} \Sigma^{(-1/2)} \mathbf{r}$ is perhaps conceptually clearer for visualising the downsampling
4. Define a new inner product on residuals as we did in the right-hand-side of (3.23). Using the m calculated in (3.40) or (3.59), this should be approximately equal to the left-hand-side of (3.23), for all θ : $\langle \mathbf{r}(\theta) | \mathbf{r}(\theta) \rangle \approx m \hat{\mathbf{r}}(\theta)^T \hat{\mathbf{r}}(\theta)$.

In practice, the two factors in (3.40) and (3.59) are generally numerically quite close to each other for the example systems explored in Table 3.1, within roughly around 3%-5%. The m_J computed in (3.59) accounts for general concentric Gaussians, and is thus more stable against downsampling noise, however, this approach was discovered *after* having produced results for the downsampling analysis using the m_{KL} in (3.40). The next section tests downsampling by setting determinants of FIMs equal. This will provide a ceiling for JS divergences computed in the following section to test the procedure; the m_J factor will perform as well or better than this (since we will define good performance as minimal Jeffreys' divergence), and m_J will be used for the results in Chapters 6, 7 & 8.

3.4 Testing the downsampling procedure

Intuitively, the fewer samples required to define a particular likelihood function, the faster likelihood evaluation and PE convergence will be; recall faster PE convergence is the sole reason for attempting downsampling. One can expect that if $N_s \approx N_f$, for example, dropping just a few samples from a dataset \mathbf{d}_f , with, say, $\dim(\mathbf{d}_f) \approx 10^7$, the downsampled posterior would be hardly distinguishable from the LISA posterior; the remaining data will still just as well constrain the likely parameters of the signal (if the noise on the remaining samples is adjusted appropriately) but the speed improvement would be slight, due to evaluating the signal model at slightly fewer points. As one drastically decreases the number of samples to only a few however, one can expect the remaining data to fit other combinations of parameters, leading to the appearance of new posterior modes; the posterior becomes more highly structured. Consider decimation by a degree such that the signal is sampled below its Nyquist frequency; aliasing should occur at this point. As mentioned, random sampling (as opposed to decimation) does very well at ameliorating this problem. However, extremely low numbers of samples from random downsampling can still cause a posterior to ‘break down’ (that is, to begin to differ qualitatively from the truth by an unacceptable degree). What we would then like to determine is: how can we quantify this breakdown of the qualitative structure of the posterior, and what is the smallest number of samples required to faithfully reproduce the true posterior, before breakdown occurs?

Changing dependencies of the posterior function (e.g. by downsampling) must ultimately change the posterior to some degree. It is not obvious how to establish criteria for when a dissimilarity may be declared ‘acceptable’; this could depend on the applications of the analysis and what a researcher hopes to gain by studying a posterior. Fortunately, somewhat paradoxically, it turns out that we can be spared from trying to formulate some sort of hard boundary between ‘acceptable’ and ‘unacceptable’ (where the posterior breaks down) by the inherent, well-known *difficulty* of evaluating posterior functions. As discussed in Section 2.4, the difficulty means we must use certain algorithms (samplers) to numerically *approximate* posterior functions. Because the results are approximations that are acquired from processes employing some level of randomness, there is always a degree of variation, or noise, present between them, defining a noise floor: we could

not hope to achieve (measure) any less variation between the downsampled posteriors than this, since the downsampled posteriors must also always be ‘discovered’ with these approximation techniques. These similarity benchmarks are more concretely defined in Section 3.4.3. For now, we note that the aim is to: *find the lowest number of random samples used to define posteriors, such that variation between them remains indistinguishable from the variation resulting from posterior approximation algorithms.*

In Section 3.4.1 we discuss some further background on posterior estimators and noise, other posterior noise sources, and limitations of the experimental approach we take (as mentioned previously, we answer the question of finding the approximate optimal downsampling rate by experiment). In Section 3.4.2 we introduce useful quantifiers with which similarities of posterior functions can be measured. We have mentioned already that Section 3.4.3 poses the convergence criteria, and finally in Section 3.4.4 we present the model and systems used for testing, and the results of the analysis.

3.4.1 Posterior distribution noise sources

We provide a brief description of noise sources present in posterior distribution estimates derived by applying samplers to noisy data. Note that these are just conceptual introductions and the mathematical representations will be presented in Section 3.4.3.

3.4.1.1 Sampler noise

It generally takes an inordinate amount of effort to evaluate the posterior over a lattice to any useful degree of accuracy in a parameter space greater than two dimensions, and posterior estimators called *samplers* are used to explore the parameter space, building an approximate representation of the function until convergence criteria are met instead (see Section 2.4). The posterior estimates produced by samplers can generally vary to a small degree from sampler settings, but rather significantly between different samplers, and by the random processes that samplers use to explore the parameter space. We tested the MCMC sampler `emcee` [73] and nested samplers `CPnest` [77] and `Nessai` [76] to ensure these issues are not limited to the main sampler we chose to perform the main analysis with: `Nessai`. Using the same sampler with the same settings to produce multiple posterior estimates from the same data, the only source of variation present will be that which results from the sampler’s inherent randomness. This shall be referred to

as *sampler noise* and is unavoidable. Because of this, the sampler noise will constitute the ‘noise floor’.

3.4.1.2 Downsampling data sample selection noise

Each data sample will generally contain different information in the form of constraints on the model parameters. If N_s samples are selected from an original set of N_f samples, where $N_s \gg 1$ (and the signal model is slowly evolving) then the selection will be likely to be a good representative sample of the information present in the full dataset. However, there will of course be some variations on the parameter constraints depending on the precise selection of the N_s samples. Thus, different manifestations of the likelihood function (by which is meant ‘different definitions of’, from the different random sample choices/different realisations of \mathbf{k} or \mathbf{D}) will exhibit some degree of variation, which we shall refer to as *downsampling noise*.

As N_s becomes smaller, the selection of samples becomes less likely to contain a good representation of the information present in the original dataset. We might then expect there to be, roughly, some number of samples below which the downsampled posterior fails to accurately represent the true posterior.

3.4.1.3 Physical noise

Real data contains detector noise. Different realisations of detector noise tend to shift the position of the maximum likelihood estimate (MLE), resulting in different posterior distributions. If we try to compare posteriors generated using different noise realisations, we will see a significant difference due to this MLE offset. This would interfere with our judgement of the amount of downsampling noise and efficacy of downsampling, which we seek to quantify. By using a zero detector noise realisation, we effectively find the average posterior of many posteriors produced using different random detector noise realisations [95, 98], that is, it is in some sense the *expected posterior distribution* (the PSD/noise matrix itself remains unchanged). Hence we shall use a zero detector noise realisation, to give a fair basis for analysing the results on downsampling noise.

3.4.2 Determining closeness of distributions

The JS Divergence provides a measure of distance between distributions particularly well-suited for our problem (see discussion in Section 3.2.4). As mentioned however, the posterior distributions themselves are computationally expensive. For high dimensional parameter spaces, the JS divergence is vastly more expensive still; it first requires the posteriors to be known, then needs to perform further mathematical operations as per its definition. We must resort to approximation techniques.

3.4.2.1 Combined distance between marginal distributions

Consider using the JS divergence of *marginalised distributions*. Marginalisation is certainly not ideal; as it is a projection, information is lost. However it can still be useful as a measure to show when variations between the *profiles* of posteriors become very small, and thus point to convergence. A great benefit is that JS divergences between 1-dimensional distributions are inexpensive to compute. One has an infinite choice of ways in which to project a distribution with ≥ 2 dimensions into 1-dimension (by parameter mixing). We are generally interested in the unmixed model parameter constraints however, so it is sensible to choose the projections that are the marginalisations onto the usual, unmixed model parameters.

Since the JS divergence can be used to define a metric and thus a distance, we can combine distances of these marginalisations to form a new distance. We can then define a metric between distributions p, q as:

$$d_{\text{JS}}(p, q) \equiv \sqrt{D_{\text{JS}}(p || q)}. \quad (3.60)$$

Let us now define a generalised Pythagorean distance function

$$\bar{d}_{\text{JS}}(p, q) \equiv \sqrt{\frac{1}{k} \sum_{i=1}^k d_{\text{JS}}^2(O^i(p), O^i(q))} \quad (3.61)$$

where the operator

$$O^i(p) \equiv \int_{\mathbb{R}^{k-1}} p(\theta) d\theta^1 \dots d\theta^{i-1} d\theta^{i+1} \dots d\theta^{k-1} \quad (3.62)$$

projects the distribution p onto the i^{th} parameter axis. The function \bar{d}_{JS} satisfies the requirements of a metric and thus is a metric. If we square the new metric to define a new

divergence \overline{D}_{JS} , we will have that

$$\overline{D}_{\text{JS}}(p \parallel q) \equiv \frac{1}{k} \sum_{i=1}^k D_{\text{JS}}(O^i(p) \parallel O^i(q)) . \quad (3.63)$$

The range of D_{JS} is $[0, 1]$, the factor of $1/k$ introduced above thus ensures that the range of \overline{D}_{JS} is $[0, 1]$ also. We shall refer to this divergence, which will be instrumental for indicating the convergence of downsampled posteriors, as the combined marginalised Jensen-Shannon (CMJS) divergence.

3.4.3 Convergence criteria

We are now almost ready to propose criteria with which to determine the minimum number of samples required for the downsampled posterior to be an accurate representation of the true posterior. Recall that we have two types of posterior noise to consider: sampler noise and downsampling noise. Note that henceforth we will treat ‘sampler noise’ as equivalent to and synonymous with the ‘noise floor’. Recall our aim is to find when the posterior, as a function of the number of samples, converges to a *stable* solution to an accuracy determined by the noise floor. It has not yet been determined how to quantify noise between distributions, however. We shall do this now.

3.4.3.1 Noisy posterior distributions

With the JS divergence as a measure of distance between distributions, we can begin to characterise some basic statistical properties of the variation in the closeness of distributions, arising from some particular noisy process (i.e. the sampler’s internal operations, or the random downsampling). Call the mean of a set of divergences, X , $\mu[X]$, and the standard deviation $\sigma[X]$.

We now define two one-parameter (parameterised by N_s) families of sets of divergences, the statistical properties of which we should expect to have certain behaviours when the posteriors, as functions of N_s , stabilise:

1. For each number N_s of samples used to define a posterior, if we generate K posterior estimations, then we can form a population of $\binom{K}{2}$ pairs (‘ K choose 2’) of estimates, and thus a set of $\binom{K}{2}$ JS divergences. We can then define the mean and

standard deviation of the noise floor distance sets X_{fl} as functions of N_s , $\mu[X_{\text{fl}}(N_s)]$ and $\sigma[X_{\text{fl}}(N_s)]$, respectively, where

$$X_{\text{fl}}(N_s) = \bigcup_{i < j \leq K} \bar{D}_{\text{JS}} \left(p_{\text{fl}}^i(N_s) \parallel p_{\text{fl}}^j(N_s) \right), \quad (3.64)$$

where p_{fl}^i is the i^{th} of the K estimate posteriors, all derived from the exact same data set (i.e. using only one instance of \mathbf{k} or \mathbf{D}).

If the posterior (estimates) stabilise at some N_s , then the means and standard deviations of the downsampling noise sets, defined similarly as

$$X_{\text{ds}}(N_s) = \bigcup_{i < j \leq K} \bar{D}_{\text{JS}} \left(p_{\text{ds}}^i(N_s) \parallel p_{\text{ds}}^j(N_s) \right), \quad (3.65)$$

where now each posterior estimate, p_{ds}^i , is derived from a *different* random selection of data points (i.e. K unique instances of \mathbf{k} or \mathbf{D}), should behave in the same way as those of the noise floor.

2. We can also find the average distance of a set of posterior distributions from a test distribution. Whereas the first family of sets can be said, in a sense, to be ‘bootstrapped’, the use of a test distribution provides a sort of anchor from which to compare divergences. Let us then define the further sets

$$Y_{\text{fl}}(N_s) = \bigcup_{i \leq K} \bar{D}_{\text{JS}} \left(p_{\text{fl}}^i(N_s) \parallel q \right), \quad (3.66)$$

and

$$Y_{\text{ds}}(N_s) = \bigcup_{i \leq K} \bar{D}_{\text{JS}} \left(p_{\text{ds}}^i(N_s) \parallel q \right), \quad (3.67)$$

where p_{fl}^i and p_{ds}^i are the same posteriors as those used in (3.64) and (3.65) respectively, and q is some test distribution.

If the posterior (estimates) stabilise at some N_s , then the means and standard deviations of the downsampling noise, defined similarly, should behave in the same way as those of the noise floor.

What we hope to show is that, as N_s is increased, downsampling noise becomes indistinguishable from sampler noise at a certain N_s and is thus independent of N_s (i.e. when choosing *different* samples is indistinguishable from choosing *more* samples). This does

not necessarily tell us that the posterior has truly converged of course, only that it is varying with data sample selection, on average, by the same amount for each selection. To address this, we can demand as a further criterion that the downsampled FIM has converged on, and is approximately equal to the true FIM. Since the FIMs are independent of sampler noise (and can be compared comprehensively, without losing information through marginalisation) this reassures us that the variation is indeed small; it is not likely for the general structure of the posterior to have not yet stabilised but for the FIM to exhibit little variation, so the posterior noise can in the most part be attributed to sampler noise. For very low sampler noise, we shall consider it very highly unlikely that both the FIMs have converged and the downsampling noise has stabilised (the means are approximately constant) for the downsampled posterior to have not yet accurately converged to the true posterior. In the case of high sampler noise however, imposing the condition that the standard deviation of the downsampling noise must be of the order of that of the sampler noise ensures that the variations are of similar magnitude and thus not likely to be from a changing structure of the posterior.

3.4.3.2 Summary

We have introduced conditions above, which, together, will constitute our criteria for accepting that a posterior has accurately converged to the truth at a certain N_s . Thus, the smallest N_s for which all of the following conditions are met shall be considered the optimal downsampling rate:

1. $\mu[X_{ds}(N_s)] \approx \mu[X_{fl}(N_s)] \approx \text{const.},$
2. $\mu[Y_{ds}(N_s)] \approx \mu[Y_{fl}(N_s)] \approx \text{const.},$
3. $\sigma[X_{ds}(N_s)] \approx \sigma[X_{fl}(N_s)] \approx \text{const.},$
4. $\sigma[Y_{ds}(N_s)] \approx \sigma[Y_{fl}(N_s)] \approx \text{const.},$
5. $\mathbf{F}_s(N_s) \approx \mathbf{F}_f.$

3.4.4 The distances evaluated

For experimentally probing the optimum downsampling rate, we defined a set of 12 different BHB systems to use as a testbed for the downsampling procedure. The GW model used is a simplified version of the presently known best accuracy GW waveforms; for approximants of waveforms expressed in *post-Newtonian* (PN) expansions, of which the

#	N_{LISA}	N_{Nyq}	SNR	Dur.	t_c^{end}	m_1	m_2	d	ι	ϕ	Free parameters
	-	-	-	y	y	M_\odot	M_\odot	Mpc	rad	rad	
1	1.58×10^8	3.04×10^6	5.89	0.5	0.1	148	27	230	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d
2	3.15×10^8	3.66×10^6	8.27	1	0.2	246	37	594	0	0.4	$\mathcal{M}_c, q, t_c, d, \iota$
3	3.15×10^9	2.08×10^7	9.08	10	0.5	195	83.5	1182	$\pi/3$	$\pi/2$	q, t_c, d, ι, ϕ
4	1.58×10^9	7.69×10^6	8.29	5	0.4	300	183.5	1835	$\pi/2.1$	$\pi/3$	$\mathcal{M}_c, q, t_c, d, \iota, \phi$
5	1.89×10^9	5.65×10^6	12.55	6	0.4	700	380	4688	$\pi/5$	$\pi/7$	$\mathcal{M}_c, q, t_c, d, \iota, \phi$
6	9.47×10^8	2.37×10^7	8.22	3	0.1	46	33.5	203	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d
7	1.89×10^9	3.65×10^7	35.88	6	0.1	148	27	115	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d
8	3.15×10^8	1.95×10^6	24.16	1	0.2	246	237	814	0	0.4	$\mathcal{M}_c, q, t_c, d, \iota$
9	1.58×10^9	3.94×10^7	10.30	5	0.1	46	33.5	203	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d
10	3.15×10^8	3.66×10^6	16.54	1	0.2	246	37	297	0	0.4	$\mathcal{M}_c, q, t_c, d, \iota$
11	9.47×10^8	2.37×10^7	16.44	3	0.1	46	33.5	101	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d
12	1.58×10^9	3.94×10^7	20.61	5	0.1	46	33.5	101	$\pi/4$	0.1	\mathcal{M}_c, q, t_c, d

Table 3.1: A list of the black-hole binary systems used to test the downsampling procedure. From left to right, the columns describe: BHB system number/label, number of LISA samples supposing a 10 Hz sampling frequency, the signal’s number of Nyquist samples (the minimum uniform sample rate number of samples required to recover the signal without aliasing), the SNR, signal duration (years), time between end of signal and BHB coalescence (years), primary mass (M_\odot), secondary mass (M_\odot), distance (Mpc, no-redshift), inclination angle (rad), phase (rad), free model parameters.

highest order known is 3.5PN, we use a GW phase of 3.5PN and up to the 0.5PN order term in the expansion of the waveform amplitude. We used systems of between four and six free parameters; details of the BHB systems (which can all be said to be slowly evolving by inspection of their Fisher information matrices as functions of time) are given in Table 3.1. Recall that the downsampling procedure has been developed for, and should be applicable to general slowly evolving signal models. Besides the requirement of slow evolution, the procedure is model independent and may be readily applied to more sophisticated models with more parameters.

We produced the two sets of posteriors, p_{fl}^i and p_{ds}^i , used to define the distance sets in Eqns (3.64)-(3.67), such that they each contained $K = 10$ posterior estimates. These were generated with Bilby [79] & Nessai [76], where our downsampled likelihood function

defined the posteriors, using $N_s(n) = \lfloor 2^{n/2} + \frac{1}{2} \rfloor$ samples, where $8 \leq n \leq 26$ and where $\lfloor x + \frac{1}{2} \rfloor$ rounds x to its nearest integer (the number of waveform data points required to compute the inner product is actually greater than this; N_s is the number of *decorrelated* samples that define the likelihood, see Section 3.3.1.3 for details). That is, sample selection vectors \mathbf{k} were defined containing N_s 1's (and $N_f - N_s$ 0's), where N_s is taken from

$$\mathfrak{N} = \{16, 23, 32, 45, 64, \dots, 8192\}. \quad (3.68)$$

Altogether, producing a posterior for each system specified in Table 3.1, for each $N_s \in \mathfrak{N}$, for each $i \in \{1, \dots, K\}$, and for each of p_{fl}^i and p_{ds}^i , resulted in 4,560 unique posterior estimations. For each BHB system, the q test distribution used in Eqns (3.66), (3.67) is taken to be $p_{\text{fl}}^1(N_s = 8192)$, $p_{\text{ds}}^1(N_s = 8192)$, respectively. The means and standard deviations of the sets defined in (3.64)–(3.67) were plotted and studied; a selection of these is shown in Figure 3.2. Considerable variation in the statistics remained, likely due to a low number of pairs of distributions. Perhaps a better approach would have been to define fewer BHB testbed systems and instead choosing, say, $K = 25$. Then the population size of distances/distribution pairs would have been $\binom{25}{2} = 300$ rather than the $\binom{10}{2} = 45$ pairs we used. Nevertheless, we saw excellent agreement with expectations.

This small population is especially evident in System 9 of Figure 3.2, which indicates an outlier in the sampler noise at 512 samples, where the sampler appears to find another posterior mode. An outlying posterior estimate in the sampler + downsampling noise indicated at 1448 samples appears to have done the same, and it is likely that the posterior is in fact being represented faithfully in the downsampled space and the noise is due primarily to sampler noise. A larger population of distances would have ‘smoothed out’ such artefacts of sampler noise.

Another interesting observation of Figure 3.2 is that the noise floor (solid blue line) is often *higher* than the sampler *and* downsampling noise. This is of course not generally the case (for the other BHB systems not shown) but this feature also suggests that the population sizes are too small. Despite this however, it was still quite clear from the study of the plots that the first four criteria given in Section 3.4.3.2 are already met at the $N_s = 256$ mark. In fact, convergence for 9 of the 12 systems occurs after only $N_s = 128$. However, this very low number can perhaps be explained in part by the fact that each sample’s value is indeed informed by its $2\mathcal{M}$ neighbours; all systems here have an MCS




Figure 3.2: A selection of plots of the means and standard deviations of the sets of combined, marginalised JS divergences (CMJS divergences), \overline{D}_{JS} , as functions of N_s . The top axis shows the fraction of the signal’s number of Nyquist samples that were required to be computed to evaluate the posteriors. The solid (/dashed) blue line shows the ‘anchored’ distance mean (/standard deviation) of the noise floor, the orange lines show the anchored downsampling *and* sampler noise. Similarly, the blue circles (/crosses) show non-anchored means (/standard deviations) of sampler noise, whereas the orange marks show the same for downsampling and sampler noise. The spikes in System 9 just demonstrate non-convergence of the sampler, which (fortunately) occurred once in each type of posterior.

of around 5, so that $N_s = 128$ corresponds to requiring roughly $N_M \sim 1400$ signal samples to be known to inform the strain values of the 128 decorrelated samples.

Unfortunately, there does not appear to be any particularly clear pattern or rule for determining specific signal dependent optimal downsampling rates; the convergence point does not seem to strongly depend on SNR, number of free parameters, Nyquist number of samples or signal length. However, visual inspection of the posteriors seems to indicate that posteriors that are already ‘nice’, i.e., with simple, smooth contours tend to require fewer samples than those that appear irregular and with poorly constrained pa-

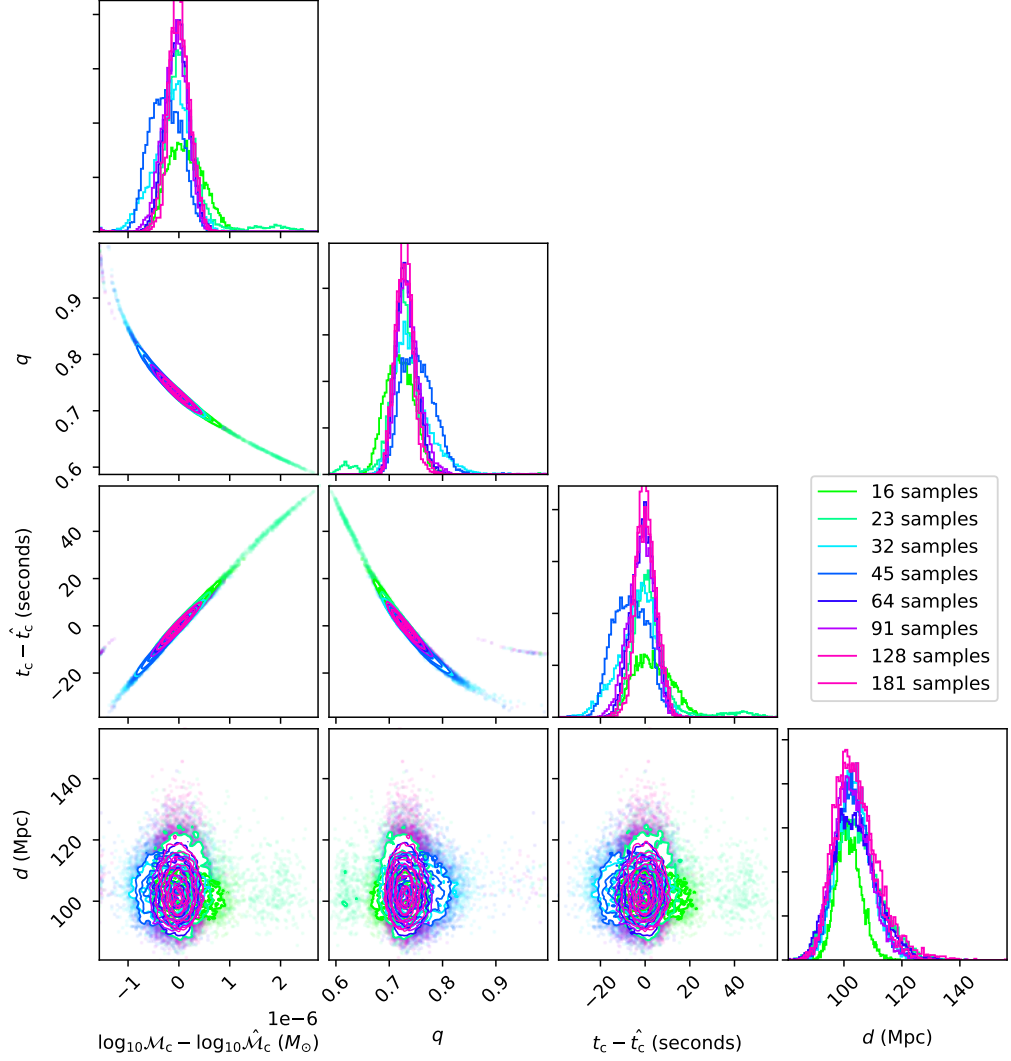


Figure 3.3: Downsampled posteriors of System 11 (Table 3.1) converge as the number of samples is increased. Only posteriors with $N_s \leq 181$ have been included to prevent a congested plot. There is perhaps zero discernible difference in these posteriors after $N_s = 128$. At low N_s , we can clearly see a ‘breakdown’ in the posterior, the few samples remaining being unable to constrain the parameters in the same way as the higher N_s posteriors. System 11 allows for the clearest version of this plot, note however that all other BHB test systems showed the same behaviour.

rameters. As a visual example of posterior convergence, we have attempted to show how the posterior of System 11 converges as the number of samples is increased in Figure 3.3.

We also studied how the FIMs converged to the true FIM defined by the full dataset, as per the fifth convergence criterion. In Figure 3.4 it is shown how the FIM of System 2 converges as N_s is increased. Similar plots for all the BHB systems we considered showed the same result: the (critical elements of the) FIMs converged to $\lesssim 3\%$ of the true FIM after around $N_s = 2^9 = 512$. Again, since this always occurs at the same N_s , independent

of the signal, there is unfortunately no obvious specific signal based reason or rule to be inferred.

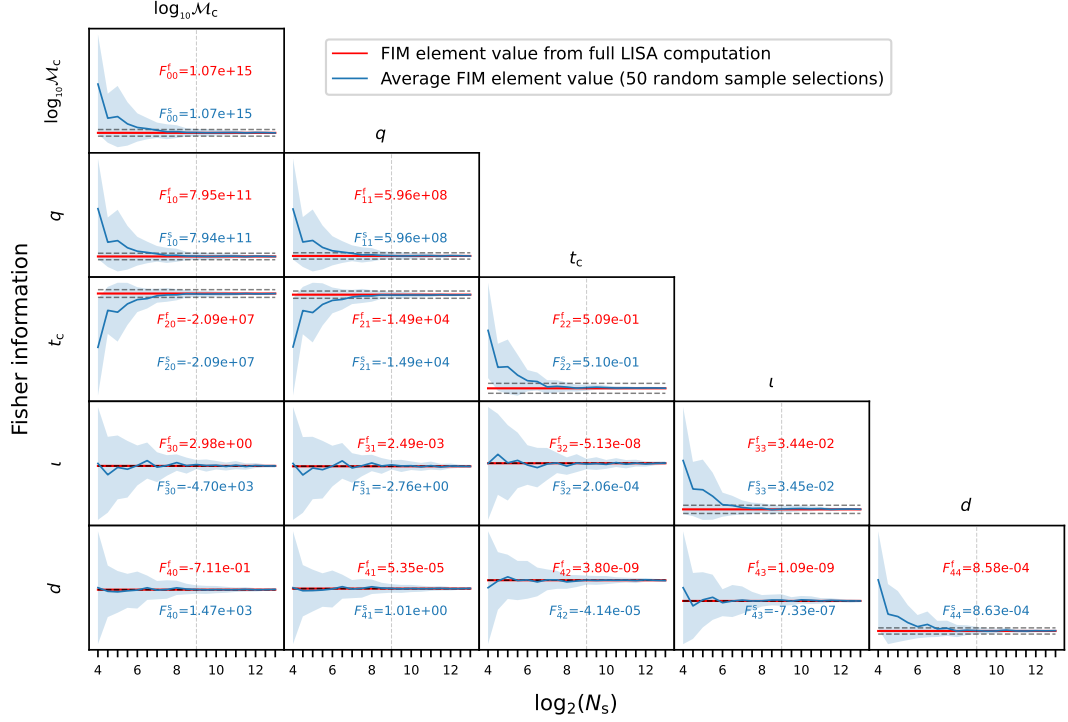


Figure 3.4: The lower triangular part of the FIM of System 2 (with $N_f \approx 3.5 \times 10^6$). Mean and standard deviation of 50 downsampling selections/realisations of \mathbf{F}_s are shown in blue for each $N_s \in \mathfrak{N}$. The diagonal FIM values are generally the most important, which we can see already agree with the full LISA FIM after around $2^9 = 512$ samples. Note that at very low samples, the Fisher information on the diagonals is overestimated. This is likely because (thinking of the FIM as the inverse PCM that describes the 5-ellipsoid that captures the shape of the peak of the likelihood) minimising the (CMJS) divergence is equivalent to maximising the overlap of the ellipsoids; if they are poorly aligned with dissimilar aspect ratios, then the downsampled ellipsoid should be made smaller to fit inside the true ellipsoid, smaller ellipsoids being equivalent to higher information. The y-axis scale is missing for clarity, but the dashed lines above and below the true Fisher information values in red show the upper and lower 5% intervals, respectively (i.e. $1.05 \times F_{ij}^f$ and $0.95 \times F_{ij}^f$ respectively). Also note that the off-diagonals in d and t are quite insignificant; if one reparameterises this matrix so that the diagonals are equal to 1 for example (divide the i^{th} row and i^{th} column by the square root of the i^{th} diagonal) one finds that these off-diagonals are approximately zero, where being out by a few orders of magnitude is still approximately zero and thus still correctly represents the true FIM.

Finally, as we shall see in the following Section (in Figure 3.5) the fastest PE completion time occurs for those posteriors defined using $N_s = 256$. The relationship between PE completion time and number of samples is of course not linear, as we shall discuss

in further detail, but empirically we have gained strong evidence that, for slowly evolving signal models and according to the convergence criteria we set out in Section 3.4.3.2, the optimal downsampling number of samples (independent of the signal duration) is approximately equal to 512.

3.5 Speed comparison and further remarks

3.5.1 Posterior estimate evaluation times

For each BHB test system we defined in Table 3.1 and for each $N_s \in \mathfrak{N}$, we produced 20 posterior estimates (recall we used $K = 10$ for each posterior type, p_{fl}^i and p_{ds}^i). The averages of each of those sets of 20 evaluation times are plotted in Figure 3.5. There is clearly a drop in evaluation time as the number of samples is dropped from 2^{13} to around 2^8 , for all BHB systems. For most systems, continuing to drop the number of samples shows starkly that evaluation time begins to increase again, this is almost certainly because the posteriors begin to qualitatively break down, the number and/or irregularity of posterior modes increases, and the sampler has a more difficult landscape to navigate. Roughly after the point at which the posterior stabilises, increasing the number of samples only increases the amount of data sample evaluations required to compute the likelihood, obviously (by definition) without affecting the likelihood, thus the posterior evaluation time should increase roughly linearly with increasing N_s . We see in Figure 3.5 that this is indeed the case, where a line of best fit is drawn through those average times for which $N_s \geq N_s(\text{min(PE time)})$. Write the linear average PE time model

$$\begin{aligned} T(N_s) &= aN_s + c \\ &\equiv b(2\mathcal{M} + 1)N_s + c \approx bN_{\mathcal{M}}(N_s) + c, \end{aligned} \quad (3.69)$$

using (3.24), so that

$$b = a(2\mathcal{M} + 1)^{-1},$$

where T is the average evaluation time, a is the line's gradient and c is a constant time offset that is attributed to the 'overheads' of the sampler (internal processes such as initialisation, computing the Fisher matrix, *normalising flow* transformations in the case of Nessai, etc.).




Figure 3.5: The average convergence times for all BHB systems defined in Table 3.1. The relationship is strongly linear after $N_s \gtrsim 2^8$ since the posterior has stabilised and is at its simplest; for $N_s \lesssim 2^8$, posterior estimation time increases as the number of samples decreases, very likely due to the structure of the posterior breaking down, becoming highly multi-modal and more difficult for the sampler to explore.

The linear model of posterior evaluation time can be used to estimate the time it would take to evaluate the posterior using all the samples (when $N_{\mathcal{M}} = N_{\text{Nyq}}$) from which we can roughly gauge the amount of time saved by employing our downsampling procedure. Simulated analyses of LISA data are often done in the frequency domain because of the great simplicity of computing the inner product that is afforded by the independence of the data samples in the frequency domain. The comparison of evaluation times shall not be especially precise because of this (the number of numerical operations required to compute signal templates will differ slightly between time domain and frequency domain signal models). Furthermore, the time domain analysis we have undertaken here requires a few more simple mathematical operations (see Section 3.3.1.3) in the form of element-wise products. It is reasonable to assume the cost of these operations are quite negligible in relation to the cost of evaluating the signal at the required points however, and particularly if one uses more realistic GW signal models defined more precisely, con-

#	N_{Nyq}	\mathcal{M}	a	b	c	$\Delta T_{\mathcal{M}}$	ΔT_{N_s}	$\Delta T_{\mathcal{M}}/\Delta T_{N_s}$
1	3.04×10^6	5	3.08×10^{-3}	2.80×10^{-4}	5.18	852.14	1.54	553.57
2	3.66×10^6	5	3.69×10^{-3}	3.36×10^{-4}	9.93	1229.43	1.85	665.74
3	2.08×10^7	5	1.24×10^{-3}	1.13×10^{-4}	0.83	2347.10	0.62	3787.27
4	7.69×10^6	5	4.41×10^{-3}	4.01×10^{-4}	6.75	3081.46	2.21	1397.48
5	6.51×10^5	4	5.96×10^{-3}	6.62×10^{-4}	8.95	431.06	2.98	144.71
6	2.37×10^7	6	1.45×10^{-3}	1.11×10^{-4}	1.89	2636.23	0.72	3639.11
7	3.65×10^7	5	4.08×10^{-4}	3.71×10^{-5}	0.28	1355.24	0.20	6642.85
8	1.95×10^6	5	1.54×10^{-3}	1.40×10^{-4}	1.34	272.63	0.77	354.82
9	3.94×10^7	6	1.88×10^{-3}	1.45×10^{-4}	2.43	5698.05	0.94	6065.19
10	3.66×10^6	5	3.26×10^{-3}	2.96×10^{-4}	6.62	1085.55	1.63	665.74
11	2.37×10^7	6	1.31×10^{-3}	1.01×10^{-4}	1.31	2384.35	0.66	3639.11
12	3.94×10^7	6	1.54×10^{-3}	1.18×10^{-4}	2.23	4664.40	0.77	6065.19

Table 3.2: Details of the linear best fit model of posterior evaluation times. From left to right, the columns show: BHB system number (see Table 3.1), number of Nyquist samples for signal, maximum correlated neighbouring samples (see Section 3.3.1.3), the three constants in Equations (3.69), the estimated original posterior evaluation time *minus the overhead time* (hours), $\Delta T_{\mathcal{M}} = \Delta T(N_{\mathcal{M}} = N_{\text{Nyq}})$ as defined below in Equation (3.70), the average downsampled posterior (using $N_s = 500$) evaluation time minus the overhead time (hours), $\Delta T_{N_s} = \Delta T(N_s = 500)$ and finally the ratio of these times (this is the approximate increase of the posterior evaluation rate, after disregarding the constant overheads).

sisting of many more terms in the expression of the strain model, $h(t)$. With these caveats mentioned, we can see, in Table 3.2, the linear PE evaluation time model parameters, along with the comparison between the expected evaluation time of the original posterior with the downsampled posterior.

Note that in Table 3.2 we use the definition

$$\Delta T \equiv T - c, \quad (3.70)$$

so that, as well as comparing the estimated evaluation times of the original and downsampled posteriors (minus overheads) which of course is of particular interest, we can also see that the quotient is simply given by

$$\frac{\Delta T(N_1)}{\Delta T(N_2)} = \frac{aN_1}{aN_2} = \frac{N_1}{N_2}, \quad (3.71)$$

and is independent of the slope. The increase in rates are thus equal for systems 2 & 10, 6 & 11, and 9 & 12 since this estimate is obtained from the definition of the linear model.

3.5.2 Discussion

The evaluation rate increases are significant, as can be seen in the rightmost column of Table 3.2; the posteriors for the signals we considered are produced from between 144 to 6642 times faster (System 5 and 7, respectively) than the estimated time of completion of the original posteriors defined with N_f samples. Importantly, the posteriors undergo very little distortion after downsampling, so we have produced strong empirical evidence here that the downsampling procedure is a successful approach to cheaply reproducing the likelihood function in simulated environments, and is robust in doing so, at least for the specific types of signal we have tested here.

This speed up certainly comes primarily from the much faster *likelihood* evaluation times, which is consistent with (3.71) and the linear evaluation time model as a function of number of samples. Thus for more intricate signal models with likelihood functions that are generally more costly to compute, the approximation techniques like the one expounded here, and others, such as ROQ, become essential tools for experimentation in the PE environment with LISA-like signals that are prohibitively expensive to compute. Indeed it will probably become worth investigating whether and how these techniques can be combined, to produce an even more rapid feedback arena for studying CBC signals and their various waveform modifications in LISA.

We have shown that our example signals have been successful subjects of the downsampling procedure defined in this work, and that we should expect to be able to rely on posteriors generated using the procedure as accurate reproductions of the original posterior, as much as we can rely on posteriors recreated using the posterior samples returned by a sampler. However, it must be stressed that this has only been tested, and is only expected to produce very accurate results for the cases of slow evolution signals; it is plausible for the accuracy to drop when signals depend on some parameter whose average information changes significantly on timescales roughly of order of the average time between samples, or faster.

The overlaid posterior plots (one of which is shown in Fig. 3.3) for all of our signals *appear* to converge after around only 200 samples, and indeed this would be consistent

in general with the CJMS distance convergences, some of which are shown in Fig. 3.2. As a result, it may seem reasonable to suppose convergence could be taken to be sooner than $N_s = 2^9 = 512$. However, these apparently earlier convergences are likely due to losing some of the detail of the posterior structure by performing the marginalisations for the CMJS distances and in generating the corner plots. The fact that the FIMs converge later than the CMJS distances and corner plots reinforces this view, and we should be cautious to trust the various marginalisation convergences without support from elsewhere, such as the converged FIMs. The FIMs contain information, albeit in a small, localised region, on all of the parameters simultaneously. Whilst the accuracy is likely already more than high enough anyway using around 500 samples for most applications (for example, considering the other posterior noise sources), recall that one can always choose to define the likelihood function using more samples as further assurance the accuracy will be very high.

This does prompt further interest in the approach previously mentioned in Section 3.3.3. That is, to downsample the dataset, then set the resulting FIM equal to the FIM of the full dataset. However, due to the uneven weighting given to the different data samples in this approach, which tightens up the approximation of the primary mode of the posterior, it is possible that secondary modes might be adversely affected more using this approach than they otherwise would. This method might therefore be suitable for other certain types of signal analysis problem, particularly, for example, where it is known that the PCM is a reasonable approximation of the posterior. It may be possible to extend the method of Section 3.3.3 however, by also choosing a few random points in the parameter space and allowing the reweighting function a higher degree polynomial that could be solved for, such that the likelihoods also agree exactly at those random points. Or one might perhaps consider setting higher order central moments of the distributions to be equal, i.e., the higher dimensional analogues of the distribution's skew, kurtosis.

The Low-Frequency Landscape

In this chapter, we work towards understanding and modelling waveform modifications of BHBs in galactic and modified gravity environments. We hope this can be later used to make progress in discovering the astronomy and physics potential of these sources, at which point the requirements and scope of the model for the analysis of LISA data will be updated. By use of the downsampling procedure explicated in the previous chapter, we will analyse data in the time-domain, where it is more straightforward to define the model; one can simply evolve a system forward in time given some initial conditions (to reproduce arbitrary time-delays in the frequency-domain, complicated frequency modulations are often required to be computed). Recently, progress has been made with frequency-domain modelling of BHBs in galactic orbits however, as in Ref. [99], by making simplifying assumptions (taking galactic orbits to be circular) that allow one to analytically derive the frequency modulations. Nevertheless, our time-domain version is particularly useful in easily allowing inclusion of additional physical effects that may be more difficult to derive in the frequency-domain, for example, eccentric orbits and *periastron precession*, and sub-leading order relativistic time-delays, which will be described in more detail in due course.

Additionally, to preempt the data analysis problems that naturally arise when the dimensionality of the parameter space becomes large (which occurs in particular for long duration signals with waveforms generated by CBCs in orbit around supermassive black-holes, see 4.2.2.1 below), we will investigate the potential of lower dimensional ‘waveform (modification) approximants’. These approximants will be defined for the modifi-

cations to the standard BHB waveform that arise due to the orbital motion of the BHB around a central mass and further relativistic effects in Section 4.2.4, and tested in Section 8.5.

If there are inconsistencies between the true physical waveform and our approximations, then, given they are small enough, our simulations should still accurately reproduce the posterior one would obtain by a ‘correct’ analysis (one free of modelling errors and oversights) of real-world data. Our model templates are used to search for a simulated injected signal which is itself produced using the model we will define. However, larger discrepancies might result in an unreliable posterior by, for example, incorrectly inferring parameter interdependencies, or showing that more information exists on a given parameter than one has in reality, so even in this early investigative stage one should try to accurately reproduce the signal (and detector) models, and be very clear about when and where they are valid. Later, in Section 6.3, we will present a list of such approximations, omissions and assumptions, with brief discussion on potential effects on the posteriors.

4.1 Waveforms and modifiers

This section aims to briefly introduce some fundamental physical and astrophysical effects that can modify the GW of an inspiralling BHB, and also to give an idea of the vastness of the range of concepts that need to be considered before one can be confident about the models that will be used to analyse the LISA data. There exist relatively comprehensive resources in the literature, appearing generally as compilations of short descriptions and studies of selections of these effects, often with derivations of the approximate impact on the waveform one can expect. These will be referenced where appropriate, as the effects are discussed.

One can attempt a categorisation of the various effects, as we shall here, although there is often significant overlap of certain effects between categories so this is rather imprecise, but has some use for a big picture view of the modelling landscape. We divide effects into: essential GR effects (‘known-knowns’), potential astrophysical environmental effects and modified gravity waveform modifications (‘known-unknowns’) and other, unknown, unmodelled effects (‘unknown-unknowns’).

4.1.1 Known-knowns

This category describes the ‘known’ physics (some foundations are *assumed* as certainties, i.e., GR and the existence of triple BH systems) that may significantly affect the waveform, such as relativistic effects and effects that derive from known astrophysical scenarios. We will suppose that a BHB may be in proximity of some other mass, in a 3-body system (see Section 4.2.1 below). Some effects we should consider include:

- *Hubble expansion*, [100]. The stretching or *redshift* of the waveform during propagation due to the expansion of space.
- *Peculiar motion* of GW emitter’s host galaxy, [100].
- *Line-of-sight motion* due to emitter’s orbit of large mass, leading to an evolving *Roemer delay*: the time-of-arrival delay of a signal due to a changing relative distance between the source and the detector from orbital motion. There are Roemer delays from Earth’s orbit around the Sun and from the GW emitter’s orbit around its host galaxy’s centre. See Section 4.2.2.2.
- *Periastron precession*. See Section 4.2.2.3.
- *Relativistic time dilation*. See Section 4.2.2.4.
- *Relativistic beaming*, [101]. The emissions of an object in relative motion are focused towards the direction of motion.
- *Gravitational lensing*, [102]. Plane GW wavefronts undergo focusing when impinging upon the ‘gravitational lens’ that is the localised spacetime curvature arising from clumps of mass, such as galaxies.
- *Tidal deformation* of binaries/eccentricity excitation, due to proximity of binary to a supermassive BH, [103].
- *Kozai-Lidov mechanism*, [104]. For certain special cases of 3-body gravitating systems, in particular a main binary with a third, more distant orbiting body, the third body’s inclination with respect to the orbital plane of the central binary is correlated with its eccentricity, and these parameters can ‘exchange’ periodically, modifying the GW emission of the main binary. The ‘outer orbit’ between the inner binary and third body may also contribute directly detectable GWs.

4.1.2 Known-unknowns

This category covers physical scenarios that may give rise to GWs for which we either lack a model, or for which we may not know the framework in which to derive the model. Generally, these sorts of concerns include: the comprehensive mass distribution (with its initial conditions); the gravitational waveform as a function of that distribution (where we are now dropping the assumption that GR is correct); other proposed physical mechanisms; or a combination of these things. We know that we don't *always* know how, whether we should in general, or when in particular, to model these effects. The details are becoming ever more clear as researchers study the various scenarios and refine expectations. However, there will exist cases where significant doubts remain about the correctness of the model. Some of these physical effects (the below is vastly condensed list of those we can consider) include:

- *Emission or propagation environmental matter effects* – Accretion onto BHB components, [105, 106]; Weak/micro lensing (BHBs located in dense star clusters), [107, 108]; Dynamical friction, [105].
- *Matter field effects* – Matter field interactions at point of GW generation, for example, from the axion field, [109].
- *Modified gravity theories* – Scalar/vector GW modes, [110]; Extra dimensions, [110]; Higher-order gravity and order dependent matter couplings, [110].

4.1.3 Unknown-unknowns

The true waveforms may well be influenced by physics not yet known or conceptualised. The waveforms (or modifications thereof) are thus themselves strictly unmodelled (see also discussion on unmodelled ‘burst’ waveforms Section 1.3.2.3). This could be generated by, for example, the true mechanisms and interplay of matter and gravity (supposing it has not yet been recognised), or new, not yet considered matter fields, for example. We may also suffer from mistakes in our model derivations.

4.2 Establishing a foundational signal model

From the previous discussion on waveform modifiers, expected and speculated, the reality of the vast potential complexity of the model is easy to see. However, as shown in Ref. [105], the actual modifications of the gravitational waveforms from the various physical and environmental effects are largely negligible for the majority of expected scenarios. We intend to reproduce and study the major modifiers of long duration signals that are expected to be entirely *necessary* for accurate signal modelling in the first instance. Since perhaps most (depending on the main binary formation scenarios) of the CBC signals in LISA are expected to be generated from binaries undergoing detectable local accelerations, which can cause very large phase modulations as we will later see, this aspect of waveform modelling will be prioritised.

It has already been confirmed that the expansion of the Universe and the peculiar motions of galaxies have no directly measurable imprint on GWs for the presently conceptualised detector networks [100]. The changing relative position of a GW emitter (a binary) due to its orbit about the centre of mass of a galaxy can however lead to major modifications of the observed waveform. The following subsections describe the physical sources and derive the major components of the model modifications.

4.2.1 Binary formation scenarios

There are two main formation scenarios for BHBs: formation in the *field*, and formation in *dense star clusters* (see Figure 4.1), though these channels alone have already been found to be insufficient to account for the GW observations to date [111] and a finer formation model is required. However, the extra detail in *formation* modelling still generally derives from these two main channels [111], so the precise details do not affect the *waveform* modelling requirements¹ (which is what we are presently concerned with) since we can still say that binaries either derive dynamically or in the field.

Formation of BHBs in the field, which covers the disk (if one exists) and halo, otherwise referred to as formation ‘in isolation’, occurs when the individual components

¹Unless, for example, those models say something such as that 3-body sources dominate GW events where interactions from a third body disrupts the waveform of the ‘close binary’ on timescales less than the LISA mission duration, which would of course strongly inform the waveform modelling requirements.

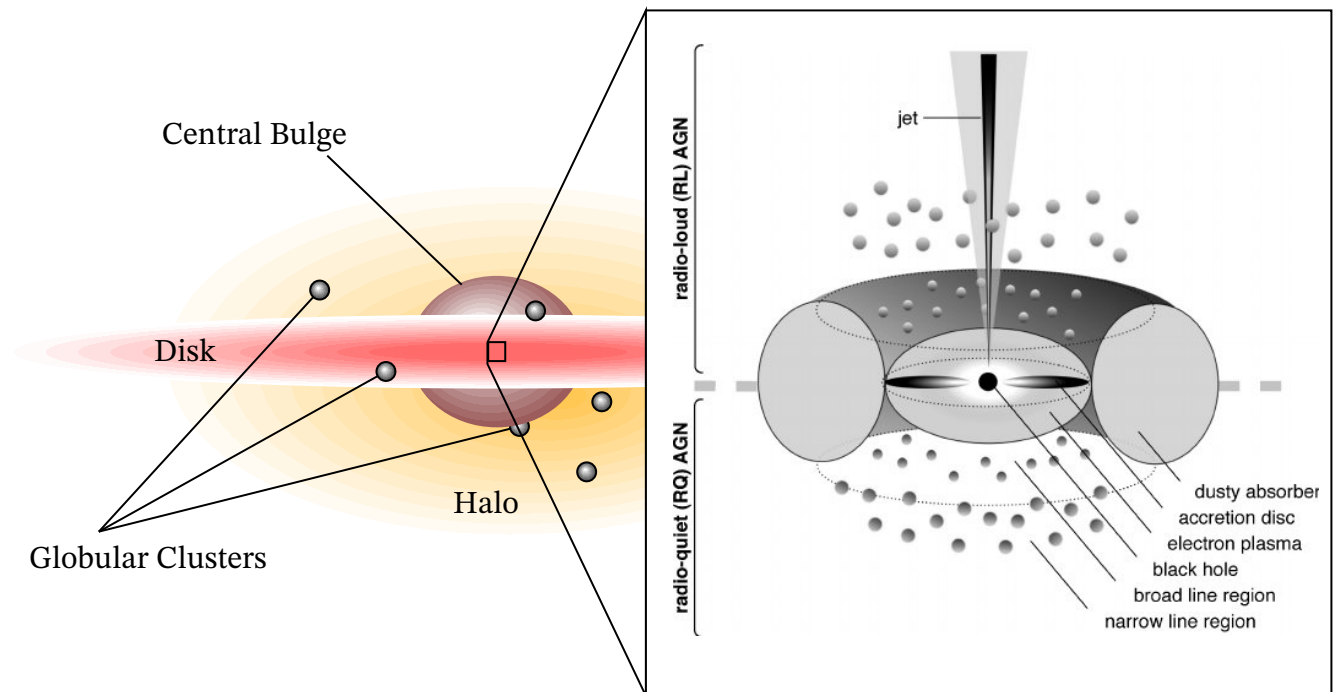


Figure 4.1: Simplified structure of a typical spiral galaxy (‘side’ view), left, with zoomed in inset (from Beckmann & Shrader (2012), [112]) of nucleus, right. The disk and halo constitute the *field* where binaries form ‘in isolation’, and the globular clusters, central bulge [113] and *nuclear star clusters* constitute dense star clusters, where binaries form ‘dynamically’ [114]. Binaries may also form dynamically in the *accretion disk* (inset; note the accretion disk, right, is distinct from the disk of the galaxy, left). Elliptical galaxies do not have a disk.

of massive stellar binaries in the field evolve into black-holes. Dynamical formation of BHBs, on the other hand, occurs via sequential interactions of massive stars or stellar remnants inside the high density, dynamical environment of star clusters [115]. Binaries from each formation scenario are therefore expected to have different properties; for example, field binaries are expected to have very low eccentricities (because of the isolated environment allowing their orbit to slowly circularise without interference), compared to those formed in dense clusters (because they form rather chaotically from close encounters with other dense objects). The statistics of many GW event observations will allow us to discriminate between competing formation models [116].

The high density star clusters include *globular cluster (GC)*s and *nuclear star cluster (NSC)*, where the latter (depending on the specific galaxy properties) resides in the *central bulge*, which also hosts the supermassive black hole (SMBH), as depicted in Figure 4.1. Increasingly more massive BHs are thought likely to arise from successive, ‘hierarchical’ mergers of progenitor BHs [114]. In the smaller globular clusters (10^3 – $10^5 M_\odot$), we can expect mergers of low-mass binaries that comprise of direct stellar remnants, but these clusters are not likely to host BHs more massive than around $100 M_\odot$; when the direct stellar remnant binaries merge, they receive a gravitational kick, often large enough to escape the GC’s relatively small escape velocity [114]. The NSC, on the other hand, are generally much more massive (10^6 – $10^8 M_\odot$), with higher escape velocities that are more difficult for the resultant, post-merger BHs to escape from, and can therefore go on to form hierarchical, more massive BHBs. SMBHs may have an *accretion disk*, a disk of gas and dust that can fall and accrete onto the SMBH. These disks can also host BHBs, and there is good support for the recent GW event GW190521 to have originated from such an environment [117].

We can therefore expect more massive BHBs to reside in the central bulge, where they undergo considerable accelerations. A study of NSC in Ref. [118] shows that NSC have ‘effective radii’ (from within which half of the NSC electromagnetic emissions originate) of around $3.3_{-1.9}^{+7.0}$ pc, and that they are generally highly elliptical [118]. Although ellipticity does not necessarily directly translate to high eccentricity stellar orbits around the central mass, more highly eccentric orbital motion is clearly likely. Furthermore, since hierarchical remnant BHs receive essentially random kicks, the orbits of binaries they go on to form will likely span a broad range of eccentricities. We shall therefore work

towards producing a model that describes eccentric orbits.

4.2.2 Orbits & time-delays

In this section we describe and derive the major fundamental and necessary properties of orbital motion and the main contributors to the gravitational waveform modifications.

4.2.2.1 Orbital parameters

The standard parameterisations of the orientations of elliptical orbits in astronomy are the well-known (‘Keplerian’) *orbital elements* [119]. Three of this set of six parameters describe the size, shape, and a reference point of an elliptical orbit, and the remainder are *Euler angles*, which define rigid, 3D spatial rotations, and thus may describe a Kepler orbit in a reference frame different to the standard source frame of the orbit. However, a transformation of the usual Euler angles commonly known as the *Tait-Bryan* angles (used heavily in aviation) in their extrinsic form (where rotation axes are not themselves




Figure 4.2: A graphical representation of the 3D rotation angles used to specify the orientation of the Kepler path of a CBC’s orbit around a central mass. The *Tait-Bryan* or ‘aircraft principal axes’ have been implemented as this allows for a useful simplification of the parameter space without loss of generality: with the aircraft pointing along the line-of-sight from detector to source, the ‘roll’ parameter is not coupled to the GW and can be ignored. This formalism is used for being very simple and intuitive for the problem at hand. The depiction of the elliptical orbit in the α - γ plane with focus at the origin shows the source frame orientation of the orbit.

Name	Symbol	Units	Range	Prior
Pitch	α	rad	$[0, \pi]$ (not cyclic)	Sine
Yaw	β	rad	$[0, 2\pi)$ (cyclic)	Uniform
Semi-major Axis	a	pc	$(f(M_*, e), a^{\max})$	Pop. Model
SMBH Mass	M_{SMBH}	$10^6 M_{\odot}$	$[M_{\text{SMBH}}^{\min}, M_{\text{SMBH}}^{\max}]$	Pop. Model
Eccentricity	e	-	$[0, 1)$	Pop. Model
Initial Eccentric Anomaly	E_0	rad	$[0, 2\pi)$ (cyclic)	Uniform

Table 4.1: Details of the Keplerian elements. The ‘Pop. Model’ prior indicates that the prior could be derived from population models, we will use uniform priors here. The lower limit of the semi-major axis is some function of the SMBH mass and the orbital eccentricity which is a bound for stable orbits. For simplicity we will ignore details of those GR effects that make an orbit unstable and just ensure that an orbit does not traverse the SMBH event horizon, thus $f(M_*, e) = R_{*,\text{Schw}} + y$, where $R_{*,\text{Schw}}$ is the Schwarzschild radius of the SMBH and y is the distance from the centre of the ellipse to the SMBH. The constants a^{\max} , M_{SMBH}^{\min} & M_{SMBH}^{\max} are also determined by astrophysical models.

rotated), are particularly useful for the current purposes from the point of view of ease of modelling, as shall become clear in the following sections.

The orientation formalism depicted in Figure 4.2 puts the line-of-sight to the GW source along the γ -axis (that is, in the direction of travel of the aircraft). Then since the background spacetime of the SMBH is taken to be spherically symmetric, a single ‘roll’ will not affect the specific time-delay modifications of the waveform that we shall be modelling here; rolling of the source frame is equivalent to the ‘inverse roll’ of the detector frame. But rolling of the detector frame is itself equivalent to changing the polarisation angle of an unmodified, ordinary BHB. Thus the roll parameter (γ) can be absorbed into the BHB polarisation parameter (ψ). Similarly, for non-accelerating BHBs, a single (instantaneous) ‘yaw’ can be absorbed into the BHB phase, and a single ‘pitch’ can be absorbed into the BHB inclination. In general however, one must consider orbital motion where these operations result in variable time dependent time-delays which cannot be disregarded. Note that the disregard of the roll parameter is not possible if in orbit around a spinning SMBH where the SMBH spin is significantly imprinted onto the BHB waveform during propagation, or if a *spinning* SMBH transfers angular momentum onto the BHB during waveform generation; the detector response to polarisation is in that case not completely degenerate with these finer waveform modifiers.

The transformation between the orbital elements i , Ω and ω (*inclination*, *longitude of the ascending node*, and *argument of periapsis*, respectively) can be found by considering, in the usual way, a decomposition of the 3D Euler rotation matrices as a product of matrices that rotate the frame around 3 separate axes [120]. For example, as per [120], the 3D Euler intrinsic rotation matrix $\mathbf{R}_{313}^{\text{in}}$ can be written as

$$\mathbf{R}_{313}^{\text{in}}(\omega, i, \Omega) = \mathbf{R}_3^{\text{in}}(\omega)\mathbf{R}_1^{\text{in}}(i)\mathbf{R}_3^{\text{in}}(\Omega), \quad (4.1)$$

for rotation matrices \mathbf{R}_k^{in} that are responsible for rotating around the ‘new x^k -axis’ (the frame orientation one obtains after all previous rotations).

Another valid set of Euler rotations is the aforementioned Tait-Bryan sequence with extrinsic 3D rotation matrix given by ordered combination of 2D rotations

$$\mathbf{R}_{123}^{\text{ex}}(\alpha, \beta, \gamma) = \mathbf{R}_1^{\text{ex}}(\gamma)\mathbf{R}_2^{\text{ex}}(\beta)\mathbf{R}_3^{\text{ex}}(\alpha), \quad (4.2)$$

where the matrix \mathbf{R}_k^{ex} is responsible for rotating around ‘fixed x^k -axis’. Note that in this formulation, the last (or leftmost, for left-acting rotation operators of column vectors) single axis rotation is the roll $\mathbf{R}_1^{\text{ex}}(\cdot)$, which can be ignored as per the preceding discussion. For reference, the axes in Figure 4.2 have corresponding indices $\alpha \sim x^3$, $\beta \sim x^2$, $\gamma \sim x^1$.

It is however still important to know the relationship between the different sets of angles; we should know what the intrinsic orbital elements (the standard parameterisations in use in astronomy) are, in terms of the extrinsic Tait-Bryan angles we have employed. To find the relationship, first note that the action of $\mathbf{R}_{313}^{\text{in}}$ on an arbitrary vector \mathbf{v} given some (ω, i, Ω) , must give precisely the same vector as that resulting from the action of $\mathbf{R}_{123}^{\text{ex}}$ on \mathbf{v} , for the unknown (α, β, γ) ; the matrices are equivalent:

$$\mathbf{R}_{313}^{\text{in}}(\omega, i, \Omega) = \mathbf{R}_{123}^{\text{ex}}(\alpha, \beta, \gamma). \quad (4.3)$$

We could compute the dependencies of (ω, i, Ω) on (α, β, γ) directly from the components of this matrix. There is a useful relationship, however, between the intrinsic and extrinsic axis rotations: a sequence of n intrinsic rotations about axes $x'_{k_1}, x''_{k_2}, \dots, x^{(n)}_{k_n}$, where $k_j \in (1, 2, 3)$ and the primes denote successive coordinate frames, by angles $\theta_1, \dots, \theta_n$ (that is, successive rotations by $\mathbf{R}_{k_j}^{\text{in}}(\theta_j)$ matrices for $j \in 1..n$) is equivalent to a corresponding sequence of n extrinsic rotations in reverse order, about axes $x_{k_n}, x_{k_{n-1}}, \dots, x_{k_1}$ by angles $\theta_n, \dots, \theta_1$ (that is, successive rotations by $\mathbf{R}_{k_j}^{\text{ex}}(\theta_j)$ matrices for $j \in n..1$) [121].

Thus $\mathbf{R}_{313}^{\text{in}}(\omega, i, \Omega) = \mathbf{R}_{313}^{\text{ex}}(\Omega, i, \omega)$. Now $\mathbf{R}_{313}^{\text{ex}}$ and $\mathbf{R}_{123}^{\text{ex}}$ are given in Ref. [120], along with the required conversion between the different sets of angles, so from Ref. [120]

$$\begin{pmatrix} \Omega \\ i \\ \omega \end{pmatrix} = \begin{pmatrix} \arctan2(-\sin \beta, \sin \alpha \cos \beta) \\ \arccos(\cos \alpha \cos \beta) \\ \arctan2(\cos \alpha \sin \beta \cos \gamma + \sin \alpha \sin \gamma, -\cos \alpha \sin \beta \sin \gamma + \sin \alpha \cos \gamma) \end{pmatrix} \quad (4.4)$$

where $\arctan2$ is the ‘2 argument arctangent’ function.

4.2.2.2 Roemer delay

Here we model the line-of-sight velocity of a BHB due to its orbit around a SMBH. We assume a Keplerian (elliptical, Newtonian) orbit here, and will add the approximate periastron precession modification arising from GR in the following section. We write a , b and e for the semi-major, semi-minor axes, and eccentricity respectively, which satisfy

$$b = a\sqrt{1 - e^2}.$$

The so-called *mean anomaly* and *eccentric anomaly*, denoted M and E , parameterise time and polar angle respectively, where we set $M_0 = M(t_0)$ and $E_0 = E(t_0)$, for an initial time t_0 , which we generally take to be the instant an observation begins, i.e., the first time stamp of the LISA signal. They are given by [122]

$$E = M + \sum_{n=1}^{\infty} \frac{2}{n} J_n(ne) \sin(nM) \quad (4.5)$$

and we reiterate that the e here is eccentricity (and not Euler’s constant), and

$$M = \sqrt{\frac{GM_{\text{SMBH}}}{a^3}}(t - t_{\text{Peri}}) \quad (4.6)$$

where M_{SMBH} is the total mass of the SMBH, t_{Peri} is the time at periastron and G is Newton’s gravitational constant. The J_n are Bessel functions of the first kind. In Euclidean coordinates in the source frame, the position of the BHB has the non-zero components

$$x_{\text{src}} = a(\cos E - e), \quad (4.7)$$

and

$$z_{\text{src}} = b \sin E. \quad (4.8)$$

With rotation matrix \mathbf{R} , and writing the source frame position as the vector $\mathbf{P}_{\text{src}} \equiv (x_{\text{src}}, 0, z_{\text{src}})^T$, the observer frame position will be given by the vector

$$\mathbf{P}_{\text{obs}} = \mathbf{R}\mathbf{P}_{\text{src}}. \quad (4.9)$$

We will be interested in the line-of-sight coordinate (our $\hat{\mathbf{z}}$ direction) and total relative velocity of the BHB, denoted v . The z -coordinate is given by

$$z_{\text{obs}} = \mathbf{P}_{\text{obs}} \cdot \hat{\mathbf{z}}, \quad (4.10)$$

The change in travel time of the GW due to its changing relative distance to the observer results in an effective changing time of arrival of the waveform. Primarily, this provides a modulation in the phase of the signal (and a negligible change in the amplitude). The observed waveform from this Roemer delay alone then is simply given by a time-shift of the source waveform

$$h_{\text{obs}}(t) \equiv h_{\text{src}}(t'(t)), \quad (4.11)$$

where the shifted time, t' , is given by

$$t'(t) = t - z_{\text{obs}}(t)/c, \quad (4.12)$$

since it takes time z_{obs}/c for the GW to traverse the extra distance.

The orbital motion has so far been considered in a Newtonian gravity framework, so the velocity of the BHB does not yet contain relativistic corrections (nor have effects on h_{src} been considered that arise from proximity to the SMBH). These effects, including relativistic corrections to the orbit, can be substantial in the environments being considered, we describe how they enter into the model in the following subsections.

4.2.2.3 Periastron precession

To derive the precise motion of the binary around a SMBH, which we model as a test particle orbiting a central mass, an exact solution of the fully non-linear geodesic equations with metric given by the EFEs (Schwarzschild or Kerr) is required. The full treatment of the problem is highly non-trivial, but in the interests of producing accurate parameter estimation results, we would like to include some GR trajectory corrections, since some extra information on the SMBH parameters are encoded in these corrections. However,

we will resort to the simple approximation of geodesics, with trajectories described by a *periastron precession* [123], well-known in the study of the Kepler orbit problem in GR.

For weakly precessing systems (where gravity is relatively weak), the apsidal precession over one full orbit of the binary around the SMBH is given by [123]

$$\Delta\varphi \approx \frac{6\pi G(M_{\text{SMBH}} + M)}{c^2 a(1 - e^2)}, \quad (4.13)$$

where M_{SMBH} is the mass of the SMBH, M is the *total* mass of the binary, a is the semi-major axis of the binary's orbit and e is its eccentricity.

At first approximation, the periapsis precesses at a constant rate (independent of the position of the binary with respect to the SMBH) i.e., that $\dot{\varphi} = \text{const.}$ Using this approximation, we have that $T\dot{\varphi} = \Delta\varphi$, where T is the period of the orbit, given by

$$T = 2\pi\sqrt{\frac{a^3}{GM_{\text{SMBH}}}}, \quad (4.14)$$

so the rate of precession of the periapsis is, taking $M_{\text{SMBH}} \gg M$

$$\dot{\varphi} \approx \frac{6\pi G(M_{\text{SMBH}} + M)}{c^2 a(1 - e^2)T} \approx \frac{3}{c^2(1 - e^2)} \frac{(GM_{\text{SMBH}})^{3/2}}{a^{5/2}}. \quad (4.15)$$

This is a good approximation for roughly circular orbits. However, most of the precession occurs in the part of the orbit closest to the SMBH [123] (i.e., the rate of precession is proximity dependent), so for highly elliptical orbits this simplification can be expected to break down.

The precession of the orbit revolves around the SMBH. However, when the origin of coordinates in which the Kepler orbit is defined is at the centre of the ellipse traced out by $\mathbf{P}_{\text{src}}(t) \equiv (x_{\text{src}}, 0, z_{\text{src}})^T$, we must first translate the coordinates such that $\mathbf{P}_{\text{src}}(t)$ is defined with respect to the SMBH, then rotate as per (4.15) before performing the inverse translation. Thus the matrix describing the precession is given by

$$\mathcal{P} = \mathcal{J}^{-1}\mathcal{R}\mathcal{J}, \quad (4.16)$$

which operates on the vector $(\mathbf{P}_{\text{src}}(t), 1)$, where \mathcal{J} and \mathcal{R} are, respectively, the translation and rotation matrices, given by

$$\mathcal{J} = \begin{pmatrix} 1 & 0 & 0 & -y \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \mathcal{R} = \begin{pmatrix} \cos \varphi & 0 & \sin \varphi & 0 \\ 0 & 1 & 0 & 0 \\ -\sin \varphi & 0 & \cos \varphi & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (4.17)$$

where y is the distance from the centre of the ellipse to the SMBH. We have

$$\mathcal{P} = \begin{pmatrix} \cos \varphi & 0 & \sin \varphi & y(1 - \cos \varphi) \\ 0 & 1 & 0 & 0 \\ -\sin \varphi & 0 & \cos \varphi & y \sin \varphi \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.18)$$

with $y^2 = a^2 - b^2$, where a, b are the semi-major and semi-minor axes, and φ is found from integrating (4.15):

$$\varphi \approx \frac{3t}{c^2(1 - e^2)} \frac{(GM_{\text{SMBH}})^{3/2}}{a^{5/2}} + C, \quad (4.19)$$

where C is a constant that can be set to zero without loss of generality; this is related to the initial periastron angle at the beginning of the observation and can be absorbed into the initial eccentric anomaly, E_0 .

Now, writing the *precessed* position vector as $\bar{\mathbf{P}}_{\text{src}}(t)$, (where the bar $\bar{}$ in general indicates precession modified quantities) we have that

$$\text{diag}(\bar{\mathbf{P}}_{\text{src}}(t), 1) = \mathcal{P} \cdot \text{diag}(\mathbf{P}_{\text{src}}(t), 1), \quad (4.20)$$

we rotate this into the observer's displacement in the same way we did for equation (4.9), so finally the observed, precessed displacement, $\bar{\mathbf{P}}_{\text{obs}}$, is given by

$$\bar{\mathbf{P}}_{\text{obs}} = \mathbf{R} \bar{\mathbf{P}}_{\text{src}}. \quad (4.21)$$

Following the steps in the previous section we also then have

$$\bar{z}_{\text{obs}} = \bar{\mathbf{P}}_{\text{obs}} \cdot \hat{\mathbf{z}}, \quad (4.22)$$

$$\bar{t}'(t) = t - \bar{z}_{\text{obs}}(t)/c, \quad (4.23)$$

again since it takes time $\bar{z}_{\text{obs}}(t)/c$ for the GW to traverse the extra distance; the periastron precession of the orbit just leads to a correction in the position, and so a corrected Roemer delay. We can write the *Roemer delay* as

$$\Delta \bar{t}_{\text{Roemer}} = \bar{t}'(t) - t = \bar{z}_{\text{obs}}(t)/c. \quad (4.24)$$

4.2.2.4 Relativistic time dilation

Approximating the SMBH as a Schwarzschild black-hole, the proper time interval along a curve is given by the line-element

$$c^2 d\tau_{\text{src}}^2 = \left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r}\right) c^2 dt^2 - \left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r}\right)^{-1} dr^2 - r^2 d\theta^2 - r^2 \sin^2 \theta d\phi^2. \quad (4.25)$$

We can always align the coordinates with the orbit direction such that the binary does not move in the ϕ -direction, so that $d\phi = 0$. The proper time interval is then

$$d\tau_{\text{src}} = \left[\left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r} \right) - \left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r} \right)^{-1} \left(\frac{dr}{c dt} \right)^2 - \left(\frac{r d\theta}{c dt} \right)^2 \right]^{1/2} dt. \quad (4.26)$$

The r - θ plane here coincides with the x - z plane used in the previous section for setting up the Kepler orbits. The dr/dt and $d\theta/dt$ terms can then be computed using (4.7) and (4.8), since

$$r = \sqrt{x_{\text{src}}^2 + z_{\text{src}}^2} = a(1 - e \cos E), \quad \text{and} \quad \theta = \cos^{-1}(x_{\text{src}}/r). \quad (4.27)$$

Take the detector to be at spatial infinity, where the coordinate time interval, dt , is equal to the detector's proper time interval, $d\tau_{\text{obs}}$. Then we integrate the above to obtain the proper time measured in the source frame, τ_{src} , as a function of the detector time, t :

$$\tau_{\text{src}}(t) = \int_{t_0}^t dt' \sqrt{\left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r(t')} \right) - \left(1 - \frac{2GM_{\text{SMBH}}}{c^2 r(t')} \right)^{-1} \left(\frac{dr(t')}{c dt'} \right)^2 - \left(\frac{r(t') d\theta(t')}{c dt'} \right)^2}. \quad (4.28)$$

We define the *Einstein delay*, $\Delta t_{\text{Einstein}}$, to be the difference between the proper time at source, τ_{src} , and the proper time at the detector, $\tau_{\text{obs}} = t$:

$$\Delta t_{\text{Einstein}}(t) \equiv t - \tau_{\text{src}}(t). \quad (4.29)$$

4.2.2.5 Shapiro delay

The *Shapiro delay* describes the propagation delay caused by the decreased ‘coordinate velocity’ of light (or, equivalently, due to the ‘extra’ physical distance light must travel) due to spacetime curvature caused by nearby mass. The ‘one-way’ trip Shapiro delay is given by (adapted from [124])

$$\Delta t \approx \frac{2GM_{\text{SMBH}}}{c^3} \left(\ln \left[\frac{z_p + (z_p^2 + r_0^2)^{1/2}}{-z_e + (z_e^2 + r_0^2)^{1/2}} \right] - \frac{1}{2} \left[\frac{z_p}{(z_p^2 + r_0^2)^{1/2}} + \frac{z_e}{(z_e^2 + r_0^2)^{1/2}} \right] \right), \quad (4.30)$$

to first order in the Schwarzschild radius $r_s = 2GM_{\text{SMBH}}/c^2$. Here, z_p , z_e and r_0 are, respectively, the distance between the orbiting body and *periastron* (the point of closest approach of the null path to the SMBH), the distance “along the line of flight” between observer and periastron, and the *impact parameter*; the distance from the gravitating body's centre to periastron. See Appendix B for more details. Note that

$$z_p = |\bar{z}_{\text{obs}}|.$$

Orbit	α (rad)	β (rad)	γ (rad)	a (pc)	$M_{\text{SMBH}} (M_{\odot})$	e	E_0 (rad)
1	0.2	0.2	0.3	9.8×10^{-3}	4×10^6	0.1435	1.14
2	0.05	0.9	0.3	6×10^{-3}	4×10^6	0.5	2.974
3	1.2	π	0.3	9.8×10^{-3}	4×10^7	0.0	0.33

Table 4.2: The Keplerian elements for our example systems, as depicted in Figure 4.3.

Using equation (4.30), which describes the delay for emissions from the far-side of the mass, we derive a general Shapiro delay (to first order in the Schwarzschild radius) for emissions from anywhere around the central mass in Appendix B. Defining $\epsilon \equiv r_0^2/z_p^2$, the far-side (FS) Shapiro delay can be written:

$$\Delta t_{\text{FS}} \approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + (1 + \epsilon)^{1/2}}{z_p \epsilon} \right] - \frac{1}{2(1 + \epsilon)^{1/2}} \right), \quad (4.31)$$

and the near-side (NS) Shapiro delay is found to be

$$\Delta t_{\text{NS}} \approx \frac{2GM}{c^3} \left(\ln \left[\frac{(1 + \epsilon)^{3/2} - (1 + \epsilon)}{z_p \epsilon^2} \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right). \quad (4.32)$$

We have that

$$\Delta t_{\text{Shapiro}} = \begin{cases} \Delta t_{\text{FS}} & \text{for } \bar{z}_{\text{obs}} \geq 0, \\ \Delta t_{\text{NS}} & \text{for } \bar{z}_{\text{obs}} < 0. \end{cases} \quad (4.33)$$

Section 6.3 discusses important shortcomings of this approximate formulation, which is appropriate on solar system scales, but which must be addressed in modelling the time-delays on galactic (SMBH) scales.

4.2.3 Summary and examples

We have computed the most obvious and crucial contributions to the time-delays. Some approximations, omissions and simplifications have been made, which are not necessarily valid, and there are some possible oversights that would be unacceptable in a real-world data analysis setting; the relevant particulars are given in Section 6.3 however, since that section deals in detail with shortcomings in the various modelling stages and lays out the scope of validity in which the results of the following chapters are derived.

Some example orbits, which we use quite extensively in our tests of, for example, parameter estimation and defining time-delay approximants, are defined in Table 4.2

and are depicted in Figure 4.3. The various contributions to the time-delays for these example orbits are shown in Figure 4.4. We finally note that our modified signal simply derives from evaluating the waveform at a delayed time given by the *total time-delay*, $t_{\text{tot}}(t)$, written as

$$t_{\text{tot}}(t) = t - \Delta t_{\text{Einstein}} - \Delta t_{\text{Shapiro}} - \Delta t_{\text{Roemer}}, \quad (4.34)$$

where the various terms are defined in equations (4.29), (4.33) and (4.24). One can see the terms are simply additive since each delay function we have derived returns some extra delay that is dependent on the (relative) position of the binary, which is always parameterised by the coordinate time t , so the order in which they are applied does not matter. The waveform with modelled time-delays is given by

$$h_{\text{tot}}(t) = h(t_{\text{tot}}(t)). \quad (4.35)$$

Finally, note that since we are modelling a source undergoing time-delay modification and that there exists a mass-redshift degeneracy, we require a ‘reference chirp mass’. At any time, we can (in principle) infer an instantaneous chirp mass: what the chirp mass *appears to be* if one neglects line-of-sight velocity (and other time-delays). Since the time-delays change in time, one must be careful in handling this, especially when working with long duration signals and spanning multiple frequency bands. Consider a signal in LISA from a binary accelerating away from the observer. The instantaneous




Figure 4.3: Depiction of the example orbits with Keplerian elements given in Table 4.2. Each orbit has an additional rotation around γ by 0.3 radians, affecting the top down view but not the line-of-sight distance or the time dilations. Note we take the distance parameter d now as the distance to the SMBH, not the binary (the change in amplitude will be negligible and is ignored). The offset from that distance is denoted Δd .

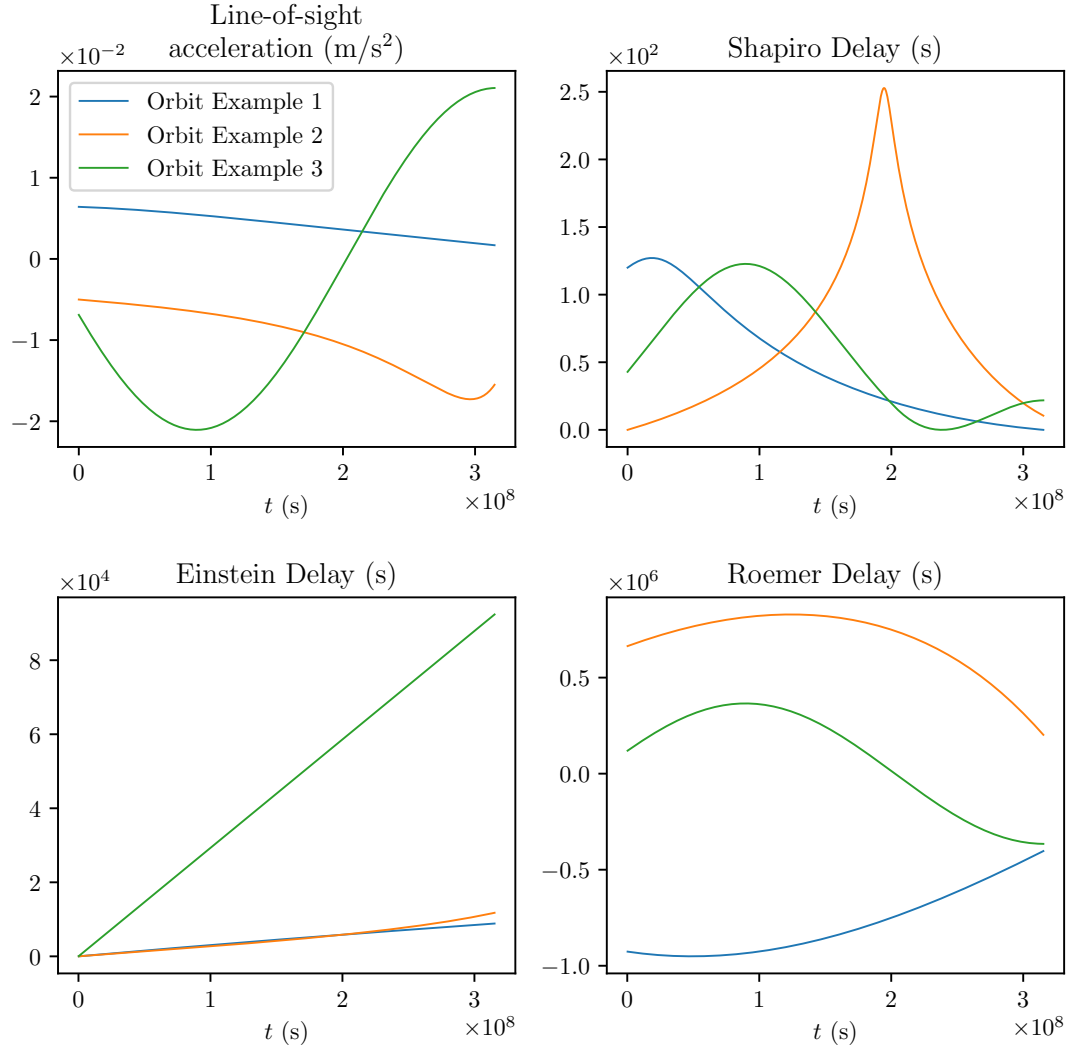


Figure 4.4: The line-of-sight acceleration along with the various contributions to the time-delays that constitute the waveform modification, for the example orbits with Kepler parameters given in Table 4.2 and depictions in Figure 4.3. The total time offset plot has been omitted here since the Roemer delay dominates the delays; the total time offsets are not visually distinguishable from the Roemer delay plot. One can however see that the other time-delays are highly significant waveform modifiers for these particular example orbits. The various time-delays are in general highly non-trivial. Note that the Einstein delay is a cumulative rather than instantaneous delay, since the BHB’s evolution is effectively permanently ‘slowed’ to distant observers, rather than instantaneously shifted, as effected by the Shapiro and Roemer delays.

chirp mass will appear to increase if acceleration is not modelled. For continuing acceleration, the chirp mass continues to appear to increase until the merger occurs in LIGO. Thus searching for a fixed source frame chirp mass signal in LISA and the *same* source frame chirp mass in LIGO will be incorrectly modelled; the time-delay modelling must be continued consistently across the whole signal. We choose to use the instantaneous chirp mass at coalescence (in LIGO) as the reference chirp mass. One must then compute the redshift (from all time-delays) at t_c and update the source frame chirp mass to reflect this for each waveform template. Note then that in PE, we are searching over the *reference chirp mass* parameter, rather than source frame chirp mass.

4.2.4 Reduced dimensionality time-delay approximants

We will see in Section 8.4 that the posteriors defined by the six Keplerian elements can be finely structured, generally consisting of a thin mode(s) containing the probability (the probability is tightly ‘condensed’ around a five-dimensional hypersurface), and thus can be difficult for samplers to explore. However, we expect that *measurable* time-delays can be accurately described by fewer degrees of freedom (than these six Keplerian elements) in large regions of the orbital element parameter space, which we will denote by \mathcal{K} , particularly when the BHB is (for example) far from the SMBH, when acceleration and the other time-delays are approximately constant or linear. In such regions, one could say there is effectively a new space of detectable parameters of a lower dimensionality: for example, for two precisely correlated parameters there is only one degree of freedom, or when a parameter becomes so weakly coupled to the signal it may be possible to simply omit it. More technically, we could define maps from \mathcal{K} to new spaces of detectable parameters, \mathcal{K}'_p , say, such that $p = \dim(\mathcal{K}'_p) < 6$. However, for brevity, we shan’t be rigorous with definitions of maps, coordinates and spaces in this section; our interest in these regions is rudimentary and without opportunity for confusion.

To aid in the parameter estimation and sampling part of the analysis process then, we can consider dividing the parameter space into regions in which the effects of some of the parameters are negligibly coupled to the signal, or become completely degenerate with another parameter, and which thus have *effectively fewer* dimensions. There are many ways one could go about approaching such a division. Here, we shall perform a simple preliminary investigation, providing rough outlines of the boundaries between regions.

Note that if and when one can describe the time-delays using a reduced set of parameters, it is often the case that there is little information on the original six Keplerian element parameters. Since they do not convey much information but are nevertheless required to be included in the model, they can generally be considered as nuisance parameters (unless further information is known about the source from other observations that already constrain the priors on \mathcal{K}).

To do this, it will be helpful to declare a single quantity that will be responsible for the various time-delays, then we can find the approximation of *that* quantity that may be defined by some model of fewer parameters. Any constant velocity of the BHB is entirely undetectable, since there is an exact degeneracy between instantaneous line-of-sight velocity and chirp mass (a line-of-sight velocity also changes the distance over time of course, which affects the amplitude, but this is negligible). A line-of-sight acceleration, however, *is* measurable in principle. For this reason, it will be useful to suppose the collected time-delays all derive from a fictitious line-of-sight acceleration, which we will call the ‘ascribed acceleration’. For total time-delay $\Delta t_{\text{tot}}(t)$, the ascribed line-of-sight distance \check{d} is simply

$$\check{d} = c\Delta t_{\text{tot}} \quad (4.36)$$

so that the ascribed acceleration \check{a} (we use the symbol \check{a} for acceleration, rather than a , which is reserved for the semi-major axis, and the check $\check{}$ to denote ‘ascribed’) is

$$\check{a} \equiv c \frac{d^2}{dt^2} \Delta t_{\text{tot}}. \quad (4.37)$$

Denote the best fit polynomial approximation of order n , of the function \check{a} over the signal by $P_n^{\text{bf}}(\check{a})$:

$$P_n^{\text{bf}}(\check{a}) \equiv \sum_{k=0}^n b_k t^k, \quad (4.38)$$

where the b_k are constant coefficients that define the best fit (by least squares) curve. Now we can evaluate the difference, or error between the ascribed acceleration and its best n^{th} degree polynomial approximation at a given point in the parameter space; if the difference is small, the approximation is good (we will see how small the difference is required to be, in e.g. equation (4.43), below). We want to know this difference over the duration of a given signal rather than at an instant, so define the *integrated (ascribed) acceleration difference (IAD)* by

$$I_n = \int_{\text{signal}} dt |\check{a} - P_n^{\text{bf}}(\check{a})|. \quad (4.39)$$

Using standard polynomials is a sensible choice; another set of basis functions could of course have been chosen with which to construct the time-delay. Since we are interested in how the waveform modifiers might affect the PE results for the ordinary LIGO parameters, one could choose as acceleration basis functions (the derivatives of) the various terms in the PN expansion of the phase. In that case, altering the amplitude of a particular PN term will alter the apparent information on the corresponding CBC parameter ordinarily encoded in that PN term. However, there is no reason to expect (indeed it is extremely unlikely) that external waveform modifiers uncoupled from the generation of the waveform at source (such as the time-delays we are considering) will have precisely the same effect as, say, a higher total spin over the signal duration. The standard polynomials, on the other hand, are acquired from Taylor expansions, and thus for minor deviations, particularly those which are constant over the duration of the signal, may be very well approximated by low-order, ordinary polynomials. We are therefore reasonably well justified in using what is an effectively arbitrary but simple and ubiquitous basis.

We can now consider asking which points in \mathcal{K} have ascribed accelerations that are indistinguishable from polynomial accelerations of degree n . Let us take a look at the IADs for the example orbits in Figure 4.3, taking the semi-major axis as a free parameter; the results are shown in Figure 4.5 for signal lengths of 10 years. As one increases the semi-major axis and so the orbital period, the 10 year section of the orbit covered by the signal becomes increasingly more similar to a straight line trajectory. We are interested in when the ascribed acceleration becomes indistinguishable from a polynomial. We will therefore compute the points at which best fit polynomial accelerations of degree n become (in)distinguishable from the ascribed accelerations, and set these as the region boundary points.

Before one can say whether the polynomial approximants are indistinguishable from the true modification, it must first be understood how one can go about saying when some waveform modification produces a measurable effect. It is reasonable to say that an effect is measurable if it shifts the primary posterior mode by a ‘significant amount’: of order of one standard deviation of any of the posterior’s marginal distributions; such a shift is clearly detectable. Since the chirp mass is precisely correlated with line-of-sight velocity, the chirp mass marginal will provide good insight into when the acceleration is measurable. The chirp mass appears at leading PN order in the waveform’s phase. The

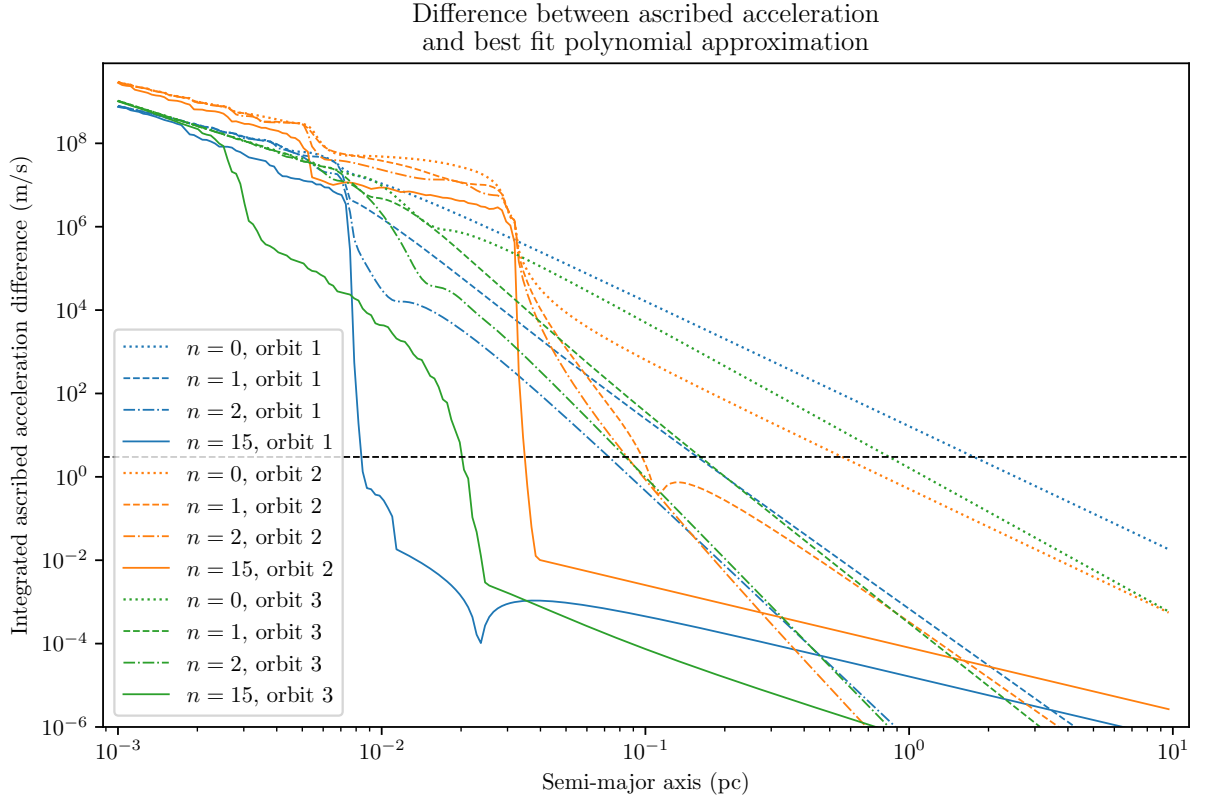


Figure 4.5: Total integrated difference between true time-delay and best fit polynomial (of order n) approximations to the delay for fixed duration signals (10 years) as a function of ‘size of orbit’ (parameterised by semi-major axis). The black dashed line at 3 m s^{-1} delimits the point at which the difference in the acceleration models is actually detectable as found in equation (4.43): at and below this line, the IAD is small enough that the orbit is indistinguishable from the polynomial approximation, thus this gives impetus for finding regions of the Kepler parameter space where we can approximate the time-delay model. One can find out the degree of polynomial required to describe the (measurable) time-delays at a given semi-major axis by locating the curve intersecting the black line immediately to the left of the given semi-major axis. We can see here that for our example orbits with semi-major axes approximately greater than 2 pc, a constant ($n = 0$) acceleration is sufficient, for orbits with semi-major axes approximately greater than 0.2 pc, one requires a linear ($n = 1$) acceleration, and for 0.1 pc, one requires a quadratic ($n = 2$) acceleration. Beyond this however, even polynomials of degree 15 (included only for the sake of interest) are not able to significantly reduce the semi-major axis: clearly the original six Keplerian elements should be used for such systems.

phase can be tracked very precisely and thus the likelihood can drop significantly for slight mismatches in the phase. For a signal consisting of \mathcal{N}_c cycles, the mismatch in the chirp mass, $\Delta\mathcal{M}_c$, between the true value and the template value gives a rough, order of magnitude estimate of the chirp mass standard deviation and is given by (adapted from Ref. [26])

$$\left| \frac{\Delta\mathcal{M}_c}{\mathcal{M}_c} \right| \sim \frac{1}{\mathcal{N}_c}. \quad (4.40)$$

Suppose that a CBC starts with a velocity of $v_0 = 0$, and after time T_{obs} of constant acceleration has a velocity of v_1 . The initial instantaneous chirp mass is $\mathcal{M}_{c,0}$, and the final instantaneous chirp mass is $\mathcal{M}_{c,1} = \mathcal{M}_{c,0}(1 + z_1)$, where $z_1 = v_1/c$ is the redshift of the signal at the end of the signal. Then if we take the chirp mass mismatch to be $\Delta\mathcal{M}_c = \mathcal{M}_{c,1} - \mathcal{M}_{c,0}$, we have that

$$|v_1| \sim \frac{c}{\mathcal{N}_c}, \quad (4.41)$$

so that, since $v_1 = aT_{\text{obs}}$, where a is the acceleration and T_{obs} is the duration of the signal, we find that the detectable velocity shift is

$$|a| \sim \frac{c}{\mathcal{N}_c T_{\text{obs}}}. \quad (4.42)$$

For stellar origin BHBs in LISA, when $T_{\text{obs}} \sim 10$ years, we ordinarily have up to around ten million cycles in band. Thus one should expect that accelerations of around $|a| \gtrsim 10^{-7} \text{ m s}^{-2}$ are measurable. Anything less will not have any significant effect on the posterior, and anything greater must be taken into account in the model. For a signal duration of 5 years, both T_{obs} and \mathcal{N}_c are halved (if we assume the signal is approximately monochromatic), meaning a 4 \times greater acceleration is required to be detectable. The maximum instantaneous line-of-sight acceleration for a circular orbit (which always occurs when the SMBH, BHB, and observer are collinear) is shown for a range of SMBH masses and semi-major axes in Figure 4.6.

Now we can also say, roughly, that for any difference between ascribed acceleration and a polynomial approximation that is less than this detectable value (when integrated over the signal), one cannot distinguish between the acceleration models and may just as well use the polynomial approximation. Taking the difference between ascribed acceleration and polynomial acceleration to be approximately constant over the signal, distinguishability between the models occurs, for a 10 year signal for stellar origin BHBs

when

$$I_n^{10y} = \int_{\text{signal}(10\text{ yr})} dt 10^{-7} \text{ m s}^{-2} \approx 3 \text{ m s}^{-1}, \quad (4.43)$$

and for a 5 year signal when

$$I_n^{5y} = \int_{\text{signal}(5\text{ yr})} dt 4 \times 10^{-7} \text{ m s}^{-2} \approx 6 \text{ m s}^{-1}. \quad (4.44)$$

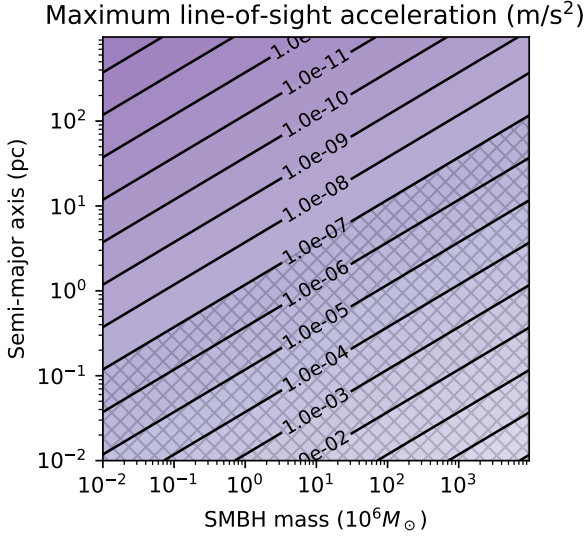


Figure 4.6: Acceleration phase diagram for circular orbits and neglecting relativistic time-delays, where observer, emitter, and central mass all lie on a straight line. The hashed region shows the semi-major axes and SMBH mass pairs that produce accelerations large enough that (for a 10 year signal) one is required to include acceleration in the gravitational waveform model.

see that, for a 10 year signal, time-delays can be taken as constant line-of-sight accelerations when the semi-major axis is around 1 parsec. If the semi-major axis is around 0.1 pc, then at least a linear (first degree polynomial) acceleration is required, and a quadratic acceleration has very little benefit over a linear one; in fact, for polynomial approximations, linear acceleration appears to be very efficient as even up to degree 15 polynomials cannot reduce the IAD much more than linear polynomials. For semi-major axes below 0.1pc, it appears likely that the full set of Keplerian elements are rather generally required to correctly model the time-delays. Note however that only three orbits were

The 10 year IAD of 3 m s^{-1} is shown in Figure 4.5 (the horizontal black dashed line). However, note that, as per the discussion in Ref. [26], the mismatch in chirp mass is not quite as strongly detectable as deduced above, since the spin information is somewhat correlated with the mass and thus, absent of knowledge of the spins, the error in the chirp mass is increased by approximately a factor of 10. Propagating this factor through the above yields a measurable IAD of around $I_n^{10y} = 30 \text{ m s}^{-1}$.

The more crucial purpose of Figure 4.5 is to show *when it is safe* to model the entire set of time-delays as a polynomial acceleration of degree n . Thus one should err on the side of caution and set the acceptable error (i.e., the IAD) to be somewhat lower than that calculated. We can

	Semi-major axis = 0.1 pc		Semi-major axis = 1 pc	
Orbit	A	B	A	B
1	$+1.58 \times 10^{-03}$	-6.57×10^{-13}	$+5.02 \times 10^{-05}$	-6.61×10^{-16}
2	-7.72×10^{-03}	-2.58×10^{-14}	-2.45×10^{-04}	-2.05×10^{-17}
3	$+1.69 \times 10^{-04}$	$+2.04 \times 10^{-13}$	$+1.68 \times 10^{-06}$	$+6.60 \times 10^{-17}$

Table 4.3: The coefficients, as defined in equation (4.45), for best fit linear approximations to the ascribed accelerations of the three example orbits of Figure 4.3 with modified semi-major axes. Units are $[A] = \text{m s}^{-2}$ and $[B] = \text{m s}^{-3}$. In Chapter 8, Figure 8.16, the posteriors for Orbit 1 are generated given the fully modelled time-delay waveform in the data, and analysed using a linear model; the parameter estimated values of A and B from the data are indeed very close to these values of the best fit polynomial coefficients (for both semi-major axes).

studied here and that no simple, general conclusion can yet be drawn about how precisely to divide the parameter space into distinct regions of different dimensionalities and different time-delay model approximants. For reference, note that the best fit linear acceleration coefficients for the example orbits set at semi-major axes of 0.1 pc and 1 pc are given in Table 4.3, for accelerations defined by

$$\ddot{\alpha} \approx P_1^{\text{bf}}(\ddot{\alpha}) \equiv A + Bt. \quad (4.45)$$

It would be instructive to at least derive some approximate boundaries for the division of the parameter space however, in terms of which regions require all six Keplerian elements and which regions are able to approximate the time-delays by various polynomials. Clearly, for BHBS very close to the SMBH where the time-delays are complicated functions, the six Kepler parameters are required. Indeed there will be some region $\mathcal{K}_6 \subset \mathcal{K}$ in which all six Kepler parameters are required to describe the ascribed acceleration. This obviously contains all points for which the semi-major axis $a \sim 0$, and so we can call this the ‘central’ region. As one moves ‘out’ from a point in this region (increasing a) one will eventually reach a point at which fewer parameters are required to describe the time-delays (although of course the dimensionality $\dim(\mathcal{K}) = 6$ of \mathcal{K} itself does not change there). These points shall constitute the boundary, $\partial\mathcal{K}_6$, of \mathcal{K}_6 . For our scheme, we will choose to divide \mathcal{K} into four regions: the central region, \mathcal{K}_6 ; a region describing the ascribed acceleration as a polynomial of degree 1 (linear acceleration), \mathcal{K}_{pbf} ; a region describing the acceleration as a polynomial of degree 0 (constant acceler-

ation), $\mathcal{K}_{p_0^{\text{bf}}}$; and finally a region that does not include acceleration as a parameter (zero acceleration), $\mathcal{K}_{p_{\text{zero}}}$.

One may find these boundaries by thinking of them as isosurfaces of constant IAD and setting the relevant constraints on equation (4.39), like those we found, for example, in equation (4.43) and equation (4.44). That is, if one is interested in the boundary $\partial\mathcal{K}_6$ between the central region \mathcal{K}_6 and the degree 1 polynomial region for a 10 year signal (of stellar origin order chirp mass) $\mathcal{K}_{p_1^{\text{bf}}}$, then this is given by all points $p \in \mathcal{K}$ such that

$$I_1(p) = \int_{\text{signal}} dt |\check{\alpha}(p, t) - P_1^{\text{bf}}(\check{\alpha}(p, t))| = 3 \text{ m s}^{-1}. \quad (4.46)$$

One can write explicitly

$$\partial\mathcal{K}_6 = \{p \in \mathcal{K} : I_1(p) = 3 \text{ m s}^{-1}\},$$

where the region itself is given by

$$\mathcal{K}_6 = \{p \in \mathcal{K} : I_1(p) \geq 3 \text{ m s}^{-1}\}.$$

For completeness, the other regions per our scheme can be written as

$$\begin{aligned} \mathcal{K}_{p_1^{\text{bf}}} &= \{p \in \mathcal{K} : I_0(p) \geq 3 \text{ m s}^{-1}, I_1(p) < 3 \text{ m s}^{-1}\} \\ \mathcal{K}_{p_0^{\text{bf}}} &= \{p \in \mathcal{K} : I_{\text{zero}}(p) \geq 3 \text{ m s}^{-1}, I_0(p) < 3 \text{ m s}^{-1}\} \\ \mathcal{K}_{p_{\text{zero}}^{\text{bf}}} &= \{p \in \mathcal{K} : I_{\text{zero}}(p) < 3 \text{ m s}^{-1}\}, \end{aligned}$$

and where we have that

$$\mathcal{K} = \mathcal{K}_6 \cup \mathcal{K}_{p_1^{\text{bf}}} \cup \mathcal{K}_{p_0^{\text{bf}}} \cup \mathcal{K}_{p_{\text{zero}}^{\text{bf}}}.$$

These regions are rather difficult to conceptualise (and to compute) given a general eccentric orbit and considering the complete definition of $\check{\alpha}$ (and of $P_1^{\text{bf}}(\check{\alpha})$, for which one can use polynomial regression when considering discrete samples, as was done for producing Figure 4.5). The difficulty is also apparent from the irregularity of the example IAD curves shown in Figure 4.5. We can however make some approximations that will allow us a rough but comprehensible picture of where the boundaries lie between the various regions we have defined. The first approximation is to restrict our attention to circular orbits only, then $e = 0$ and we are dealing with a 5-dimensional slice through \mathcal{K}_6 . This allows us to combine E_0 and β into a single parameter; recall that in Section 4.2.2.1 we saw that a 3D Euler rotation made up of a chain of extrinsic 2D rotations is the

same as a chain of intrinsic 2D rotations of the same angles but in reverse order. Thus if we have an extrinsic rotation matrix defined by angles α and β , we may write it as

$$\mathbf{R}_{12}^{\text{ex}}(\alpha, \beta) = \mathbf{R}_1^{\text{ex}}(\beta)\mathbf{R}_2^{\text{ex}}(\alpha) = \mathbf{R}_2^{\text{in}}(\alpha)\mathbf{R}_1^{\text{in}}(\beta). \quad (4.47)$$

However, since we are just dealing with a circular orbit, the initial eccentric anomaly, E_0 , by equation (4.5), reduces to a rotation around the orbital axis by an initial reference angle and can be defined by matrix $\mathbf{R}_1^{\text{in}}(E_0)$, acting in precisely the same way as the β rotation. It can therefore be absorbed into β ; since the matrices are left-acting rotation operators, then acting with matrix $\mathbf{R}_{12}^{\text{ex}}(\alpha, \beta)$ on the left of $\mathbf{R}_1^{\text{in}}(E_0)$ gives

$$\begin{aligned} \mathbf{R}_{12}^{\text{ex}}(\alpha, \beta)\mathbf{R}_1^{\text{in}}(E_0) &= \mathbf{R}_2^{\text{in}}(\alpha)\mathbf{R}_1^{\text{in}}(\beta)\mathbf{R}_1^{\text{in}}(E_0) \\ &= \mathbf{R}_2^{\text{in}}(\alpha)\mathbf{R}_1^{\text{in}}(\beta + E_0) \\ &= \mathbf{R}_{12}^{\text{ex}}(\alpha, \beta + E_0). \end{aligned} \quad (4.48)$$

The parameter space has thus reduced from the 6-parameter set $(\alpha, \beta, a, M_{\text{SMBH}}, e, E_0)$ to the 4-parameter set $(\alpha, a, M_{\text{SMBH}}, E'_0)$, where we have defined $E'_0 \equiv \beta + E_0$.

The second approximation (which is not well substantiated) is to declare the Einstein and Shapiro delays to be negligible and ignore them. This drastically simplifies the definition of $\tilde{\alpha}$ at the cost becoming less accurate. This approach to finding $\partial\mathcal{K}_6$ may only be applied safely in those regions where both $\partial_t^2\Delta t_{\text{Shap}}$ and $\partial_t^2\Delta t_{\text{Einstein}}$ are also well approximated as linear functions (linear because I_1 is defined with respect to a linear polynomial). When the orbit is circular, $\Delta t_{\text{Einstein}}$ is precisely a linear function, so that $\partial_t^2\Delta t_{\text{Einstein}} = 0$, but one only needs to glance at Figure 4.4 and equation (B.7) to see that it is not straightforward to locate those points in \mathcal{K} in which the Shapiro delay is approximately linear, even for circular orbits. Alternatively, one could place safe constraints on $\partial\mathcal{K}_6$ by demanding that the total contribution to the IAD from Δt_{Shap} be negligible. The integral of $\partial_t^2\Delta t_{\text{FS}}$ (the ascribed-acceleration of the ‘far-side’ Shapiro delay), for example, is given by:

$$\frac{d}{dt}\Delta t_{\text{FS}} = \frac{2GM}{c^3} \left(\dot{\epsilon} \frac{\frac{3}{2}(1+\epsilon)^{1/2} - 1}{(1+\epsilon)^{3/2} - (1+\epsilon)} - \frac{\dot{z}_p}{z_p} - 2\frac{\dot{\epsilon}}{\epsilon} - \frac{\dot{\epsilon}}{4}(1+\epsilon)^{-3/2} \right), \quad (4.49)$$

where, $\epsilon(t) \equiv r_0(t)^2/z_p(t)^2$, where r_0 is the impact parameter and z_p is the magnitude of the line-of-sight distance from the SMBH to the binary (see Appendix B). With this, the

individual contribution of the Shapiro delay to the IAD can be computed. If this is negligible, then it may be ignored and one can proceed to check that the IAD defined solely by the Roemer delay (i.e., assuming that $\Delta t_{\text{tot}} = \Delta t_{\text{Roemer}}$) meets the relevant criteria, as in equation (4.46), to determine whether $p \in \mathcal{K}_{P_1^{\text{bf}}}$ or $p \in \mathcal{K}_6$.

For this initial study, the GR components of the delay shall simply be neglected, since we only aim to acquire a rough sense of the structure of the boundaries of the various regions. In any case, our time-delay model is incomplete (see Section 6.3). Whilst our estimate shall not be especially accurate, it shall be a useful first step in acquiring a sense of the degree of waveform deformation to expect given the orbital parameters. For an accurate division of the parameter space with fully modelled time-delays, which should be used with the real LISA dataset, a full numerical search for the boundaries defined by equation (4.46) should be performed, taking into account all possible GR time-delay effects and environmental waveform modifiers.

We make very rough approximations of the definition of the IAD in Appendix A in order to be able to produce closed form equations of the isosurfaces defining the boundaries between our defined regions. The boundary $\partial\mathcal{K}_6$ (separating \mathcal{K}_6 and \mathcal{K}_{P_1}) is approximately given by equation (A.8):

$$3 \text{ m s}^{-1} \approx \frac{8\pi^2}{T^2} a \cos(\alpha) \left[\frac{25\pi}{T} \cos(E'_{1/2}) y^2 + \sin(E'_{1/2}) 5y + \frac{T}{2\pi} (\cos(E'_1) - \cos(E'_{1/2})) \right].$$

Then $\partial\mathcal{K}_{P_1}$ (separating \mathcal{K}_{P_1} and \mathcal{K}_{P_0}) is approximately given by equation (A.11):

$$3 \text{ m s}^{-1} \approx \frac{8\pi^2}{T^2} a \cos(\alpha) \left(\sin(E'_{1/2}) \cdot 5y + \frac{T}{2\pi} [\cos(E'_{1/2}) - \cos(E'_0)] \right).$$

Finally, the boundary $\partial\mathcal{K}_{P_0}$ (separating \mathcal{K}_{P_0} and $\mathcal{K}_{P_{\text{zero}}}$) is approximately given by equation (A.12):

$$3 \text{ m s}^{-1} \approx \frac{4\pi}{T} a \cos(\alpha) [\cos(E'_{1/2}) - \cos(E'_1)].$$

The ranges of application of the different acceleration models with boundaries given by the constraint equations (A.8), (A.11) and (A.12) are shown in Figure 4.7. Of course if one uses a different set of acceleration models (rather than standard polynomials) then these regions would be defined by different isosurface constraint equations. The results of employing these various time-delay remodellings in parameter estimation are shown and discussed in Section 8.5.

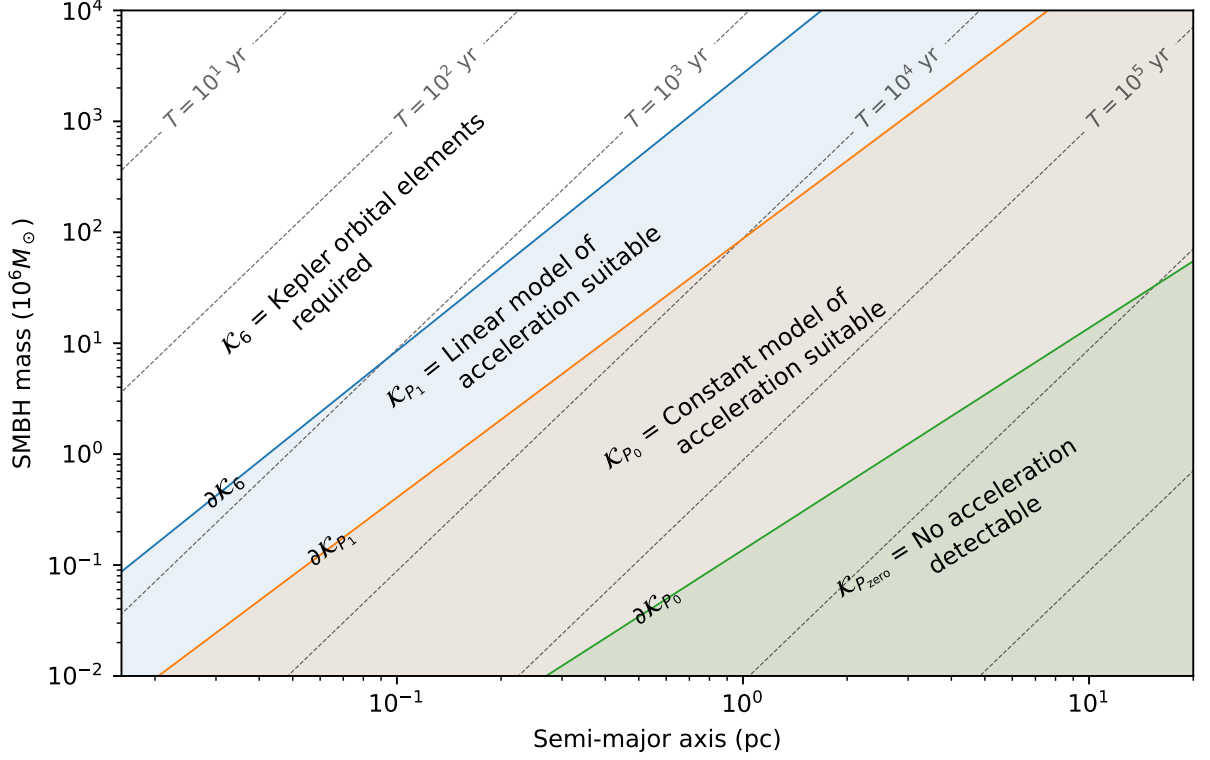


Figure 4.7: A slice through the *reduced* Keplerian element parameter space \mathcal{K} (circular orbits with E_0 and β combined), with $(e, \alpha, E'_{1/2}) = (0, \pi/4, \pi/4)$, showing the various regions we have defined: \mathcal{K}_6 , \mathcal{K}_{P_1} , \mathcal{K}_{P_0} and $\mathcal{K}_{P_{\text{zero}}}$, and their approximate boundaries, which accurately describe the line-of-sight accelerations (as far as can be detected) using 6, 2, 1 and 0 parameter models, respectively. The regions are defined by very rough approximations (as detailed in Appendix A) to the time-delays, however, with our fiducial orbital system in Chapter 8, when using the fully-modelled time-delays we find very good agreement with this approximation. This lends support to the notion of time-delay models that incrementally approximate the fully modelled time-delay, and which are applicable at different regions of the Kepler parameter space. The lines of constant T show that for the same period, the Kepler parameters are required/better detected for more massive SMBHs.

Quadratic Gravity Theories

In modern field theory, the EFEs can be obtained by applying the stationary action principle to the *Einstein-Hilbert action*, given by

$$\mathcal{S}_{\text{EH}} = \int_{\mathcal{M}} d^4x \sqrt{-g} \left(\frac{1}{2\kappa} R + \mathcal{L}_m \right), \quad (5.1)$$

where g is the determinant of the metric $g_{\mu\nu}$, \mathcal{M} is the spacetime manifold, \mathcal{L}_m is a suitable matter Lagrangian, and R and κ are, as before, the Ricci scalar and Einstein constant, respectively.

However, there exists significant interest in alternative gravity theories, in large part since this is a potential avenue to explore for the explanation of the phenomena presently attributed primarily to dark matter and dark energy. There is also huge effort directed at the quantisation of gravity, as this is widely believed to be the best approach to unifying the fundamental forces. Any contending theory of gravity must at least comply with the following criteria; that is, a new theory must:

1. be consistent with all existing measurements and tests of gravity,
2. predict deviations where GR has not been tightly constrained,
3. admit a well-posed initial value problem (the physics should make sense as a classical field theory; for example, evolving some initial data should not lead to unphysical solutions).

Additionally, it is desirable for a theory to be amenable to testing and hence to be solvable (at least perturbatively), for example, we would like to be able to derive its GW solutions.

There are few known theories that fulfil these criteria, but two in particular are well-studied and hold some promise. These are the so-called *Einstein-dilaton Gauss-Bonnet* gravity and *Dynamical Chern-Simons* gravity. These theories both feature scalar fields coupled to gravity. Their field equations are acquired by variation of their respective actions.

5.1 Einstein-dilaton Gauss-Bonnet

The so-called *Gauss-Bonnet (GB)* gravities are theories of gravity which are of interest due to their simple algebraic and geometric properties and definitions. From a classical point of view, they are the only theories with field equations of at most second derivatives in the metric [125]. This is because the particular combination of terms leads to a cancellation of higher-order derivatives in their equations of motion. A corollary of the *Chern-Gauss-Bonnet* theorem is that, on a manifold \mathcal{M} of dimension 4, one has the following identity

$$\chi(\mathcal{M}) = \frac{1}{8\pi^2} \int_{\mathcal{M}} d^4x \sqrt{-g} \mathcal{G} ,$$

[126] where $\chi(\mathcal{M})$ is the topological invariant known as the *Euler-Poincaré characteristic* of \mathcal{M} and the GB term (in 4 dimensions)

$$\mathcal{G} \equiv R^2 - 4R_{\mu\nu}R^{\mu\nu} + R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} . \quad (5.2)$$

There is an obvious appeal in considering somehow combining the Euler-Poincaré characteristic with the gravitational action. However, in 4 dimensions the GB term is a total derivative, so that by Stokes' theorem its integral may be rewritten as a surface term, then (for manifolds without boundary) its contribution to the action identically vanishes. The manifolds are indeed usually considered to be without boundary or are otherwise supplemented with a counter term to cancel any boundary terms that are introduced (for example, the *Gibbons-Hawking-York* term is added to the action for deriving GR from a manifold with boundary [127]), since variation of the action with respect to the metric induces boundary terms that are not well-defined and would cause the stationary action principle to break down [127]. Therefore, the GB gravities are considered higher dimensional theories only, since only there does the addition of the GB term to the action result in some change to the equations of motion.

In order to acquire extra dynamics beyond GR in the 4-dimensional case, one can couple the GB term to a new scalar field $\vartheta_{\text{EdGB}}(x)$ (called the *dilaton*). The resulting theory is called *Einstein-dilaton Gauss-Bonnet (EdGB) gravity*, and its action is written [128]

$$\mathcal{S}_{\text{EdGB}} = \int_{\mathcal{M}} d^4x \sqrt{-g} \left(\bar{\kappa} R + \alpha_{\text{EdGB}} \vartheta_{\text{EdGB}} \mathcal{G} - \frac{1}{2} (\nabla \vartheta_{\text{EdGB}})^2 \right), \quad (5.3)$$

where α_{EdGB} is a coupling constant and where it is implied that this applies to the vacuum region only, such that $\mathcal{L}_m = 0$. To keep things slightly tidier, the constant factor of R in the action is often written in the literature as $\bar{\kappa} \equiv \frac{1}{2\kappa} = (16\pi)^{-1}$, where the κ here is the *dimensionless* Einstein constant.

5.2 Dynamical Chern-Simons

The so-called *Chern-Simons* theories are powerful 3-dimensional topological quantum gauge theories that have found applications in many areas of physics. They allow difficult problems of quantum theories to be expressed as purely topological questions that can be explicitly answered [129]. In parallel, it is known that (2+1)-dimensional gravity is always locally flat and has no propagating degrees of freedom (or GWs) i.e., the curvature is constant [130]. It is therefore a topological theory and indeed it has a Chern-Simons formulation, at least perturbatively (Witten remarked that there may be technical reasons for the duality to fail in the non-perturbative setting [131]). Nevertheless, the idea of being able to formulate gravity as a topological theory is still very enticing, since placing questions of geometry on an equal topological footing to those of quantum theories may allow a formulation of quantum gravity, and has thus led to significant efforts to formulate a working example.

The Chern-Simons theories are defined via the stationary action principle, where the action is the integral of the *Chern-Simons 3-form*. One of the earliest attempts to define a theory with local gravitational degrees of freedom as, or in regard to, a topological theory is the 3-dimensional Chern-Simons gravity theory [132]. This is technically misnamed however, since it is a modification of the Einstein-Hilbert action in 3-dimensions by *addition* of the Chern-Simons 3-form, and not the pure topological Chern-Simons theory. Its extension to 4 dimensions was presented around 20 years later in 2003 [133]. Although these modified GR theories are not strictly Chern-Simons, they turn out to have surpris-

ing features that derive from and connect various physical theories (see [134]) which motivate their continued study.

To make these theories dynamical Chern-Simons (dCS), the coupling constant for the so-called *Pontryagin* term, *RR , that is added to the Einstein-Hilbert action for the 4-dimensional theory [134] is promoted to a scalar field $\alpha_{\text{dCS}} \rightarrow \alpha_{\text{dCS}}\vartheta_{\text{dCS}}(x)$, where

$${}^*RR \equiv R_{\mu\nu\rho\sigma} {}^*R^{\mu\nu\rho\sigma},$$

and where ${}^*R^{\mu\nu\rho\sigma} = \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}R_{\alpha\beta}{}^{\rho\sigma}$ is the *dual* of the Riemann tensor and $\epsilon^{\mu\nu\alpha\beta}$ is the (totally-antisymmetric) Levi-Civita symbol. The 4-dimensional dCS gravity action then reads [128]

$$\mathcal{S}_{\text{dCS}} = \int_{\mathcal{M}} d^4x \sqrt{-g} \left(\bar{\kappa}R + \frac{1}{4}\alpha_{\text{dCS}}\vartheta_{\text{dCS}} {}^*RR - \frac{1}{2}(\nabla\vartheta_{\text{dCS}})^2 \right). \quad (5.4)$$

5.3 Small deviations From GR

In this section we briefly introduce the formalism for deriving solutions of the equations of motion for the modified gravity theories assuming small deviations from GR. We provide a summary of the main results required to define modifications to GWs that arise from non-vanishing couplings of the relevant scalar field to gravity, defined in equations (5.3) and (5.4). The full details of the derivations are lengthy and can be found in Refs. [135, 136] and references therein. The modified GWs themselves shall be discussed in the following section.

One can compare the actions in equations (5.3) and (5.4) to the most general (effective) quadratic gravity action given in Equation (1) of [136]

$$\begin{aligned} \mathcal{S}_{(\text{eff})} \equiv \int_{\mathcal{M}} d^4x \sqrt{-g} & \left(\bar{\kappa}R + \alpha_1 f_1(\vartheta)R^2 + \alpha_2 f_2(\vartheta)R_{\mu\nu}R^{\mu\nu} + \alpha_3 f_3(\vartheta)R_{\mu\nu\delta\sigma}R^{\mu\nu\delta\sigma} \right. \\ & \left. + \alpha_4 f_4(\vartheta)R_{\mu\nu\delta\sigma} {}^*R^{\mu\nu\delta\sigma} - \frac{\nu}{2} \left[(\nabla\vartheta)^2 + 2V(\vartheta) \right] + \mathcal{L}_m \right), \end{aligned} \quad (5.5)$$

to see that the EdGB action, from equations (5.2) and (5.3), is given by setting

$$\begin{aligned} \alpha_1 f_1 &= -\frac{1}{4}\alpha_2 f_2 = \alpha_3 f_3 = \alpha_{\text{EdGB}}\vartheta_{\text{EdGB}}, \\ \alpha_4 f_4 &= 0, \end{aligned}$$

and comparing to equation (5.4), the dCS action is given by setting

$$\begin{aligned}\alpha_1 f_1 &= \alpha_2 f_2 = \alpha_3 f_3 = 0, \\ \alpha_4 f_4 &= \frac{1}{4} \alpha_{\text{dCS}} \vartheta_{\text{dCS}},\end{aligned}$$

where for both theories, $\nu = 1$ and $V = 0$.

Then it is shown in [136] that the metric deformation one acquires is actually proportional to $\xi_i \equiv \alpha_i^2 / \bar{\kappa}$, rather than the α_i one might initially have surmised by observing the definition of the actions, since it is this quantity which scales the solutions. On dimensional grounds, since a perturbation parameter should be dimensionless, the proper dimensionless perturbation parameter must be appropriately rescaled, which is achieved using the mass of the system (the only available scale) so as to be finally defined by [136]

$$\zeta_i \equiv \xi_i / m^4 = \alpha_i^2 / (\bar{\kappa} m^4) = 16\pi \alpha_i^2 / m^4, \quad (5.6)$$

where m is the total mass of the system.

The quadratic gravity metrics can be written as a decomposition into a background metric that satisfies the EFEs, plus a perturbation due to the contribution from the quadratic gravity terms. The metric is written as an expansion [135] in powers of the perturbation parameter, ζ , though it should be emphasized that the parameter itself is not necessarily required to be small; since each ζ^n multiplies a power of the metric perturbation $(h_{\mu\nu})^n$ in the equations of motion, higher orders in ζ can be ignored when one assumes the relevant scalar field is weakly coupled to gravity and higher order metric perturbations are negligible. The parameter then also acts as a book-keeping parameter since it keeps track of the order of the metric perturbation. For this reason, the small coupling approximation can be obtained by demanding that $\zeta < 1$ [128].

5.4 Schematic frequency domain waveform

As in [128], we can write the GW signal in the frequency domain as

$$\tilde{h}(f) = A(f) \exp(i\phi(f)), \quad (5.7)$$

where the modified phase ϕ can be written as

$$\phi(f) = \phi_{\text{GR}}(f) + \beta u^b, \quad (5.8)$$

where ϕ_{GR} is the GR phase, and the second term contains contributions from the effects of modified gravity theories, where $u \equiv (\pi \mathcal{M} f)^{1/3}$. Here, β is the amplitude coefficient, a function of the component masses and spins [128]. The exponent b is related to the PN order p_{PN} by $p_{\text{PN}} = (b + 5)/2$ [110], and for the different modified gravity theories, takes the values

$$b_{\text{dCS}} = -1, \quad \text{and} \quad b_{\text{EdGB}} = -7. \quad (5.9)$$

The leading order dCS term comes in at 2PN and the leading order EdGB term comes in at -1 PN. The amplitude coefficients for the two theories, β_{dCS} and β_{EdGB} , can be found in [128] and are reproduced here to see the dependence of the modified gravity parameters on the masses and spins. We have

$$\beta_{\text{dCS}} = \frac{\zeta_{\text{dCS}}}{\eta^{14/5} m^2} \left[\frac{15075}{114688} \left(m_2^2 \chi_1^2 - \frac{350}{201} m_1 m_2 \chi_1 \chi_2 + m_1^2 \chi_2^2 \right) - \frac{5}{8192} (m_1 s_2^{\text{dCS}} - m_2 s_1^{\text{dCS}})^2 \right], \quad (5.10)$$

and

$$\beta_{\text{EdGB}} = -\frac{5}{7168} \frac{\zeta_{\text{EdGB}} (m_1^2 s_2^{\text{EdGB}} - m_2^2 s_1^{\text{EdGB}})^2}{\eta^{18/5} m^4}, \quad (5.11)$$

where $m = m_1 + m_2$ is the total mass and the *dimensionless spin parameters* are given by

$$\chi_i = \frac{\vec{S}_i \cdot \hat{L}}{m_i^2}. \quad (5.12)$$

Here, the \vec{S}_i is the (dimensionless) spin angular momentum vector of the i^{th} black-hole and \hat{L} is the orbital angular momentum of the binary. We also have the *charges* $s_i^{\text{dCS,EdGB}}$ of the respective scalar fields [137] sourced by the i^{th} black-hole, which are given by

$$s_i^{\text{dCS}} = \frac{2 + 2\chi_i^4 - 2\varsigma_i - \chi_i^2(3 - 2\varsigma_i)}{2\chi_i^3}, \quad (5.13)$$

and

$$s_i^{\text{EdGB}} = \frac{2\varsigma_i(1 - \varsigma_i)}{\chi_i^2}, \quad (5.14)$$

where $\varsigma_i = \sqrt{1 - \chi_i^2}$. In our analysis, these parameters are somewhat simplified by our decision to study a model with a single spin component, i.e. setting $\chi_1 = \chi_2$, from which we can see that $\varsigma_1 = \varsigma_2$ and $s_1^{\text{dCS,EdGB}} = s_2^{\text{dCS,EdGB}}$. From equations (5.10) and (5.13), one finds that there is no extra dCS modification to the GR waveform when BHs have zero spin. For EdGB gravity, one can see from equations (5.11) and (5.14) that there is no




Figure 5.1: The modified gravity β -parameter phase diagram for EdGB and dCS gravity. These plots are derived from equations (5.10) and (5.11) given $\zeta_{\text{EdGB,dCS}} = 1$. Note that the modified gravity parameters are independent of mass scaling. The EdGB parameter is largely independent of effective (aligned) spin, whereas the dCS gravity dependence on effective (aligned) spin strongly depends on the mass ratio. If data contains a GR signal (i.e., with $\zeta = 0$), then $\beta = 0$. But $\beta = 0$ can also be achieved in EdGB gravity if $q = 1$, and in dCS gravity if $\chi_1 = \chi_2 = 0$. In Chapter 7, we will see that this can be a significant source of bias in parameter estimation for EdGB gravity.

EdGB modification to the GR waveform when the spins and masses are precisely equal. Ideal sources of GWs for tests of these modified gravity theories are therefore generally black-holes with large spins forming binaries with highly asymmetric masses (see Figure 5.1). Due to how the amplitude of the modification scales with the mass as in equation (5.6), it is also more beneficial for constraining these gravity theories (at a given SNR) to observe GWs emitted from low-mass binaries.

5.5 Phase shifted Fourier waveforms in time domain

We now have the modified gravity waveforms in the frequency domain, but we would like to include these modifications in the time domain model to be able to include the effects in a complete time-evolving model, exhibiting both acceleration and modified gravity effects simultaneously. In this section we will use geometric units where $G = c = 1$. Start by writing the frequency-domain waveform model as

$$\tilde{h}(f) = \tilde{h}_{\text{GR}}(f) \cdot e^{i\beta u^b}. \quad (5.15)$$

Then the time-domain model is given by the inverse Fourier-Transform:

$$\begin{aligned}
 h(t) &= \mathcal{F}^{-1} \{ \tilde{h} \} (t) \\
 &= \mathcal{F}^{-1} \{ \tilde{h}_{\text{GR}} \cdot e^{i\beta u^b} \} (t) \\
 &= \langle \mathcal{F}^{-1} \{ \tilde{h}_{\text{GR}} \} * \mathcal{F}^{-1} \{ e^{i\beta u^b} \} \rangle (t) \\
 &= \langle h_{\text{GR}}(t) * \int_{-\infty}^{\infty} df e^{i\beta u^b} e^{2\pi i f t} \rangle (t),
 \end{aligned}$$

where $*$ is the convolution operator and we use the notation $\langle \cdot * \cdot \rangle (t)$, to emphasize the dependencies of the convolved functions and the resultant function, and the third line follows from the convolution theorem. By the definition of convolution

$$h(t) = \int_{-\infty}^{\infty} d\tau h_{\text{GR}}(t - \tau) \int_{-\infty}^{\infty} df e^{i\beta u^b + 2\pi i f \tau}. \quad (5.16)$$

It will be convenient to introduce a constant k , which may take on an arbitrary value, and with which we make the following definitions:

$$B \equiv b/3, \quad \hat{f} \equiv \left(\frac{\beta}{k} \right)^{\frac{1}{B}} \pi \mathcal{M}_c f, \quad \hat{\tau} \equiv \frac{2}{\mathcal{M}_c} \left(\frac{k}{\beta} \right)^{\frac{1}{B}} \tau. \quad (5.17)$$

Then we can make the relevant substitutions to partially rewrite the integrals

$$h(t) = \frac{k}{2\pi} \int_{-\infty}^{\infty} d\hat{\tau} h_{\text{GR}}(t - \tau) \int_{-\infty}^{\infty} d\hat{f} e^{ik(\hat{f}^B + \hat{\tau}\hat{f})}. \quad (5.18)$$

The choice of constant k of course does not alter $h(t)$, however, if $k \gg 1$, then we can use the *stationary phase approximation (SPA)* to evaluate the \hat{f} integral. In order to be able to apply the SPA, there must be an extremum of the phase function, which we write here as

$$G(\hat{f}) \equiv \hat{f}^B + \hat{\tau}\hat{f}, \quad (5.19)$$

with respect to \hat{f} , i.e., its derivative must vanish somewhere. We therefore demand that $B \neq 0, B \neq 1$. Note also that \hat{f}^B can take multiple values, depending on the value of \hat{f} and B : there may exist non-zero imaginary components of \hat{f}^B , so it should be understood that we require the purely real solution be used if it exists (otherwise, the phase modification factor $e^{i\beta u^b}$ also modifies the amplitude).

The points of stationary phase, which we denote $\hat{f}_0 = \hat{f}_0(\hat{\tau})$, occur where

$$\begin{aligned}
 0 &= G'(\hat{f}) = B\hat{f}_0^{B-1} + \hat{\tau} \\
 \Rightarrow \hat{f}_0 &= \left(\frac{-\hat{\tau}}{B} \right)^{\frac{1}{B-1}}, \quad (5.20)
 \end{aligned}$$

and thus if there is no (real) solution for a given B and $\hat{\tau}$, then the phase has no stationary point and the \hat{f} integral evaluates to zero. We must also ensure that validity of SPA continues for $B < 1$ as $\hat{\tau} \rightarrow 0$; in this case too, there is no stationary point of the phase and the integral evaluates to zero.

Now for some B and $\hat{\tau}$, Taylor expanding the phase function around a (stationary) point $\hat{f}_0 = \hat{f}_0(\hat{\tau})$ gives:

$$G(\hat{f}) = G(\hat{f}_0) + G'(\hat{f}_0)(\hat{f} - \hat{f}_0) + \frac{1}{2}G''(\hat{f}_0)(\hat{f} - \hat{f}_0)^2 + \dots \quad (5.21)$$

Recall, however, that the first derivative vanishes there, so we write the integral as

$$h(t) = \frac{k}{2\pi} \int_{-\infty}^{\infty} d\hat{\tau} h_{\text{GR}}(t - \tau) \int_{-\infty}^{\infty} d\hat{f} e^{ik(G(\hat{f}_0) + \frac{1}{2}G''(\hat{f}_0)(\hat{f} - \hat{f}_0)^2 + \dots)} \quad (5.22)$$

It is emphasized again that $\hat{f}_0 = \hat{f}_0(\hat{\tau})$, so that as one sweeps through the $\hat{\tau}$ integral, the stationary point of the phase (if one exists) is ‘picked out’ of the real line in each \hat{f} integral. Now as $k \rightarrow \infty$, we can drop higher order terms in the phase and write this as

$$h(t) = \frac{k}{2\pi} \int_{-\infty}^{\infty} d\hat{\tau} h_{\text{GR}}(t - \tau) e^{ikG(\hat{f}_0)} \int_{-\infty}^{\infty} d\hat{f} e^{i\frac{k}{2}G''(\hat{f}_0)(\hat{f} - \hat{f}_0)^2} \quad (5.23)$$

If we define $s \equiv \sqrt{\frac{k}{2}G''(\hat{f}_0(\hat{\tau}))(\hat{f} - \hat{f}_0)}$, we can see that this integral is a well-known standard result, $\int_{-\infty}^{\infty} ds e^{is^2} = \sqrt{\pi}e^{i\pi/4}$, so that we have

$$\begin{aligned} \int_{-\infty}^{\infty} d\hat{f} e^{i\frac{k}{2}G''(\hat{f}_0)(\hat{f} - \hat{f}_0)^2} &= \left(\frac{k}{2}G''(\hat{f}_0)\right)^{-1/2} \int_{-\infty}^{\infty} ds e^{is^2} \\ &= \left(\frac{k}{2}B(B-1)\hat{f}_0^{B-2}\right)^{-1/2} \sqrt{\pi}e^{i\pi/4}, \end{aligned}$$

since $G''(\hat{f}) = B(B-1)\hat{f}^{B-2}$. Thus we now have that

$$\begin{aligned} h(t) &= \sqrt{\frac{k}{2\pi}} e^{i\pi/4} \int_{-\infty}^{\infty} d\hat{\tau} h_{\text{GR}}(t - \tau) e^{ikG(\hat{f}_0)} (B(B-1)\hat{f}_0^{B-2})^{-1/2} \\ &= (B(B-1))^{-1/2} \sqrt{\frac{k}{2\pi}} e^{i\pi/4} \int_{-\infty}^{\infty} d\hat{\tau} A_{\text{GR}}(t - \tau) \hat{f}_0^{1-B/2} e^{i\phi_{\text{GR}}(t-\tau) + ikG(\hat{f}_0)}, \end{aligned}$$

since the time-domain waveform can be written in the form

$$h_{\text{GR}}(t) = A_{\text{GR}}(t) e^{i\phi_{\text{GR}}(t)}. \quad (5.24)$$

At this point let us revert back to the original variables with equations (5.17) for better readability. Note that we still write the stationary points with the zero subscript $_0$, i.e., $\hat{f}_0 \rightarrow f_0$. Then equation (5.20) becomes

$$f_0 = (\pi\mathcal{M}_c)^{-1} \left(\frac{-2\tau}{B\beta\mathcal{M}_c} \right)^{\frac{1}{B-1}}. \quad (5.25)$$

We also have that

$$\hat{f}^{1-B/2} d\hat{t} = \frac{2}{\mathcal{M}_c} (\beta k)^{-1/2} (\pi \mathcal{M}_c f)^{1-B/2} d\tau,$$

so that

$$\begin{aligned} h(t) &= (2\pi\beta B(B-1))^{-1/2} \frac{2e^{i\pi/4}}{\mathcal{M}_c} \int_{-\infty}^{\infty} d\tau A_{\text{GR}}(t-\tau) (\pi \mathcal{M}_c f_0)^{1-B/2} e^{i\phi_{\text{GR}}(t-\tau) + i\beta(\pi \mathcal{M}_c f_0)^B + 2\pi i f_0 \tau} \\ &= C \int_{-\infty}^{\infty} d\tau \frac{A_{\text{GR}}(t-\tau)}{f_0^{B/2-1}} \exp\{i(\phi_{\text{GR}}(t-\tau) + \beta(\pi \mathcal{M}_c f_0)^B + 2\pi f_0 \tau)\}, \end{aligned} \quad (5.26)$$

where we defined $C \equiv e^{i\pi/4} \sqrt{\frac{2\pi}{\beta B(B-1)(\pi \mathcal{M}_c)^B}}$. Substituting for f_0 from equation (5.25) yields

$$h(t) = C \int_{-\infty}^{\infty} d\tau \frac{A_{\text{GR}}(t-\tau)}{f_0^{B/2-1}(\tau)} \exp\{iF(t-\tau, \tau)\}, \quad (5.27)$$

where

$$\begin{aligned} F(t-\tau, \tau) &\equiv \phi_{\text{GR}}(t-\tau) + \beta \left(\frac{-2\tau}{B\beta \mathcal{M}_c}\right)^{\frac{B}{B-1}} + \frac{2\tau}{\mathcal{M}_c} \left(\frac{-2\tau}{B\beta \mathcal{M}_c}\right)^{\frac{1}{B-1}} \\ &= \phi_{\text{GR}}(t-\tau) + D\tau^{\frac{B}{B-1}}, \end{aligned}$$

and where $D \equiv 2 \left(\frac{B-1}{B}\right) \left(\frac{-2}{\beta B \mathcal{M}_c^B}\right)^{\frac{1}{B-1}}$.

The integral in equation (5.27) is still a convolution, which, being symmetric in its arguments, allows one to exchange $t-\tau \leftrightarrow \tau$. Doing so above, the waveform becomes

$$h(t) = C \int_{-\infty}^{\infty} d\tau \frac{A_{\text{GR}}(\tau)}{f_0^{B/2-1}(t-\tau)} \exp\{iF(\tau, t-\tau)\}. \quad (5.28)$$

To compute this integral, we can perform a second application of the SPA. We are not afforded the same level of rigour as we had for the first integral, where we introduced the constant k , which could take the phase to arbitrarily high rates of oscillation; if we do this sort of change of variables here with τ , then the amplitude part of the integrand is also able to vary at these arbitrarily high rates, where the SPA could break down. However, the SPA is known to be remarkably accurate in many such situations (despite the SPA having been developed for application in asymptotic limits [138]) including in our particular case of gravitational waveforms from compact binaries (see, e.g., [139]) since it is known that the amplitude is very slowly evolving in comparison to the cycle duration.

The stationary point of this phase for a given t occurs at the τ_0 such that $\dot{F}(\tau, t-\tau)|_{\tau_0} = 0$, where the dot denotes differentiation with respect to τ . We have

$$\dot{F}(\tau, t-\tau) = 2\pi f_{\text{GR}}(\tau) - \frac{BD}{B-1} (t-\tau)^{\frac{1}{B-1}}. \quad (5.29)$$

Thus the stationary point is the $\tau_0 = \tau_0(t)$ satisfying

$$t = \tau_0 + \left(2\pi f_{\text{GR}}(\tau_0) \frac{B-1}{BD} \right)^{B-1}, \quad (5.30)$$

or on substituting for D :

$$\begin{aligned} t &= \tau_0 - \frac{B\beta[\pi\mathcal{M}_c f_{\text{GR}}(\tau_0)]^B}{2\pi f_{\text{GR}}(\tau_0)} \\ &= \tau_0 - \frac{1}{2} B\beta\mathcal{M}_c [\pi\mathcal{M}_c f_{\text{GR}}(\tau_0)]^{B-1}. \end{aligned} \quad (5.31)$$

Now to apply the SPA to equation (5.28), we Taylor expand the phase function around that stationary point, which then reads

$$F(\tau, t - \tau) = F(\tau_0, t - \tau_0) + \dot{F}(\tau_0, t - \tau_0)(\tau - \tau_0) + \frac{1}{2}\ddot{F}(\tau_0, t - \tau_0)(\tau - \tau_0)^2 + \dots, \quad (5.32)$$

where $\dot{F} = 0$ and

$$\ddot{F}(\tau, t - \tau) = 2\pi \dot{f}_{\text{GR}}(\tau) + \frac{2}{B-1} \left(\frac{-2}{\beta B \mathcal{M}_c^B} \right)^{\frac{1}{B-1}} (t - \tau)^{\frac{1}{B-1}-1}. \quad (5.33)$$

Then we can write

$$\begin{aligned} h(t) &= C \int_{-\infty}^{\infty} d\tau \frac{A_{\text{GR}}(\tau_0)}{f_0^{B/2-1}(t - \tau_0)} \exp \left\{ i \left(F(\tau_0, t - \tau_0) + \frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) (\tau - \tau_0)^2 \right) \right\} \\ &= C \frac{A_{\text{GR}}(\tau_0)}{f_0^{B/2-1}(t - \tau_0)} e^{iF(\tau_0, t - \tau_0)} \int_{-\infty}^{\infty} d\tau e^{i \frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) (\tau - \tau_0)^2}, \end{aligned}$$

where we have assumed the amplitude over the small region around the point of stationary phase can be approximated as a constant, and hence can be taken outside the integral.

In the same fashion as the evaluation of the \hat{f} integral, to evaluate the τ integral, define $r \equiv \sqrt{\frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) (\tau - \tau_0)}$, and use the well-known standard result $\int_{-\infty}^{\infty} dr e^{ir^2} = \sqrt{\pi} e^{i\pi/4}$ for

$$\int_{-\infty}^{\infty} d\tau e^{i \frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) (\tau - \tau_0)^2} = \left(\frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) \right)^{-1/2} \sqrt{\pi} e^{i\pi/4}.$$

Then

$$h(t) = C \frac{A_{\text{GR}}(\tau_0)}{f_0^{B/2-1}(t - \tau_0)} \left(\frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) \right)^{-1/2} \sqrt{\pi} e^{i\pi/4} e^{i\phi_{\text{GR}}(\tau_0) + iD(t - \tau_0)^{\frac{B}{B-1}}},$$

where recall $D \equiv 2 \left(\frac{B-1}{B} \right) \left(\frac{-2}{\beta B \mathcal{M}_c^B} \right)^{\frac{1}{B-1}}$ and of course $\tau_0 = \tau_0(t)$ defined in equation (5.31). Working back through all substitutions

$$\begin{aligned}
 h(t) &= e^{i\pi/4} \sqrt{\frac{2\pi}{\beta B(B-1)(\pi \mathcal{M}_c)^B}} A_{\text{GR}}(\tau_0) \left((\pi \mathcal{M}_c)^{-1} \left(\frac{2(\tau_0 - t)}{B\beta \mathcal{M}_c} \right)^{\frac{1}{B-1}} \right)^{1-B/2} \\
 &\quad \times \left(\frac{1}{2} \ddot{F}(\tau_0, t - \tau_0) \right)^{-1/2} \sqrt{\pi} e^{i\pi/4} e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} \\
 &= 2\pi e^{i\pi/2} A_{\text{GR}}(\tau_0) f_{\text{GR}}(\tau_0)^{1-B/2} e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} [\beta B(B-1)(\pi \mathcal{M}_c)^B \ddot{F}(\tau_0, t - \tau_0)]^{-1/2} \\
 &= 2\pi i A_{\text{GR}}(\tau_0) f_{\text{GR}}(\tau_0) e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} [\beta B(B-1)(\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^B]^{-1/2} \\
 &\quad \times \left(2\pi \dot{f}_{\text{GR}}(\tau_0) + \frac{2}{B-1} \left(\frac{-2}{\beta B \mathcal{M}_c^B} \right)^{\frac{1}{B-1}} \left[-\frac{1}{2} \beta B \mathcal{M}_c (\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^{B-1} \right]^{\frac{1}{B-1}-1} \right)^{-1/2} \\
 &= 2\pi i A_{\text{GR}}(\tau_0) f_{\text{GR}}(\tau_0) e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} [\beta B(B-1)(\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^B]^{-1/2} \\
 &\quad \times \left[2\pi \dot{f}_{\text{GR}}(\tau_0) + \frac{2}{B-1} \left(\frac{-2}{\beta B} \right) (\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^{-B} (\pi f_{\text{GR}}(\tau_0))^2 \right]^{-1/2} \\
 &= 2\pi i A_{\text{GR}}(\tau_0) f_{\text{GR}}(\tau_0) e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} \\
 &\quad \times \left[2\pi \beta B(B-1)(\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^B \dot{f}_{\text{GR}}(\tau_0) - 4(\pi f_{\text{GR}}(\tau_0))^2 \right]^{-1/2} \\
 &= 2\pi A_{\text{GR}}(\tau_0) f_{\text{GR}}(\tau_0) e^{i\phi_{\text{GR}}(\tau_0) + iD(t-\tau_0)\frac{B}{B-1}} \\
 &\quad \times \left[4(\pi f_{\text{GR}}(\tau_0))^2 - 2\pi \beta B(B-1)(\pi \mathcal{M}_c f_{\text{GR}}(\tau_0))^B \dot{f}_{\text{GR}}(\tau_0) \right]^{-1/2},
 \end{aligned}$$

where we substituted for C and f_0 in the first line, used equation (5.31) to substitute $t - \tau_0$ and repositioned terms in the second line, substituted for \ddot{F} in the third line, rearranged in the fourth and fifth lines, and took a minus sign outside the square root in the last line giving us another imaginary unit. Thus finally we have that:

$$\begin{aligned}
 h(t) &= \pm A_{\text{GR}}(\tau_0) \left[1 - \frac{\beta}{2\pi} B(B-1)(\pi \mathcal{M}_c)^B f_{\text{GR}}^{B-2}(\tau_0) \dot{f}_{\text{GR}}(\tau_0) \right]^{-1/2} \\
 &\quad \times e^{i\phi_{\text{GR}}(\tau_0) - i\beta(B-1)[\pi \mathcal{M}_c f_{\text{GR}}(\tau_0)]^B},
 \end{aligned} \tag{5.34}$$

where we explicitly write the positive/negative options from the square root, and where $\tau_0(t)$ is defined in equation (5.31). This solution appears to be reasonable since if one

takes $\beta \rightarrow 0$, then we can see from equation (5.31) that $\tau_0 \rightarrow t$ so that

$$h(t)_{\beta \rightarrow 0} \rightarrow \pm A_{\text{GR}}(t) e^{i\phi_{\text{GR}}(t)} = \pm h_{\text{GR}}(t),$$

and the GR waveform is recovered (taking the positive solution). Equation (5.34) is the general solution of the phase and amplitude modifications of a time-domain signal of the form $A^{\text{td}} e^{i\psi_{\text{td}}}$ that arise from the transform of a phase modified Fourier domain signal of the form $A^{\text{fd}} e^{i\psi_{\text{fd}}}$.

Recall we discussed that the amplitude modifications are often ignored, and that we shall do so here in our model too, since phase is by far the dominant measurable property of the waveform. We will therefore write our approximate model as

$$h(t) = A_{\text{GR}}(\tau_0) e^{i\phi_{\text{GR}}(\tau_0) - i\beta(B-1)[\pi\mathcal{M}_c f_{\text{GR}}(\tau_0)]^B}. \quad (5.35)$$

Now we must find $\tau_0 = \tau_0(t)$ for when $\beta \neq 0$. The solution of equation (5.31) (and indeed whether an exact one can be found) depends on the precise form of f_{GR} . Equation (5.31) could well be transcendental, for example, since we have not yet defined f_{GR} , though the case for GW waveforms is generally that the frequency function is written as a PN-expansion, which is a truncated power series (i.e., a polynomial) in powers of $\tau_c^{-1/8}$, where τ_c is the time to coalescence. For example, the waveform model we employ (from [27]) gives the GR frequency (eqn. 6b of [27]) as

$$2\pi f_{\text{GR}} = \omega_{\text{GR}}(t) = \frac{1}{8} \sum_{k=0}^N \hat{\omega}_k \theta^{k+3}, \quad (5.36)$$

where $\theta(t) = [\eta\tau_c/(5M)]^{-1/8}$, $\eta = m_1 m_2 / M^2$ is the *symmetric mass ratio*, $M = m_1 + m_2$ is the total mass, $N/2$ is the ‘PN-order’, $\tau_c = t_c - t$ is the time to coalescence, t_c is the coalescence time and the $\hat{\omega}_k$ are PN term coefficients. Even so, it is still perhaps more efficient to numerically approximate the solution by iteration methods rather than trying to find zeros of the polynomial. Rearranging equation (5.31) as

$$\tau_0 = t + \frac{1}{2} B \beta \mathcal{M}_c [\pi \mathcal{M}_c f_{\text{GR}}(\tau_0)]^{B-1}, \quad (5.37)$$

then defining an iterative approximation of τ_0 as τ_n^\dagger , one has

$$\tau_{n+1}^\dagger \equiv t + \frac{1}{2} B \beta \mathcal{M}_c [\pi \mathcal{M}_c f_{\text{GR}}(\tau_n^\dagger)]^{B-1}, \quad (5.38)$$

so that one approaches equation (5.37) in the limit of large n , that is, $\tau_n^\dagger \rightarrow \tau_0$ as $n \rightarrow \infty$. The iteration procedure of course induces a computational cost. In practice it is often

observed that these sorts of sequences converge quickly, however this must be solved for each value of t present in the signal, and recall that signals in LISA can consist of up to a billion samples.

One alternative approach is to suppose that the GR relation for CBCs [140]

$$\dot{f}_{\text{GR}} \approx \frac{96}{5} \pi^{8/3} \left(\frac{G\mathcal{M}_c}{c^3} \right)^{5/3} f_{\text{GR}}^{11/3}, \quad (5.39)$$

approximately holds and is accurate far from merger. We can integrate this to find $f_{\text{GR}} = f_{\text{GR}}(T)$, where we use the stand-in variable T for the argument to avoid confusion with t in (5.31) which should be treated as a constant there. Then continuing with the use of geometric units, $G = c = 1$, for compatibility with the previous section, we obtain

$$f_{\text{GR}}(T) = \frac{1}{\pi} \left(\frac{5}{256} \frac{1}{t_c - T} \right)^{3/8} \mathcal{M}_c^{-5/8}. \quad (5.40)$$

Substituting into (5.31), we have that

$$\begin{aligned} \tau_0 &= t + \frac{1}{2} B \beta \mathcal{M}_c \left(\frac{5\mathcal{M}_c}{256} \frac{1}{t_c - \tau_0} \right)^{3(B-1)/8} \\ \Rightarrow (\tau_0 - t)^{3(1-B)/8} (\tau_0 - t_c) &= -\frac{1}{2} B \beta \mathcal{M}_c \left(\frac{5\mathcal{M}_c}{256} \right). \end{aligned} \quad (5.41)$$

This can then be solved exactly for $b_{\text{dCS}} = -1$, with solution

$$\tau_0 = \frac{1}{2} \left(\pm p \sqrt{4(t_c - t) + p^2} + 2t - p^2 \right), \quad (5.42)$$

where $p \equiv \frac{1}{2} B \beta \mathcal{M}_c \left(\frac{5\mathcal{M}_c}{256} \right)$. However, there is still no formulaic solution for $b_{\text{EdGB}} = -7$.

Finally, we will pursue the small parameter approximation, taking β to be a small parameter and Taylor expanding in β . In one sense, equations like (5.37) can be viewed as being circularly defined, or as an infinite recurrence relation. The iteration procedure outlined above is essentially a truncation of this recurrence relation, nesting the equation inside itself a finite number of times with some initial guess. In the small β limit, we are essentially taking the first step of the sequence ($\tau_0 \approx \tau_0^\dagger$) by writing

$$\begin{aligned} \tau_0 &= t + \frac{1}{2\pi} B \beta (\pi \mathcal{M}_c)^B \left\{ f_{\text{GR}} \left(t + \frac{1}{2\pi} B \beta (\pi \mathcal{M}_c)^B [f_{\text{GR}}(\tau_0)]^{B-1} \right) \right\}^{B-1} \\ &= t + \frac{1}{2\pi} B \beta (\pi \mathcal{M}_c)^B \left\{ f_{\text{GR}}(t) + \frac{1}{2\pi} B \beta (\pi \mathcal{M}_c)^B [f_{\text{GR}}(\tau_0)]^{B-1} \dot{f}_{\text{GR}}(t) + \mathcal{O}(\beta^2) \right\}^{B-1} \\ &= t + \frac{1}{2} B \beta \mathcal{M}_c (\pi \mathcal{M}_c f_{\text{GR}}(t))^{B-1} \left\{ 1 + \frac{1}{2\pi} B \beta (\pi \mathcal{M}_c)^B [f_{\text{GR}}(\tau_0)]^{B-2} \dot{f}_{\text{GR}}(t) + \mathcal{O}(\beta^2) \right\}^{B-1} \\ &\approx t + \frac{1}{2} B \beta \mathcal{M}_c (\pi \mathcal{M}_c f_{\text{GR}}(t))^{B-1} \left\{ 1 + \frac{1}{2\pi} B(B-1) \beta (\pi \mathcal{M}_c)^B [f_{\text{GR}}(\tau_0)]^{B-2} \dot{f}_{\text{GR}}(t) + \mathcal{O}(\beta^2) \right\} \\ &\approx t + \frac{1}{2} B \beta \mathcal{M}_c [\pi \mathcal{M}_c f_{\text{GR}}(t)]^{B-1} + \mathcal{O}(\beta^2), \end{aligned} \quad (5.43)$$

then dropping the $\mathcal{O}(\beta)$ terms. Now, to $\mathcal{O}(\beta)$ in the phase, we have that

$$\begin{aligned}\phi_{\text{GR}}(\tau_0) &= \phi_{\text{GR}}\left(t + \frac{1}{2}B\beta\mathcal{M}_c[\pi\mathcal{M}_cf_{\text{GR}}(t)]^{B-1}\right) \\ &= \phi_{\text{GR}}(t) + \frac{1}{2}B\beta\mathcal{M}_c[\pi\mathcal{M}_cf_{\text{GR}}(t)]^{B-1} \times 2\pi f_{\text{GR}}(t) \\ &= \phi_{\text{GR}}(t) + B\beta[\pi\mathcal{M}_cf_{\text{GR}}(t)]^B,\end{aligned}\tag{5.44}$$

in which case equation (5.35), to $\mathcal{O}(1)$ in the amplitude, becomes

$$h(t) = A_{\text{GR}}(t) \cdot e^{i\phi_{\text{GR}}(t) + i\beta u_{\text{GR}}^b} = h_{\text{GR}}(t) \cdot e^{i\beta u_{\text{GR}}^b},\tag{5.45}$$

where recall $u_{\text{GR}}^b = (\pi\mathcal{M}_cf_{\text{GR}})^B$ and $B \equiv b/3$. We shall take equation (5.45) as our modified gravity waveform and perform our analyses with this model.

5.5.1 Validity of time-domain approximants

In Ref. [128], the upper bound on the EdGB coupling constant was found by analysing six GW events (chosen specifically for their capacity to constrain EdGB and dCS gravity) to be approximately $\sqrt{\alpha_{\text{EdGB}}} \lesssim 1.7$ km (at 90% credibility). The dCS parameter could not be constrained further than previously existing constraints, and remains at $\sqrt{\alpha_{\text{dCS}}} \lesssim 8.5$ km. Consider, as examples, the 2 fiducial CBCs that we will later study in detail in Chapters 7 & 8; a ‘low-mass’ and a ‘high-mass’ system (see Table 7.1). For these systems, by equation (5.6), the constraints translate to the dimensionless quadratic gravity parameters as

$$\begin{aligned}|\zeta_{\text{EdGB}}^{\text{low-mass}}| &\lesssim 3.62 \times 10^{-6} & |\zeta_{\text{dCS}}^{\text{low-mass}}| &\lesssim 2.27 \times 10^{-3} \\ |\zeta_{\text{EdGB}}^{\text{high-mass}}| &\lesssim 1.92 \times 10^{-9} & |\zeta_{\text{dCS}}^{\text{high-mass}}| &\lesssim 1.20 \times 10^{-6}.\end{aligned}$$

Given the mass and spin parameters of the fiducial systems, from equations (5.10) and (5.11) we find that

$$\begin{aligned}|\beta_{\text{EdGB}}^{\text{low-mass}}| &\lesssim 3.57 \times 10^{-9} & |\beta_{\text{dCS}}^{\text{low-mass}}| &\lesssim 1.12 \times 10^{-4} \\ |\beta_{\text{EdGB}}^{\text{high-mass}}| &\lesssim 2.41 \times 10^{-12} & |\beta_{\text{dCS}}^{\text{high-mass}}| &\lesssim 6.19 \times 10^{-8}.\end{aligned}$$

These β values are indeed mostly small, and (perhaps with the exception of the low-mass dCS value) could be said to satisfy the requirement that $\beta \ll 1$ that is necessary to obtain the modified gravity waveform as in equation (5.45). It has been mentioned a number of times already that the measurement is very sensitive to phase, so one cannot be too careful when approximating the phase or its argument. This is most certainly

true for real-world data analysis, and strongly desired in simulations if one wishes to reproduce posteriors to high accuracy. For simulated settings, where the goal is, for example, to acquire some general properties of the structure of the posterior, a modicum of leniency is sometimes permissible (since the signal is generated by the model), however, it should be stressed that the posteriors we acquire can only be representative of the modified gravity model provided by [128] if the majority of the probability for the $\beta_{\text{EdGB,dCS}}$ parameter is contained within a small $\beta_{\text{EdGB,dCS}}$ limit (centred at zero), translating directly to requiring that $\zeta_{\text{EdGB,dCS}} \ll 1$. This can now be relatively strongly assured in practice using these results from Ref. [128] as prior bounds on $\zeta_{\text{EdGB,dCS}}$ for future Bayesian PE.

We shall see the results in Chapters 7 & 8: our findings reflect those of [128], as we show that the EdGB gravity parameter can be very accurately measured with some detector combinations we consider, since the perturbative regime (small β approximation) is enough to strongly constrain it. On the other hand, the dCS gravity parameter is poorly constrained: the posteriors “overflow” the small β limit and more detailed understanding and modelling of the waveform is required before any posteriors defined by the dCS GW model can be trusted. For our particular model approximant, the results primarily mean that our derivation via the small β_{dCS} limit is not applicable in our PE context, since we will find non-vanishing likelihood at relatively large β_{dCS} values.

In the broader context, this also seeds concern that the currently known dCS waveform may not be precise enough to be able constrain dCS gravity with GW measurements even using DECIGO, which can yield SNRs in the tens of thousands. A better knowledge of the waveform will increase our chances of constraining dCS gravity, since, as explained in [128], higher-order PN terms can only decrease degeneracies between the dCS parameter and other parameters, and thus serve only to improve constraints of it; by omitting the higher-order quadratic gravity terms, one acquires a lower bound in the parameter constraints. It is presently unclear how well each of these extra terms will perform in further constraining the dCS parameter as they are not yet known. If $\beta_{\text{dCS}} \approx 0$, the extra terms may ultimately not be required when extremely sensitive future ground-based detectors (as yet not conceptualised) observe ideal compact binaries, capturing the strong gravity regime merger where dCS modifications are greatest. For the time-domain model in particular, as a result of the very high SNRs one might achieve, these observations

could constrain the dCS parameter to a high enough degree for the parameter estimation results to constrain all the probability of β_{dCS} within bounds of the small parameter limit, where the $\mathcal{O}(\beta_{\text{dCS}}^2)$ terms can be ignored.

5.6 Acceleration/modified gravity bias

It is understood that acceleration of a BHB and other environmental matter effects, such as accretion, can be partially degenerate with modified gravity parameters [99]. We will show here how this degeneracy arises, and provide some more useful general results for later discussion.

Integrating equation (5.36) for the phase, we obtain

$$\phi_{\text{GR}}(t) = \phi_0 + \frac{5M}{\eta} \left(\hat{\omega}_5 \ln(\tau_c) + \sum_{\substack{k=0 \\ k \neq 5}}^{N/2} \frac{\hat{\omega}_k}{5-k} \theta^{k-5} \right), \quad (5.46)$$

where recall $\theta(t) = [\eta\tau_c/(5M)]^{-1/8}$, η is the symmetric mass ratio, M is the total mass, $N/2$ is the PN-order, $\tau_c = t_c - t$ is the time to coalescence and the $\hat{\omega}_k$ are PN term coefficients (given in [27]). The above is the GR phase. The EdGB modification, for example, comes in at PN-order -1 ; this would appear as a new term in the sum over k at $k = -2$. That is, we would have a new term in the phase proportional to θ^{-7} . To confirm the frequency domain modification modifies the time-domain waveform appropriately in PN terms, we can inspect the phase of our waveform approximation in equation (5.45):

$$\phi(t) = \phi_{\text{GR}}(t) + \beta u_{\text{GR}}^b, \quad (5.47)$$

where $u_{\text{GR}}^b = (\pi\mathcal{M}_c f_{\text{GR}})^{b/3}$. Using (5.40), this becomes

$$\begin{aligned} \phi(t) &= \phi_{\text{GR}}(t) + \beta \left(\frac{5}{256} \frac{\mathcal{M}_c}{\tau_c} \right)^{b/8} \\ &= \phi_{\text{GR}}(t) + \beta \left(\frac{\mathcal{M}_c}{256} \frac{\eta}{M} \right)^{b/8} \theta^b, \end{aligned} \quad (5.48)$$

and so clearly we have agreement that the new term in the phase is proportional to θ^{-7} for EdGB modifications, since $b_{\text{EdGB}} = -7$.

Any physical effect that modifies the phase across the signal in a similar way can cause the signal to appear as though generated by EdGB gravity if not accounted for in the model. Such biases for some modified gravity theories are indeed to be expected for

those signals generated by, for example, binaries undergoing acceleration relative to the observer [99].

For the most simple acceleration model, that of constant acceleration, which we model by modifying the detector time, the new detector time can be written

$$t \rightarrow t_{\text{acc}}(t) = t + t_0 + \frac{v_0}{c}t + \frac{a}{2c}t^2, \quad (5.49)$$

for constants t_0, v_0, a . Using $t = t_c - \tau_c$ and expanding the t^2 term, then ignoring the constants, since they can be absorbed into the reference time, and the (extra) terms linear in t , since they describe a constant velocity which can be absorbed into the chirp mass, the phase of a constantly accelerated binary can be written

$$\phi(t) \rightarrow \phi_{\text{acc}}(t) \equiv \phi(t_{\text{acc}}) = \phi(t + a\tau_c^2/(2c)). \quad (5.50)$$

If $a \ll 1$, we can write this as

$$\begin{aligned} \phi_{\text{acc}}(t) &= \phi(t) + \frac{\pi a}{c} \tau_c^2 f_{\text{GR}} \\ &= \phi(t) + \frac{a}{c} \left(\frac{5}{256} \right)^{3/8} \mathcal{M}_c^{-5/8} \tau_c^{13/8}, \end{aligned} \quad (5.51)$$

using (5.40). The extra term in the phase due to constant acceleration is proportional to $\tau_c^{13/8}$, itself proportional to θ^{-13} . Then, since the exponent b (of θ) is related to the PN order p_{PN} by $p_{\text{PN}} = (b + 5)/2$ and $b_{\text{acc}} = -13$, we have that $p_{\text{acc}} = -4$. More generally, a term in the perturbative time modification proportional to τ^n gives a term in the phase proportional to $\tau^{n-3/8}$ with $p_{\text{PN}} = -4(1 + n)$.

Perturbatively at least, the constant acceleration is therefore exactly degenerate with, for example, *macroscopic extra dimensions* models and the *time varying G* model [110], which both enter the waveform at -4PN . It is stressed that this does not necessarily hold exactly in the non-perturbative regime. For example, for accelerations, we can produce a very accurate model just by determining t_{acc} , and this will affect all terms in the PN-expansion in the same way (we modify the waveform in *propagation* by taking $t \rightarrow t_{\text{acc}}$; this transformation is applied to all terms), but different modified gravity theories generally introduce the various higher-order PN-terms with different coefficients (modifications in *generation*), breaking the exact degeneracies.

The p -PN term ($p \neq 5/2$) has a time dependence $\propto \tau_c^{(5-2p)/8}$, so that terms of PN order $p < 5/2$ dominate as $\tau_c \rightarrow \infty$ and terms of PN order $p > 5/2$ dominate as $\tau_c \rightarrow 0$.

One might expect that the high magnitude powers of τ_c respectively dominate strongly enough for $\tau_c \gg 1$ and $\tau_c \ll 1$ that the different PN terms are generally not strongly correlated. However, we will see in Chapter 8 that acceleration effects at -4PN , which at first seems quite distinct from the -1PN contribution of EdGB gravity, is still able to cause strong EdGB bias when the acceleration is not modelled, and vice versa.

Analysis Framework Details

In this short chapter, we finalise and bring together the modelling and signal analysis results of the previous chapters, pointing out important properties of the data analysis framework and issues to be aware of in sampling, and making clear the shortcomings of the model we are using. So far, we have primarily dealt with low-frequency band modelling and analysis methods, which we now bring into a complete *multi-band* analysis framework that simultaneously includes high-frequency band data. We further emphasise simplifying choices and analytical methods (i.e., the neglected and marginalising parameters) that we will employ, and some features of the physics that should be borne in mind when making conclusions based on results obtained. It is hoped this will provide a helpful overview of the scope and limitations of the results presented in the final chapters.

6.1 Multi-band data analysis

Multi-band data analysis is simply the analysis of a signal generated by a single source that is present in multiple frequency *bands* (intervals of the frequency domain) defined by distinct detectors in a network. A waveform might smoothly pass through bands, for example, the chirping inspiral of a binary which sweeps from low to high frequencies, or a source might emit waves that can be decomposed into different frequencies such that different frequency components are contained within separate frequency bands. A possible source of the latter type could be the harmonic overtones one finds in the ring-

down of a CBC merger, with different harmonics present in different frequency bands. This thesis is concerned exclusively with the chirping BHBs, with a primary goal being to elucidate reasonable expectations of multi-band GW data analysis.

It is hoped that the chirping behaviour of BHBs can be exploited; since different amounts of information on, and correlations between parameters are carried by the signal through different frequency bands, it is expected that the combination of data from multiple frequency bands will help to: better constrain poorly constrained parameters, break parameter degeneracies and test GR (with or without modelling alternatives to GR; a merger time in LIGO inconsistent with the predictions of GR based on the expected merger time from LISA could suggest a rate of energy loss to GWs inconsistent with GR [141]). Multi-band data may be particularly useful for testing modified gravity theories [110] by constraining negative PN orders in the waveform (see Section 5.6), as well as potentially being able to detect the low-frequency waveform modifiers we mentioned in Chapter 4. It has been shown that multi-band observations may number in the hundreds [142] for signals individually resolvable in LISA, and that, using future ground-based detectors to look back at what will be ‘historical’ LISA data, the number of multi-band detections could be increased by a factor of $\gtrsim 4$ [85].

6.1.1 Multi-band detector networks

In part because of simplicity and familiarity, we have chosen three detectors to use for the analyses presented in Chapters 7 & 8. They are: LISA, DECIGO, and LIGO (see Section 2.1), and the multi-band networks shall consist of all (appropriate) combinations of these detectors. This choice is because the LISA mission is confirmed, LIGO is operational and its successors will operate at roughly the same band (although with lower noise floor), and DECIGO, which is a rather speculative mission at this stage, is intended to ‘plug the gap’ of frequencies between LISA and LIGO (see Figure 2.3). The very high SNRs one obtains from DECIGO make it somewhat of an unfair comparison of observing at the given frequency range. However, we will see that DECIGO is necessary for making any progress in constraining some parameters, and point out some interesting features of including this detector in a network.

6.1.2 Multi-band parameter estimation

By probability rules, the probability of two independent events, A and B , with probabilities $P(A)$ and $P(B)$, both occurring is

$$P(A \cup B) = P(A) \cdot P(B). \quad (6.1)$$

This is also true for continuous variables, so since the noise in one detector is independent from the noise in another detector, the joint probability of datasets \mathbf{d}_i observed by detectors with data probabilities $p_i(\mathbf{d}_i | \theta)$

$$p(\mathbf{d}_1, \dots, \mathbf{d}_n | \theta) = \prod_{i=1}^n p_i(\mathbf{d}_i | \theta) \quad (6.2)$$

and the posterior probability of θ , from Bayes' theorem equation (2.6), is just

$$p(\theta | \mathbf{d}_1, \dots, \mathbf{d}_n) = \frac{p(\theta) p(\mathbf{d}_1, \dots, \mathbf{d}_n | \theta)}{p(\mathbf{d}_1, \dots, \mathbf{d}_n)}. \quad (6.3)$$

This is simple enough, but how should we use this in practice? This depends on the research goals and the means by which one 'knows', or is 'in possession of', the likelihood function. For example, if one intends to perform large numbers of comparisons of posteriors from different detector networks, it would be useful to evaluate the likelihoods from each detector given some signal and store them as scalar functions (since the evaluations can be very computationally costly). Then the scalar functions can be multiplied together (then normalised) as required for the desired multi-detector network posterior.

However, obtaining the scalar functions is not straightforward. Two common means of obtaining (approximations of) them are by evaluation directly over a lattice of points defined on the parameter space, or by using sampling methods to obtain samples with a local sample density at a point approximating the scalar value of the posterior. One can then obtain the scalar using techniques such as *kernel density estimate (KDE)*, then, converting and storing this as a scalar would additionally require some sort of lattice representation, so that multiplying posteriors together can be done pointwise on the lattice.

When the parameter space dimensionality becomes even moderately high (around three parameters), to maintain a reasonable detail (i.e., resolution of the lattice) the direct lattice evaluation approach quickly becomes computationally intractable. The KDE approach, on the other hand, becomes unreliable and depends upon careful manual calibration of the KDE settings for each posterior. Thus, for high-dimensional problems, we

are rather limited in how we can go about producing multi-band posteriors; for accurate representation of the posteriors, we must *evaluate the posterior by sampling from the joint likelihood of the multiple datasets from the detectors in the desired network*.

6.2 Exploring the framework

Initial experimentation with the analysis framework brought to light a number of interesting issues and features. We point out the more important ones here (some issues do not need to be raised, for example, sampler configuration settings which have since become default settings of the sampler package).

6.2.1 Comparison of a posterior with its PCM approximation

Recall from Section 3.3.2 that the parameter covariance matrix (PCM) at the injection values is the inverse of the Fisher matrix there. It defines an ellipsoid with a boundary that coincides with the 1-sigma isoprobability contour of the Gaussian that is the leading-order Taylor expansion of the posterior (at the injection values). Because of the simple definition and ease of computation and study of the PCM, it is very often used as an approximation of the posterior, which is then used to infer PE expectations, particularly in place of posteriors that are very expensive to compute, such as LISA posteriors.

Figure 6.1 shows a posterior of a reduced parameter model superimposed on top of its PCM posterior. We include this to illustrate that the PCM can often be a poor representation of the true posterior; the variances and covariances are not in close agreement, and higher-order structure is of course completely absent from the PCM. The example in Figure 6.1 is in fact not such a pathological example and shows roughly order of magnitude agreement of variances, which is generally considered a reasonable gauge of the true posterior [88]. However, these differences can grow significantly when the SNR drops, the number of free parameters increases, or when one combines two PCM posteriors (as in multi-band parameter estimation). It can also be difficult to accurately compute the FIM in the first place, for instance, in cases where extremely thin modes exist, since one defines the FIM by computing numerical derivatives (of either the likelihood itself or the waveform). The following chapters therefore rely on results derived from the accurate posteriors derived using the downsampling procedure given in Chapter 3.

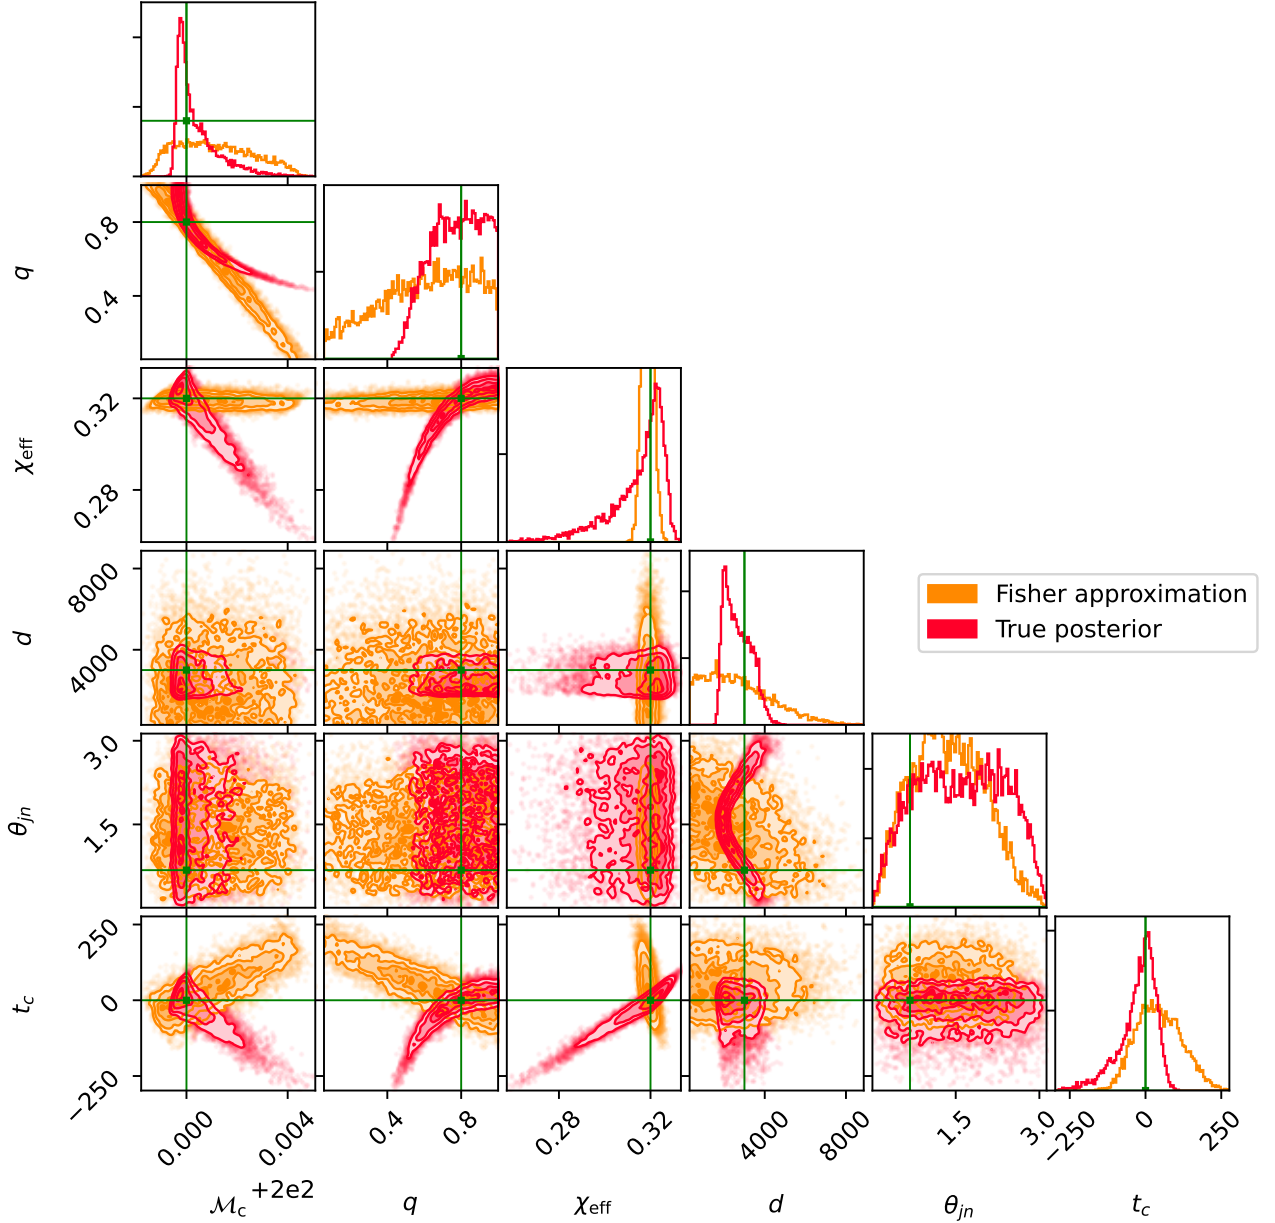


Figure 6.1: Comparison of the posterior approximated by the parameter covariance matrix (PCM: the inverse of the Fisher information matrix) and the true posterior for an example binary system, with injection parameters given by the green lines. The posterior is marginalised over polarisation ψ and with other parameters fixed, generated by a ‘high-mass’ (Table 7.1) binary observed by LISA (10 years) only. In this particular example, the one-parameter marginals arguably show fair agreement to order-of-magnitude. However, other example systems were found that clearly show that the PCM approximation can be a poor representation of the one-parameter marginalisations. We can see the PCM can be a poor tool for assessing parameter covariances, since some covariances are almost opposite in some cases. For precision analyses as we strive for in Chapters 7 & 8, one should not rely on Fisher information approximations.

6.2.2 Threshold SNRs

The performance of LISA in providing useful contributions to parameter constraints is strongly dependent on the signal’s SNR in LISA, and on the number of parameters and their ability to affect change upon the ‘number of templates’ required for the ‘template bank’ that is used in parameter estimation¹ [143]. One might have hoped to acquire better results from stellar order binaries in LISA by, say, neglecting higher-order spin parameters from parameter estimation (as we shall do, see Section 6.2.3) since at early stages in the inspiral the orbit more closely resembles a Newtonian orbit, where higher-order GR effects are less prevalent. The reasons would be to reduce the number of parameters, hoping that a signal might more easily be found despite searching with an incorrect model (with some parameters fixed). However, those parameters that do not affect the signal too much (and so can be fixed/omitted more leniently) are precisely the ones that do not much affect the SNR detection threshold. Thus, in general, we should fully model (but not over-model, see Sections 4.2.4 & 8.5; it is beneficial to optimise the number of templates and avoid duplicates).

6.2.3 The large dimensionality/sampling problem

Despite the use of our downsampling procedure to speed up the evaluation of the LISA likelihood, an initial assumption of aligned spins gives a posterior in which the two remaining spin parameters (the magnitudes) form a thin mode in the spin-spin space, appearing to cause a significant impact on the convergence time. Additionally, the large dimensionality of the parameter space resulted in convergence times too long for the large set of systems for which we anticipated producing posteriors in Chapters 7 & 8.

6.2.3.1 Reduced spin parameter model

For the former problem, to speed up parameter estimation, we simply opt to reduce the model by assuming only spin aligned, equal spin black-holes, yielding binaries with a

¹One way to think of this is by considering the following: if a new parameter is introduced that, over the allowed range (i.e., its priors), defines waveforms that are not similar to the existing waveforms (i.e., those existing in the template bank) then the number of templates must significantly increase; if, on the other hand, the waveforms defined by the new parameter has little or no effect on the waveform, then the number of templates rises marginally or not at all.

single spin component, χ_{eff} (see Section 6.3.1 for more details). This is the spin parameter most strongly coupled to the signal in the early inspiral stage of the GW. The higher-order spin parameters appear in the waveform in the higher-order PN terms, and thus couple more strongly towards merger as one enters the highly non-linear regime of GR; at earlier stages of the inspiral, the binary orbit approaches a Newtonian orbit where GR effects become less important. One can therefore expect that LISA will not strongly constrain the higher order spin effects. Omitting (fixing) the spin precession, χ_p , increases the PE convergence speed by around a factor of two.

6.2.3.2 Numerical marginalisation of phase in time-domain

For the latter problem, we can remove uninteresting (or *nuisance*) parameters from the parameter space by marginalisation. The coalescence phase is one such parameter. The phase marginalised probability is given by

$$p'(\boldsymbol{\theta}' | \mathbf{d}) = \int_0^{2\pi} d\phi_c p(\boldsymbol{\theta} | \mathbf{d}) = \int_0^{2\pi} d\phi_c \frac{p(\boldsymbol{\theta}) p(\mathbf{d} | \boldsymbol{\theta})}{p(\mathbf{d})}, \quad (6.4)$$

where $\boldsymbol{\theta}'$ is the tuple of parameters excluding the coalescence phase, ϕ_c . If the prior on ϕ_c is $(2\pi)^{-1}$, then the prior on $\boldsymbol{\theta}$ can be written

$$p(\boldsymbol{\theta}) = \frac{p'(\boldsymbol{\theta}')}{2\pi},$$

where $p'(\boldsymbol{\theta}')$ is the prior on $\boldsymbol{\theta}'$. Thus

$$p'(\boldsymbol{\theta}' | \mathbf{d}) = \frac{p'(\boldsymbol{\theta}')}{p(\mathbf{d})} \int_0^{2\pi} d\phi_c \frac{p(\mathbf{d} | \boldsymbol{\theta})}{2\pi}. \quad (6.5)$$

Therefore, one can obtain the marginal probability simply by substituting a marginalised likelihood. Now since we can write the waveform in the form $h(t) = A(t) \cos(\phi(t) + \phi_c)$, then from equation (2.18) we have that

$$\begin{aligned} p(\mathbf{d} | \boldsymbol{\theta}) &= [(2\pi)^N \det(\boldsymbol{\Sigma})]^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{h}(\boldsymbol{\theta}))^\top \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \mathbf{h}(\boldsymbol{\theta})) \right] \\ &= [(2\pi)^N \det(\boldsymbol{\Sigma})]^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{A} \cos(\boldsymbol{\phi}' + \phi_c))^\top \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \mathbf{A} \cos(\boldsymbol{\phi}' + \phi_c)) \right], \end{aligned}$$

where $\boldsymbol{\phi}' = \boldsymbol{\phi}'(\boldsymbol{\theta}')$. By trigonometric identities, we have

$$\begin{aligned} (\mathbf{d} - \mathbf{h}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \mathbf{h}(\boldsymbol{\theta})) &= (\mathbf{d} - \mathbf{A} \cos(\boldsymbol{\phi}') \cos(\phi_c) + \mathbf{A} \sin(\boldsymbol{\phi}') \sin(\phi_c))^T \\ &\quad \times \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \mathbf{A} \cos(\boldsymbol{\phi}') \cos(\phi_c) + \mathbf{A} \sin(\boldsymbol{\phi}') \sin(\phi_c)) \\ &= A \cos^2(\phi_c) + B \cos(\phi_c) \sin(\phi_c) + C \cos(\phi_c) \\ &\quad + D \sin^2(\phi_c) + E \sin(\phi_c) + F, \end{aligned}$$

where

$$\begin{aligned} A &= (\mathbf{A} \cos(\boldsymbol{\phi}'))^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \cos(\boldsymbol{\phi}') & D &= (\mathbf{A} \sin(\boldsymbol{\phi}'))^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \sin(\boldsymbol{\phi}') \\ B &= -2(\mathbf{A} \cos(\boldsymbol{\phi}'))^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \sin(\boldsymbol{\phi}') & E &= 2\mathbf{d}^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \sin(\boldsymbol{\phi}') \\ C &= -2\mathbf{d}^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \cos(\boldsymbol{\phi}') & F &= \mathbf{d}^T \boldsymbol{\Sigma}^{-1} \mathbf{d}. \end{aligned}$$

Then the marginalised likelihood is

$$\int_0^{2\pi} d\phi_c \frac{p(\mathbf{d} | \boldsymbol{\theta})}{2\pi} = \int_0^{2\pi} d\phi_c \frac{e^{-\frac{1}{2} (A \cos^2(\phi_c) + B \cos(\phi_c) \sin(\phi_c) + C \cos(\phi_c) + D \sin^2(\phi_c) + E \sin(\phi_c) + F)}}{2\pi [(2\pi)^N \det(\boldsymbol{\Sigma})]^{1/2}}. \quad (6.6)$$

There does not appear to be a straightforward method to obtain a closed form expression of this integral, but a numerical integration can be achieved easily and performed very cheaply, as we shall show. The computationally expensive part is in evaluating $\boldsymbol{\phi}'$. There is a moderately expensive operation in computing inner products in the definition of the constants (given some $\boldsymbol{\theta}'$) A to F . However, notice that all the constants are easily computed by various inner products of the vectors $\boldsymbol{\Sigma}^{(-1/2)} \mathbf{d}$, $\boldsymbol{\Sigma}^{(-1/2)} \mathbf{A} \sin(\boldsymbol{\phi}')$, and $\boldsymbol{\Sigma}^{(-1/2)} \mathbf{A} \cos(\boldsymbol{\phi}')$, so only these three matrix operations are required, and the downsampling significantly reduces the cost of these matrix operations. For many likelihood evaluations, store the sines and cosines (and their squares) of a range of ϕ_c values in memory so as not to reevaluate them on each likelihood function call.

One can now approximate the integral by a discrete sum. We tested the numerically ‘pre-marginalised’ posteriors against ‘post-marginalised’ posteriors (that is, we acquired the full posterior from the non-marginalised model, and marginalised this) and found that the number of terms in the sum must be decreased to around 20 before the approximation begins to fail. The likelihood evaluation time only begins to increase noticeably (by a few percent) when the number of terms in the sum reaches around 10^4 . We used 10^3 terms, giving an accurate marginalised likelihood with negligible likelihood evaluation time increase. The PE convergence time, however, with the phase parameter removed, is significantly decreased by a factor of around three.




Figure 6.2: An example of the difference in the convergence of a sampler. Here, two attempts were made at running the sampler `Nessai` on precisely the same posterior, but with different ‘seeds’ which initialise the (pseudo-)random processes involved in sampling. Posteriors were defined using the ‘low-mass’ system from Table 7.1 using LISA+LIGO. The first attempt failed to locate the posterior mode with maximum likelihood (returning instead a LIGO posterior distorted by extra LISA data which it failed to find fitting parameters for, see discussion on Figure 7.2 for details on this phenomenon). The second attempt succeeded in fitting the model to the data. In simulations at least, one can avoid this failure by choosing tight priors.

6.2.4 Real-world data analysis: selection of priors

It is prudent to point out problems one may encounter in the analysis of real-world LISA data. Figure 6.2 shows superimposed resulting sets of posterior samples from running a sampler twice on precisely the same data. Due to the random nature of the sampler's path of exploration of the posterior (and the finite time allowance for exploration) it is partly to chance whether the primary mode will be found. One can clearly see the first attempt at evaluating the posterior in Figure 6.2 was unsuccessful in locating the injected parameters, whereas the second attempt succeeded. The selection of priors plays a large part in determining whether the primary mode will be located; the sampler is 'searching for a needle in an n -dimensional haystack'. This becomes very difficult when n is large; tighter priors ameliorates this issue.

6.3 Modelling caveats

Finally, we must point out some of the shortcomings of our assumptions and models, with which the results of Chapters 7 & 8 are produced. These are either in the form of missing (or fixed) parameters which are not searched over in parameter estimation, and thus lead to the posteriors we obtain being slices through what would be considered the 'true' posterior, or an incomplete modelling of the physical system (both in defining the waveform and the detector), which can potentially lead to unrepresentative results. These caveats are besides the long list of possible deformations of the waveform that were mentioned in Section 4.1.

An important modelling simplification employed in this thesis is that there is only one BHB signal present in the data. This is of course highly unrealistic; there will be many overlapping BHB inspiral signals in LISA (as well as other signals such as mergers of SMBHs) and many BHB mergers/ringdowns in future ground-based detectors such as CE and ET, and our multi-band data analysis results are derived without the problem of matching a LISA inspiral with the correct LIGO merger. CE, for example, is expected to observe up to thousands of BHB mergers per day (on the order of one merger per second). Depending on the model and signal, LISA's error on any given inspiral's coalescence time ranges from seconds to hours or more, thus there shall be multiple CE mergers with coalescence times compatible with a LISA inspiral coalescence time.

However, many mergers should be ruled out by sky position constraints alone, which will be strong enough from both detectors to disentangle the signals. Ultimately, one must simply perform a many-signal (a GW signal consisting of $n > 1$ BHBs) PE analysis on the data, which may or may not be able to pair up a given inspiral with its corresponding merger. If the network also contained DECIGO however, then one could track the signal all the way through inspiral to merger.

A further limitation of note is that all of our simulated signals will leave LISA around one month before merger. For a given detectable BHB signal in LISA however, the merger may take place up to decades after the inspiral leaves the LISA band, in which case the constraints on the coalescence time from LISA are considerably weakened. In those intervening decades, whilst some modified gravity theories may have already been highly constrained, the possibility of using GWs to study nuclear galactic structure using BHBs will continue to rest solely with LISA and other low-frequency observatories, and of course continuing to gain more precise knowledge on the contents of the LISA data will continue to improve PE of every detectable source, gradually improving LISA science many years after its mission.

6.3.1 Missing parameters

6.3.1.1 Spin parameters

One of the standard parameters of BHBs (see Section 1.4.2) that is not included in our model for the sake of simplicity and which has already been mentioned is the spin precession parameter. Whilst the effective spin is often well-constrained, the spin precession is usually much less well-constrained [24]. Occasionally, a parameter can be poorly constrained due to very weak coupling to the signal, in which case one may be justified in ignoring its effects. However, this is generally not the cause of poor constraints of a parameter, for example, highly correlated parameters can lead to poor constraints. It is also worth remarking that the spin-spin coupling enters the waveform at 2PN [26], the same as the leading-order dCS contribution, so one should also expect significant correlation between the spin and dCS parameters. Omission of spin precession, therefore, limits the scope of validity of the results presented here, and shall need to be dealt with appropriately in future work.

6.3.1.2 Eccentricity

Furthermore, the BHB model should account for orbital eccentricity. This is especially true in the lower frequency parts of the signal; as one progresses through the chirping signal increasing time and frequency, eccentric BHB orbits tend to ‘circularise’ since the points of closest approach of elliptical orbits have the highest accelerations and strongly dissipate (via GW radiation) orbital angular momentum there, periodically, such that the orbit becomes (quasi-)circular. Early on in the signal, significant eccentricities may be present. The eccentricity is thought to be strongly dependent on the formation scenarios of the binaries [116].

High eccentricities generally decrease the instantaneous SNR at a given frequency, and it has been shown that the eccentricities that can be reasonably expected are high enough to have a significant impact on the detection rate of binaries in LISA [144]. Whilst this (absence of detections) can be useful to some degree for constraining formation scenarios, the eccentricity can be expected to severely limit the prospects of multi-band GW astronomy of stellar-origin BHBs with LISA, and is therefore another important missing aspect of the model that limits the results presented in this thesis and should be considered in future work. At the same time as eccentricity decreasing the SNR, it may increase the SNR detection threshold (see Section 6.2.2) further hampering detection efforts.

6.3.1.3 SMBH spin

We have also omitted the SMBH spin, which may imprint significantly upon the GW if the periapsis of the line-of-sight of the GW approaches the SMBH closely enough. GWs follow null geodesics, which, for those close to the event horizon of Kerr black-holes, can undergo large diversions as compared to the Schwarzschild case. The frame-dragging effect from the spin of an SMBH may also directly modify the orientation and eccentricity of binaries in close orbit [145], as well, indeed, as its position/orbital path, making significant imprints on the evolution of the binary, and thus the GW.

6.3.2 Omitted physical effects

6.3.2.1 Hubble expansion

We have employed a Minkowski universe model here for clarity. The GW model in an expanding universe exhibiting redshifting phenomena is acquired by simply replacing the GW model's chirp mass \mathcal{M}_c and distance d by, respectively, the *redshifted chirp mass* $\mathcal{M}_{c,z} = \mathcal{M}_c(1+z)$ and the *luminosity distance*, $d_L = d(1+z)$, where z is the redshift [146]. See Section 6.3.3 for details on redshifted GR time-delay modified waveforms.

6.3.2.2 Shapiro and geometric time-delay

Second order Shapiro delay effects can become significant on SMBH scales, and in general in dealing with black-holes, rather than extended bodies such as the Sun. The first and second order terms in r_s of the Shapiro delay, where r_s is the Schwarzschild radius, can be interpreted as the ‘variable [coordinate] speed of the light ray’, whereas the second is due to the change in path length [124]. The latter is also known as the *geometric time-delay*. It is most noticeable that the second order terms in r_s are in fact of $\mathcal{O}\left(\frac{r_s^2}{cr_0}\right)$ [124], where r_0 is the impact parameter (see Section 4.2.2.5), so that as $r_0 \rightarrow r_s$, the second order terms tend towards having similar magnitudes as the first order terms. Also note that, for example, our example Orbit #3 in Table 4.2 gives a second order Shapiro delay term $\mathcal{O}\left(\frac{r_s^2}{cr_0}\right) \sim 4 \times 10^{-2}$ s (at the largest r_0). The order of magnitude of the second order terms is inversely proportional to r_0 , so as the orbit progresses and r_0 decreases from its maximum, the order of magnitude of second order Shapiro delay will increase, and thus potentially significant time-delay terms will be present.

6.3.3 Redshift and mass/distance scale degeneracy

The peculiar motion of the binary in relation to the SMBH of its host galaxy is encoded in the waveform and is in principle measurable. However, the host galaxy has its own peculiar motion and Hubble recession, and associated redshift, and this remains degenerate with the mass and distance scale (local to the source) as we shall show here. This remaining redshift can be written as a Roemer delay. We ignore the Hubble redshift here for clarity and work in flat spacetime but the results here will hold for Hubble redshift

also. Including the peculiar velocity time-delay $\Delta t_{\text{pec}} = v_{\parallel}t/c = zt$, where $v_{\parallel} = cz$ is the line-of-sight velocity of the galaxy and z is the redshift, with the total time-delay we derived in equation (4.34), we obtain the new (redshifted) observer time

$$\begin{aligned} t_{\text{tot,pec}} &= t_{\text{tot}} - \Delta t_{\text{pec}} \\ \Rightarrow (1+z)t_{\text{tot,pec}} &= t_{\text{tot}} \\ \Rightarrow t_{\text{tot,pec}} &= \frac{\tau_{\text{src}} - \Delta t_{\text{Shapiro}} - \Delta t_{\text{Roemer}}}{1+z}, \end{aligned} \quad (6.7)$$

where we used equations (4.34) and (4.29).

If τ describes the local rate of evolution of a binary of chirp mass \mathcal{M}_c , then $\tau/(1+z)$ describes the observed evolution of either a binary receding at velocity cz or instead a binary with chirp mass $(1+z)\mathcal{M}_c$ by the mass-redshift degeneracy. Consequently, we have that

$$h(\tau/(1+z), \mathcal{M}_c) = h(\tau, \mathcal{M}_c(1+z)). \quad (6.8)$$

That is, the (unmodified) waveform is invariant under simultaneous transformations

$$\begin{aligned} \tau &\rightarrow \tau' = \tau(1+z) \\ \mathcal{M}_c &\rightarrow \mathcal{M}'_c = \mathcal{M}_c(1+z). \end{aligned}$$

But it is also true that the time-delay modified waveform remains invariant under the additional transformations of the SMBH mass and semi-major axis

$$\begin{aligned} M_{\text{SMBH}} &\rightarrow M'_{\text{SMBH}} = M_{\text{SMBH}}(1+z) \\ a &\rightarrow a' = a(1+z). \end{aligned}$$

To see this, apply the above four transformations simultaneously to the redshifted waveform:

$$\begin{aligned} h(t_{\text{tot,pec}}(M_{\text{SMBH}}, a), \mathcal{M}_c) &= h\left(\frac{\tau_{\text{src}} - \Delta t_{\text{Shapiro}}(M_{\text{SMBH}}, a) - \Delta t_{\text{Roemer}}(M_{\text{SMBH}}, a)}{1+z}, \mathcal{M}_c\right) \\ &\rightarrow h\left(\frac{\tau'_{\text{src}} - \Delta t_{\text{Shapiro}}(M'_{\text{SMBH}}, a') - \Delta t_{\text{Roemer}}(M'_{\text{SMBH}}, a')}{1+z}, \mathcal{M}'_c\right). \end{aligned} \quad (6.9)$$

The proper time is in fact dependent on M_{SMBH} and on a (since r depends on a), but one can see in equation (4.29) that since $a \rightarrow a'$ and hence $r \rightarrow r(1+z)$ (and performing a

change of variable $t' \rightarrow t'(1+z)$ in the integral), then the factors of $(1+z)$ cancel out of the square root. Thus one is left with

$$\frac{\tau'_{\text{src}}}{1+z} = \frac{\tau_{\text{src}}(1+z)}{1+z} = \tau_{\text{src}}.$$

The Shapiro delay term, using the Shapiro delay equation (4.30) is then simply

$$\frac{\Delta t_{\text{Shapiro}}(M'_{\text{SMBH}}, a')}{1+z} = \Delta t_{\text{Shapiro}}(M_{\text{SMBH}}, a),$$

and the Roemer delay (with periastron precession) term can be seen in the defining Roemer delay and periastron precession equations (4.6)–(4.8) and (4.14)–(4.15), to be given by

$$\frac{\Delta t_{\text{Roemer}}(M'_{\text{SMBH}}, a')}{1+z} = \Delta t_{\text{Roemer}}(M_{\text{SMBH}}, a),$$

since the factors of $(1+z)$ cancel out. Putting this all together in equation (6.9), we have that

$$\begin{aligned} h(t_{\text{tot,pec}}(M_{\text{SMBH}}, a), \mathcal{M}_c) &\rightarrow h(\tau_{\text{src}} - \Delta t_{\text{Shapiro}}(M_{\text{SMBH}}, a) - \Delta t_{\text{Roemer}}(M_{\text{SMBH}}, a), \mathcal{M}'_c) \\ &= h(t_{\text{tot}}(M_{\text{SMBH}}, a), \mathcal{M}_c(1+z)) \\ &= h(t_{\text{tot,pec}}(M_{\text{SMBH}}, a), \mathcal{M}_c), \end{aligned}$$

where we used the usual mass-redshift degeneracy, equation (6.8), and equation (6.7) in the last line. Therefore the waveform is invariant under the four simultaneous transformations given above, and thus a peculiar velocity (and/or cosmological redshift) for orbiting binaries is degenerate with both the mass & distance scales at the source: positively redshifted signals will appear to derive from heavier binaries around heavier SMBHs on wider orbits, and vice versa.

Multi-band GW Astronomy Part I: Isolated Binaries

One might say the ‘first-order’ effects of multi-band GW astronomy are those observed when, for example, extending the LIGO band model down to lower frequency bands (or, conversely, extending a LISA band analysis to include LIGO band data) with the assumption that no waveform modifiers from environmental influences or modifications of GR come into play, in either the generation or propagation of the GW across the spectrum of the waveform. That is, supposing ordinary GR and supposing the GW emitter is in an otherwise empty universe. When the GW emitter is a CBC, we will refer to this as an *isolated GR binary*. We shall also study binaries in the EdGB and dCS gravity theories, and refer to these as (isolated) EdGB and dCS binaries.

To gain an understanding of the capabilities, advantages and potential problems that might be encountered in multi-band detector network data analysis for GW astronomy, we explore both the determination of source parameters for mission specification detectors, and the general behaviour of posterior probability distributions as functions of some parameter, some property of the signal, or detector network. The present and following chapters shall present an aggregate of results of various specific investigations. As our ‘base systems’ (fiducial BHB systems that provide a starting point and which will undergo some modifications) we take a ‘low-mass’ and a ‘high-mass’ binary, very roughly similar to GW150914 and GW190521, respectively, whose details are provided in Table 7.1 (see Table 1.1 for description of these parameters).

Parameter	Units	Low mass	High mass
Chirp Mass (\mathcal{M}_c)	M_\odot	30.4	200
Mass Ratio (q)	-	0.82	0.80
Luminosity Distance (d_L)	Mpc	410	3000
Effective Spin (χ_{eff})	-	0.32	0.32
Inclination (i)	rad	0.68	0.68
Orientation (ψ)	rad	0.659	0.659
Right Ascension (α_\oplus)	rad	1.375	1.375
Declination (δ_\oplus)	rad	-0.7108	-0.7108
Quadratic Gravity Parameter ($\zeta_{\text{EdGB,dCS}}$)	-	0.0	0.0
Time until coalescence at LISA decommission time (τ_f)	y	0.1	0.1

Table 7.1: The parameters of the ‘low-mass’ and ‘high-mass’ fiducial BHBs. Recall that the phase (ϕ) will be marginalised and is therefore omitted above. We also omit the coalescence time, t_c since it is simply a reference value and not of great interest. We are operating in a Minkowski universe, but for completeness, note that 410 Mpc is around 1.3 billion light years, which, with Hubble parameter $H \approx 70$ km/s/Mpc, corresponds to redshift $z \approx 0.1$, and 3000 Mpc is approximately 10 billion light years, or redshift $z \approx 0.7$.

There is somewhat of an imbalance between these systems in that we choose to study the low-mass binary given a 4 year observing time with LISA, whereas the signal of the high-mass system is given 10 years in LISA. The reason for giving the high-mass system a 10 year LISA signal (and setting it at a higher distance) is to have a roughly similar exchange of the LISA SNR of the high-mass system (~ 13.8 , see Table 7.2) and LIGO SNR of the low-mass system (~ 14.5 , see Table 7.5). Since, for example, we know that LISA can constrain the chirp mass very well, and LIGO can constrain the coalescence time very well, the hope is to acquire additional insight into how the strength of constraints on different parameters given by the different frequency bands affects the overall posterior structure.

The parameters of our fiducial systems are fairly typical, perhaps with the exception of the effective spin, χ_{eff} , which is towards the higher magnitude of spin as compared to the present catalogue of detections [24] to aid in investigating the modified gravity theories (recall that the higher spin systems carry more information on modified gravity),

and of course the chirp mass of the high-mass system, which is very high as compared to the catalogued events. Such high-mass merger events are not expected at very high rates, for example, reference [147] estimates the rate of mergers of $\sim 200M_{\odot}$ total mass systems (with $q = 1, \chi_{\text{eff}} = 0.8$) binaries at $0.056 \text{ Gpc}^{-3} \text{ y}^{-1}$. LISA is of course sensitive to the inspiral part however, not the merger, and so the probability of detecting such a binary is far greater with LISA in principle, even if one does have to wait some decades, say, for the merger to appear in ground based observatories. Our high-mass binary is even more massive than this however, the individual masses are around $\sim 250 M_{\odot}$ and $\sim 200 M_{\odot}$; the LIGO SNR of such high-mass systems does not reach the standard SNR threshold [85] of ~ 8 required to confirm detection, so if such events were occurring they would not be observed by LIGO alone.

7.1 Data analysis pipeline

All of the posteriors produced for analysis in this and the following chapter are defined with a *zero-noise realisation*. This is so that the posteriors themselves are not unnecessarily noisy (since the noise effects are not presently of interest) and will not consist of scatter that could confuse the qualitative aspects of the posterior behaviour that we wish to determine. They are produced using `Nessai v0.4.1` [76] as discussed in Section 6.2.3, which we found to be the most robust and consistent sampler, and potentially the fastest, against `CPNest`, `dynesty`, `Bilby_MCMC` & `emcee`. We used `Bilby` [79] to handle the details of producing the posterior samples via `Nessai`, and using a multi (where appropriate) detector likelihood function as detailed in Section 6.1.2 with phase marginalisation as described in Section 6.2.3.2.

For LISA signals, we used a *downsampled likelihood function* with 500 samples as defined in Chapter 3. For DECIGO signals, we also used a downsampled likelihood function but with 1000 samples, since the waveform less strictly fits the requirement of being ‘slowly evolving’. Downsampling DECIGO signals was not formally tested, however, we noted extremely consistent results throughout many undocumented initial tests of producing posteriors. The DECIGO waveforms are much shorter, but still fit in many cycles ($\mathcal{N}_c \gtrsim 10^4$) due to the higher frequency there. The signals for LISA and DECIGO are defined by the TaylorT3 waveform Ref. [27] and the PSDs informing the likelihood are

given in Section 2.2. The LIGO signal model (we used IMRPhenomC) and likelihood function are provided by `lalSuite` [63].

We ran the code producing each posterior on the *Hawk* supercomputer based at Cardiff University, UK. Processors are either 2.4 GHz with 20 cores, or 2.5 GHz with 32 cores. Occasionally we used multiple cores, but always when PE convergence times are discussed, the posteriors being referred to were produced using a single core.

7.2 ‘Low-mass’ binary

GW150914 is an excellent example of a favourable GW event for LIGO (which made this very first confirmed detection all the more spectacular). The masses were such that the merger occurred towards the higher frequency end of the PSD ‘bucket’ (see Figure 2.3). It was also relatively close, making it quite loud; we will see this is necessary in order for the LISA part to be able to contribute to parameter estimation. In part because of these features, which both lead to a high SNR, and because GW150914 is so familiar, it is a good example system to study with additional consideration of the lower frequency part of the waveform in the context of multi-band data analysis.

The different detector network configurations and the signal properties for each network are given in Table 7.2, where the number of cycles in band can be estimated, as per [146], using

$$\mathcal{N}_{\text{cyc}} = \frac{1}{32\pi^{8/3}} \left(\frac{GM_c}{c^3} \right)^{-5/3} \left(f_{\text{min}}^{-5/3} - f_{\text{max}}^{-5/3} \right). \quad (7.1)$$

7.2.1 GR binary

Using the standard set of binary parameters (recall we use a slightly reduced set, see Section 6.2.3) we first determine the parameter constraints afforded by the different detectors and detector networks that are being considered (see Table 7.2). To do this, we introduce a type of plot that will be frequently used during this and the following chapters that is useful for visually comparing parameter constraints between networks. Two quantities are required for these plots: the standard deviations of the marginal distributions for each posterior, which essentially describes the precision (low standard deviation \Leftrightarrow high precision), and what we refer to as the *precision to accuracy ratio* (PAR) which is

Detector Network	SNR	Total Spanned Observed Duration	f_{\min}	f_{\max}	Total Spanned Cycles
L_s^4	3.1	4 y	16.2 mHz	65.5 mHz	$\sim 3.0 \times 10^6$
L_s^{10}	5.7	10 y	11.6 mHz	65.5 mHz	$\sim 5.6 \times 10^6$
L_g	14.5	~ 3 s	~ 20 Hz	~ 350 Hz	$\mathcal{O}(100)$
L_s^4, L_g	14.8	4.1 y	16.2 mHz	~ 350 Hz	$\sim 3.4 \times 10^6$
L_s^{10}, L_g	15.6	10.1 y	11.6 mHz	~ 350 Hz	$\sim 6.0 \times 10^6$
D	10,575.3	11.8 days	0.1 Hz	13.95 Hz	$\sim 1.6 \times 10^5$
D, L_g	10,575.3	11.8 days	0.1 Hz	~ 350 Hz	$\sim 1.6 \times 10^5$
L_s^4, D	10,575.3	4.1 y	16.2 mHz	13.95 Hz	$\sim 3.4 \times 10^6$
L_s^4, D, L_g	10,575.3	4.1 y	16.2 mHz	~ 350 Hz	$\sim 3.4 \times 10^6$

Table 7.2: Details of the ‘low-mass’, isolated fiducial system signal in the various detector networks. Recall the detector shorthands in the leftmost column from Table 2.1.

useful for highlighting bias. Define the PAR for the *average* posterior (i.e., that obtained by using a zero-noise realisation in the data) for the i^{th} parameter as:

$$\text{PAR}_{\theta_i} \equiv \frac{\Delta_{\theta_i}}{\sigma_{\theta_i}}, \quad (7.2)$$

where σ_{θ_i} is the standard deviation of the i^{th} parameter’s marginal, and Δ_{θ_i} is the *mismatch*, the difference between the marginal’s mean value and the true injection value (low mismatch \Leftrightarrow high accuracy). Thus for a PAR of less than 1, the parameter marginal is more accurate than it is precise: the truth is within 1 standard deviation of the mean. For a PAR greater than 1, the marginal is more inaccurate than it is precise: the truth is not within 1 standard deviation of the mean. In this latter case, the posterior cannot be relied upon to constrain that parameter; this is a case of bias. This occurs rarely, but for those cases in which the PAR is greater than 1, the markers on the plot will be highlighted in red, as in, for example, Figure 7.3. The PAR is an indicator of bias rather than a strict threshold, and can be less meaningful for, for instance, cyclic parameters or parameters that have symmetric modes, and can be ignored for parameters that are to be marginalised.

The standard deviations for the average (zero-noise) low-mass GR binary posteriors for the different detector networks are given in Table 7.3, and the PARs are given in Table

7. MULTI-BAND GW ASTRONOMY PART I:
ISOLATED BINARIES

Detectors	$\sigma_{\mathcal{M}_c}$	σ_q	$\sigma_{\chi_{\text{eff}}}$	σ_d	$\sigma_{\theta_{jn}}$	σ_ψ	σ_{t_c}	σ_{α_\oplus}	σ_{δ_\oplus}
L_s^4	∞	∞	∞	∞	∞	∞	∞	∞	∞
L_g	9.4×10^{-1}	1.6×10^{-1}	7.8×10^{-2}	137	0.494	0.941	5.4×10^{-3}	$1.7 \times 10^{+0}$	5.7×10^{-1}
L_s^4, L_g	1.6×10^{-4}	1.1×10^{-1}	8.6×10^{-3}	83.1	0.259	0.898	3.0×10^{-4}	9.1×10^{-3}	1.3×10^{-2}
D	2.9×10^{-6}	1.7×10^{-3}	3.5×10^{-5}	78.8	0.75	0.905	9.4×10^{-5}	1.5×10^{-6}	1.8×10^{-6}
D, L_g	2.7×10^{-6}	1.7×10^{-3}	2.9×10^{-5}	54.2	0.243	0.911	1.3×10^{-4}	1.4×10^{-6}	1.5×10^{-6}
L_s^4, D	2.8×10^{-6}	1.8×10^{-3}	2.8×10^{-5}	77.9	0.744	0.907	1.6×10^{-4}	1.6×10^{-6}	1.6×10^{-6}
L_s^4, D, L_g	2.3×10^{-6}	1.4×10^{-3}	2.5×10^{-5}	53.7	0.239	0.895	1.2×10^{-4}	1.2×10^{-6}	1.6×10^{-6}

Table 7.3: Standard deviations of the marginals of posteriors from zero-noise realisation data, for the different detector networks, using the isolated low-mass system. The ∞ values are where the (uniform) prior is returned (i.e., there is little/no information on the parameter, due to low SNR and/or weak coupling of the parameter to the model).

Detectors	$\text{PAR}_{\mathcal{M}_c}$	PAR_q	$\text{PAR}_{\chi_{\text{eff}}}$	PAR_d	$\text{PAR}_{\theta_{jn}}$	PAR_ψ	PAR_{t_c}	$\text{PAR}_{\alpha_\oplus}$	$\text{PAR}_{\delta_\oplus}$
L_s^4	∞	∞	∞	∞	∞	∞	∞	∞	∞
L_g	3.0×10^{-1}	5.2×10^{-1}	2.7×10^{-1}	0.621	0.92	0.842	5.2×10^{-2}	6.0×10^{-1}	7.3×10^{-1}
L_s^4, L_g	1.4×10^{-1}	5.6×10^{-2}	2.9×10^{-1}	0.297	0.149	0.72	1.0×10^{-1}	1.2×10^{-2}	2.5×10^{-1}
D	5.3×10^{-4}	6.8×10^{-3}	2.0×10^{-3}	0.684	1.17	0.702	5.1×10^{-3}	2.5×10^{-2}	2.9×10^{-2}
D, L_g	6.1×10^{-3}	1.4×10^{-3}	6.4×10^{-3}	0.124	0.186	0.7	1.4×10^{-3}	1.2×10^{-2}	8.0×10^{-3}
L_s^4, D	1.4×10^{-2}	1.9×10^{-2}	1.6×10^{-2}	0.703	1.2	0.725	6.5×10^{-3}	1.3×10^{-2}	1.5×10^{-2}
L_s^4, D, L_g	4.4×10^{-3}	1.7×10^{-5}	3.2×10^{-3}	0.081	0.142	0.69	5.1×10^{-3}	1.8×10^{-2}	6.6×10^{-3}

Table 7.4: The ‘precision-to-accuracy’ ratios (PARs) for the marginals of the low-mass, isolated GR binary. There is little evidence of bias here; generally the PAR is less than 1 (aside from some of the cyclic parameters, for which the PAR does not hold much meaning). For all subsequent signal studies, the precision and accuracy shall instead be presented as *marginal precision and accuracy (MPA)* plots, where the markers for parameters with $\text{PAR} > 1$ shall be displayed in red (black otherwise), as in Figure 7.3, for example.

7.4. The corresponding *marginal precision and accuracy (MPA)* plot is shown in Figure 7.1, confirming the expectation that the addition of LISA to LIGO strongly affects constraints of the chirp mass, effective spin and sky position. The mass ratio is no more precise, but around a factor of 5 more accurate (since all posteriors produced here are done so using a zero noise realisation it is appropriate to talk about accuracy). Surprisingly, LISA improves the coalescence time standard deviation by around a factor of 10; this appears to be due to the suppression of a secondary mode across a broad streak on the sky localisation in the LIGO posterior, which allows for a spread of probability in t_c . When DECIGO is included in the detector network, again as expected, almost all parameters become far more tightly constrained. The distance d is very strongly correlated with the orientation however, since these parameters both strongly inform the amplitude, they thus do not acquire particularly strong restrictions by including DECIGO. However, since sky position is well constrained by DECIGO, we can expect far better chances of host galaxy identification, with the possibility of luminosity distance then being given by galaxy catalogues, or being optically measurable [148].

It is particularly enlightening to focus on the LISA+LIGO network and consider the behaviour of the posterior as one slowly increases the sensitivity of LISA (from ‘zero sensitivity’). That is, we will ‘override’ the LISA PSD, defining a new detector with PSD given by

$$S'_n(f) = S_n(f) \times \left(\frac{\rho_{\text{opt}}}{\rho_*} \right)^2, \quad (7.3)$$

where $S_n(f)$ is the original PSD, ρ_{opt} is the injected signal’s SNR in LISA given the original PSD, and ρ_* is some target SNR of the signal: the SNR of the signal given the new detector PSD. Then, as a function of the target SNR ρ_* (the new PSD is not strictly defined for $\rho_* = 0$, so note that we take $\rho \rightarrow 0$ where necessary) we produce the posteriors using LISA(ρ_*)+LIGO and plot their marginals in Figure 7.2.

This highlights some interesting features of the behaviour of the multi-band posterior. Firstly, we see where the LISA part begins to contribute to the parameter constraints and the process of the LISA posterior ‘freezing out’. Observe that the chirp mass (and effective spin, which is fairly strongly correlated with chirp mass) freezes out first in Figure 7.2, when $\rho_* = 2.0$. The chirp mass governs the leading-order contribution to the GW frequency [27], thus asking what the primary frequency component (the frequency with largest amplitude) of the LISA signal is, is similar to asking what the chirp mass is. This

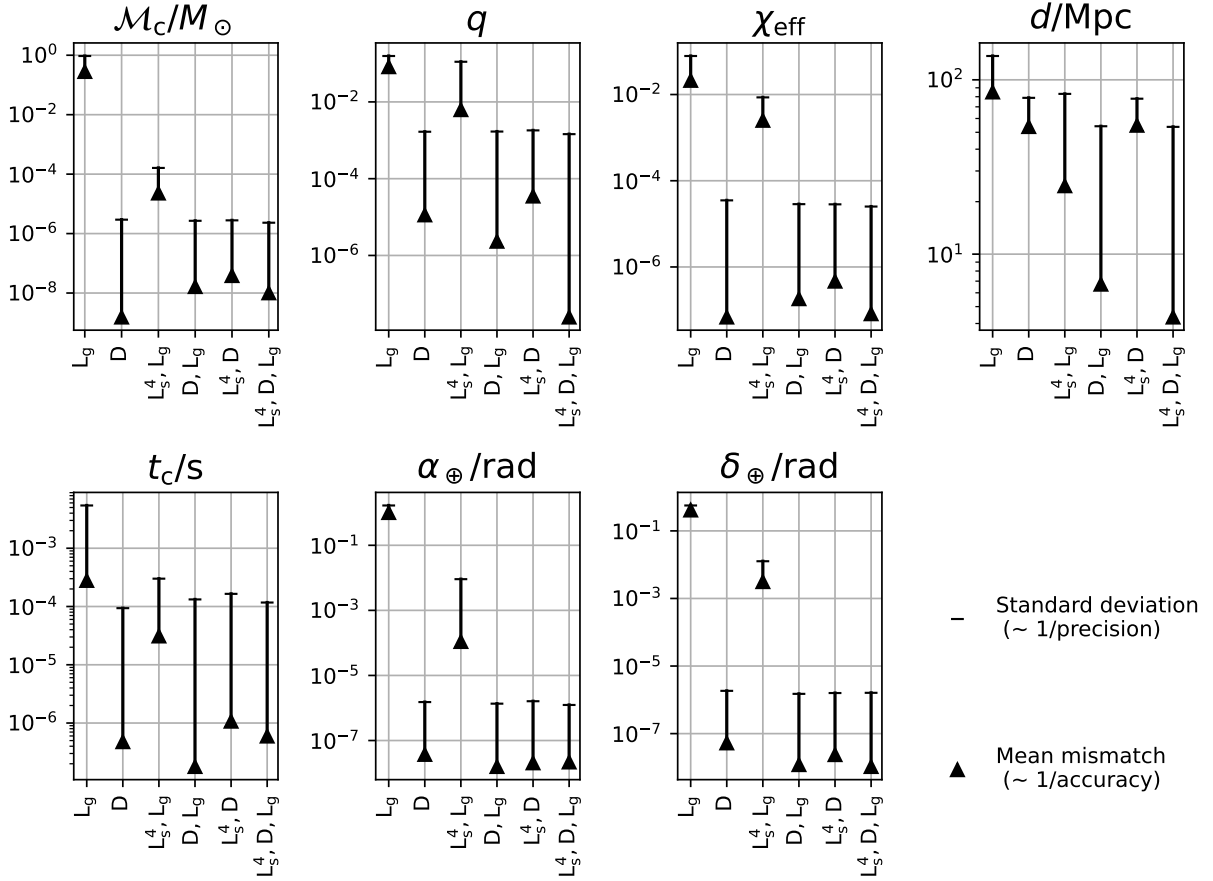


Figure 7.1: The MPA plots for the low-mass, isolated GR binary. Immediately we notice that DECIGO dominates the parameter constraints, and that there is not much additional help in constraining parameters from the other detectors for a network including DECIGO, aside from the distance. There are a few curious features: in particular, the coalescence time is better constrained by DECIGO alone, rather than when cooperating with other detectors. This could be due to intricacies in the interplay of the likelihoods from different detectors (see discussion of Figure 6.2 & Figure 7.2). We are working with noisy processes, however, and subtle results like this could be accounted for by sampler noise or downsampling noise. The LISA plot is omitted since the signal is not detected.

is despite the fact that the frequency is chirping, since: we have a strong constraint on the temporal location of the LISA signal in the inspiral from the knowledge of coalescence time from LIGO, and the leading-order time derivative of the frequency, $\dot{f} = \dot{f}(t)$, is already fixed by the instantaneous frequency and the chirp mass, as given in equation (5.39). Expressed another way, consider one particular time sample in the LISA signal, with time t_1 . Since we know t_c from LIGO, then by integrating (5.39), $f(t_1)$ is a function of \mathcal{M}_c , and by (5.39), $\dot{f}(t_1)$ is a function of $f(t_1)$ and \mathcal{M}_c (and so ultimately just of \mathcal{M}_c), hence given t_c , \mathcal{M}_c and $f(t)$, then $\dot{f}(t)$ is completely determined. Given a coalescence

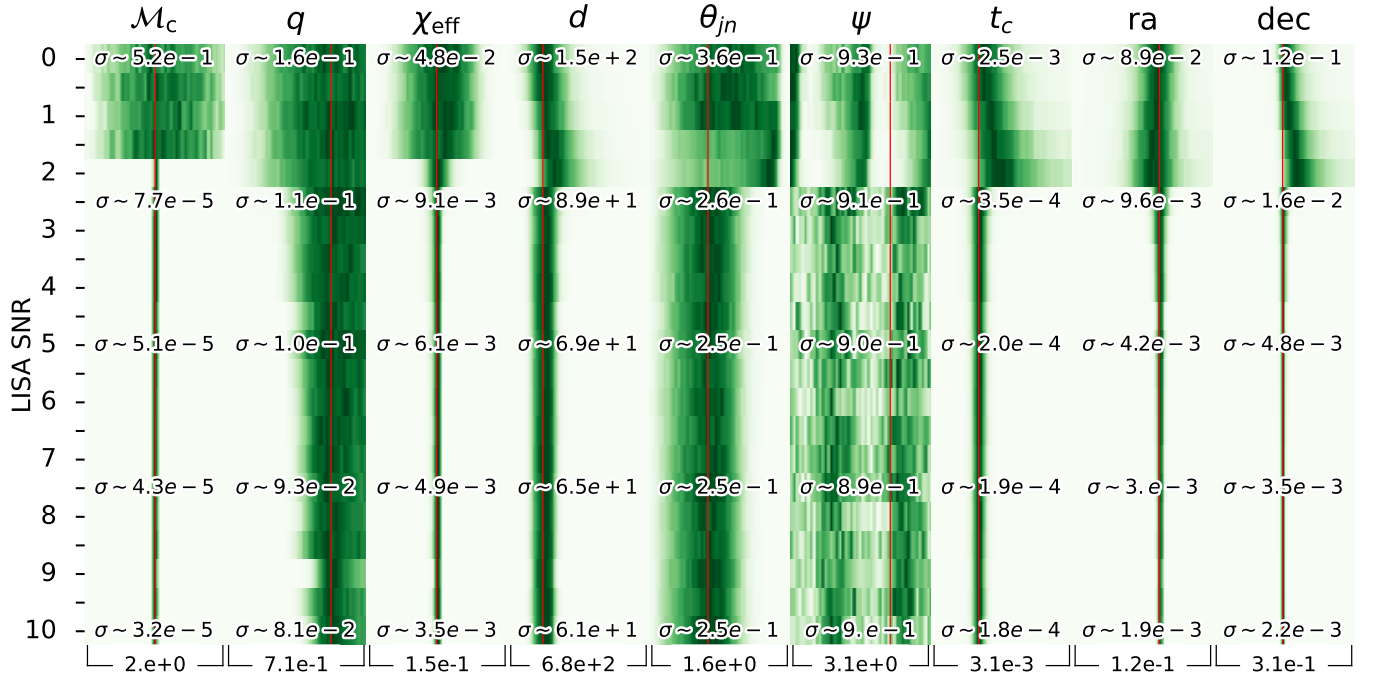


Figure 7.2: Marginals of the (zero-noise) posterior as a function of the SNR in LISA, ρ_* , for the low-mass system in LISA $_{\rho_*}$ +LIGO. Each row is a set of histograms displaying marginals of a posterior, obtained by binning the posterior samples appropriately. The columns therefore show the evolutions of parameter marginals, as the posteriors are defined with, in this example, increasing LISA SNR. The red lines show the injection parameters. At $\rho_* = 0$, the posterior is equal to the LIGO only posterior. We surmise that, due to the ‘incorrect’ LIGO marginal on the orientation, ψ , the primary, ‘correct’ posterior mode in LISA containing the signal parameters is suppressed (note that the LIGO posterior can be thought of as a prior for LISA). Thus as one increases ρ_* , the LISA signal must be accounted for in some way, and this appears to be achieved by the visible shifts in the distance, coalescence time, and sky parameters. The chirp mass and effective spin (which strongly inform the primary frequency component of the signal) freeze out when $\rho_* \approx 2$ as the LISA signal becomes loud enough to override the LIGO suppression, and the other parameters are located as $\rho_* \approx 2.5$.

time, the *instantaneous* frequency and its derivative at another earlier time are entirely dependent on chirp mass. But the frequency at all other times is determined by an instantaneous frequency at any one given time, therefore, asking what the instantaneous frequency at a time t_1 is, to repeat the original claim, is the same as asking what the chirp mass is (given that t_c is known).

Now because the two aforementioned questions are similar, it is reasonable to apply the arguments of Bretthorst [94] on the resolving power of the primary frequency component of a signal to the resolving power of the chirp mass. Bretthorst lays out a number of conditions required in order to make his arguments about resolving the frequency, which are not met with our data, but the situation here nevertheless remains roughly similar. The conclusion reached by Bretthorst is that:

"[...] if you have sufficient signal-to-noise for signal detection more data are important for resolution; otherwise more signal-to-noise will detect the signal with less data."

This is because the evidence of a (single frequency component) signal (i.e., claiming a detection) strongly depends on the SNR, whereas the constraints on the error estimate depend more strongly on the number of samples and duration of the signal. We see this sort of behaviour in Figure 7.2; an increase in SNR is required until the point of detection, at which point there is a very rapid change in σ (since the ‘height’ of the peak of the posterior is similar to the exponential of ρ_*^2 , see Ref. [94]), and after which we see that $\sigma \sim 1/\sqrt{\rho_*}$.

Secondly, we notice further interesting features of Figure 7.2: the posterior appears to ‘become worse’ as the SNR increases, at very low SNRs in LISA. The LISA signal is too quiet for the LISA posterior mode containing the injection parameters to contain a significant enough proportion of the probability to have noticeable impact on the multi-band posterior. The marginal of the orientation, ψ , remains roughly constant until $\rho_* = 2.5$, equal to the LIGO only marginal (where $\rho_* = 0$). As the SNR is increased, and while the LISA signal still provides insufficient information, it appears that the gradual shift in the other parameters away from the truth is required to account for the extra signal power, whilst the LIGO posterior is providing strong constraints on the values of ψ , that is, until the LISA signal becomes strong enough to supersede LIGO’s preference at $\rho_* = 2.5$. This is an important feature to be aware of in real-world analysis.

7.2.2 Quadratic gravity

In this section we will investigate the effect on the posterior of introducing the quadratic gravity coupling constants $\zeta_{\text{EdGB,dCS}}$ as free parameters. Recall that in our derivation of the quadratic gravity GW model, we adopted the small β approximation. As per [128], the quadratic gravity models presented are valid in the small coupling approximation: this occurs when $\zeta_{\text{EdGB,dCS}} < 1$, however, we obtained the waveform in equation (5.45) by the small parameter limit, demanding that $\beta_{\text{EdGB,dCS}} \ll 1$ and thus that $\zeta_{\text{EdGB,dCS}} \ll 1$. We therefore require posteriors in which all of the probability for the parameter $\zeta_{\text{EdGB,dCS}}$ is contained within a small region centred at zero in order for the posterior to have been evaluated within the valid range for the approximant.

7.2.2.1 Einstein-dilaton Gauss-Bonnet

We now produce and examine the precision and accuracy plots for marginals given the EdGB gravity model with parameters given in Table 7.1. This is shown in Figure 7.3. One immediate feature of interest to point out is that the coalescence time constraint appears to *improve* by adding this extra parameter. We will notice this particular phenomenon recurring when considering acceleration parameters, too: as per Section 5.6, recall that the EdGB parameter and a constant acceleration are correlated, so that some non-zero EdGB coupling can be confused with a non-zero acceleration. A ‘one-off’, small improvement in a given posterior should not cause too much furrowing of the brow since noisy processes are employed for obtaining these posteriors and some variance is expected, however this improvement appears to be a consistent feature across the set of results.

This is somewhat unusual as one ordinarily expects all parameters to either retain the same or acquire larger error when introducing an extra parameter, amounting to more uncertainty in the constraints from the extra freedom in the model. However, it is also very easy to imagine a probability distribution, \mathcal{D} , from which a new distribution defined by marginalising, \mathcal{D}'_{m} , can have lower variance than one defined by taking a ‘slice’, \mathcal{D}'_{s} . Consider, for example, that \mathcal{D} is a unit disk of uniform probability, centred at the origin of the x, y -plane (and zero probability outside of the disk). The $y = 0$ slice, \mathcal{D}'_{s} , yields a top-hat function, whereas the marginal over y , \mathcal{D}'_{m} , is given by $2\sqrt{1-x^2}$ (both zero outside of the range from -1 to 1) the latter thus clearly having lower variance despite having

been more obviously obtained from a higher dimensional posterior. It is possible that a similar situation is occurring with the coalescence time from the inclusion of the extra EdGB gravity parameter (with \mathcal{D}'_s corresponding to the ‘GR-only’ model with $\zeta_{\text{EdGB}} := 0$, and \mathcal{D}'_m corresponding to the free ζ_{EdGB} model).

We can also see from Figure 7.3 that, in general, there are tight constraints on ζ_{EdGB} , particularly from DECIGO. Note that the posterior railed up against uniform prior bounds for ζ_{EdGB} in detector network L_s^4, L_g . The posteriors are particularly troublesome to plot (see Figure 7.4); the ζ_{EdGB} parameter marginal appears as a ‘spike’ at the injection parameter that drops away to a non-zero ‘platform’, but as one increases the (uniform) prior bounds, the platform remains at some non-zero constant. Thus in continuing to widen the priors, there eventually becomes relatively very little probability left in the spike. The majority of probability resides in the platform and the posterior slowly begins to resemble the prior (whilst simultaneously causing significant sampling difficulty resulting in very long waits for convergence). However, in Section 5.5.1 we discussed the fact that we can place reasonably stringent priors on ζ_{EdGB} from the results of Ref. [128] (the $\sqrt{\alpha_{\text{EdGB}}}$ parameter is assumed common between all GW events). Recall that the constraints placed on the dimensionless parameter for this system of $\zeta_{\text{EdGB}}^{\text{low-mass}} \lesssim +3.62 \times 10^{-6}$. The posterior we used is therefore similar, but not entirely in keeping with that which would have been obtained given more appropriate priors.

Using only DECIGO, the constraints are strong enough from just one observation to give an improvement on the results from Ref. [128] by a factor of 10. Using LISA alongside DECIGO improves the DECIGO constraints by a further factor of 10. This is despite the fact that generally observing around the merger (in the highly-non-linear regime) is best for constraining higher-order gravity theories, and can be credited to the extremely high sensitivity of DECIGO. Its ability to pin down parameters strongly suppresses subsidiary modes in the posterior, such as those obvious in Figure 7.4. The analysis of data from other detectors, such as LISA, is prevented from returning strongly biased regions of likelihood since they are essentially excluded. It appears that DECIGO removes the ‘platform’, at which point the ‘spike’ in the ζ_{EdGB} parameter becomes highly effective at constraining the parameter.

It is clear from Figure 7.4 where the $\text{PAR}_q > 1$ in Figure 7.3 comes from: the ‘platform’ of ζ_{EdGB} , which contains most of the probability, corresponds to the region at $q \rightarrow 1$,

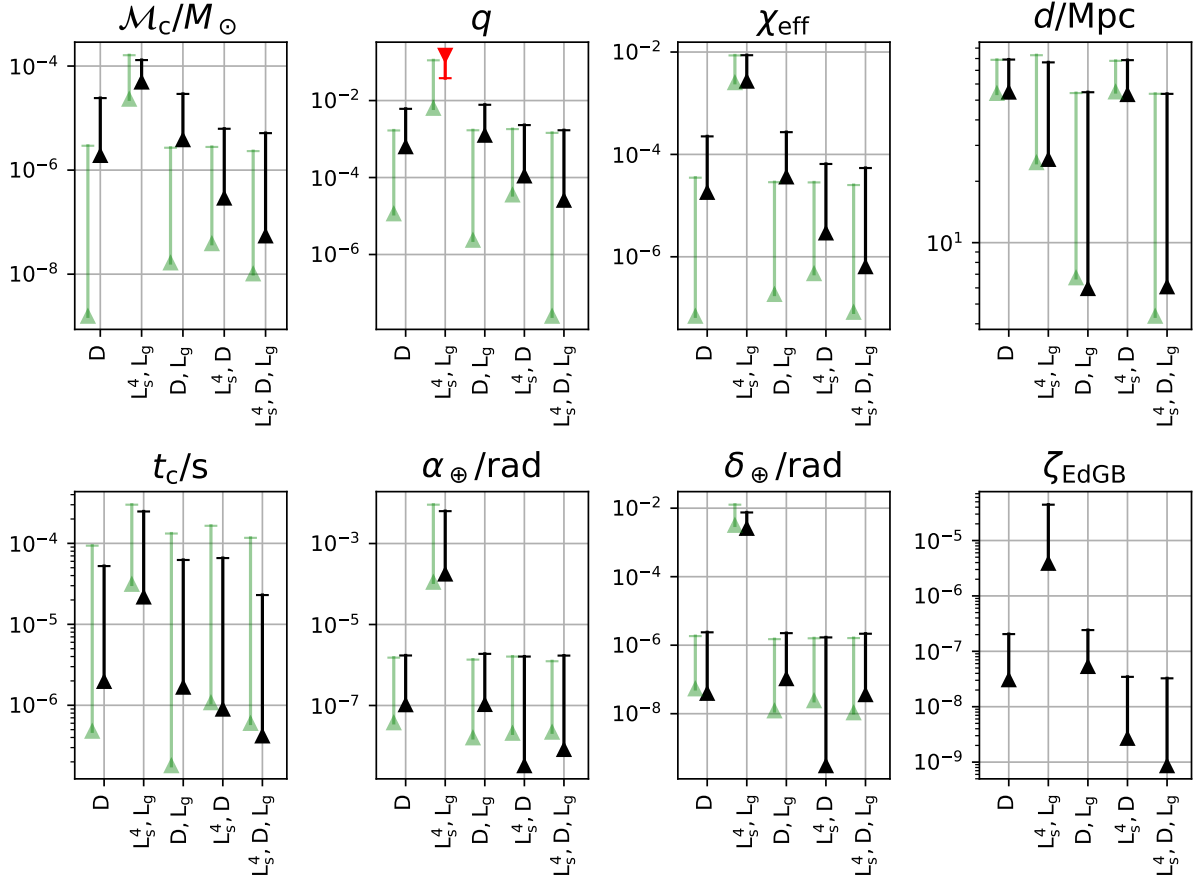


Figure 7.3: The MPA plots for the low-mass, isolated EdGB binary. The posterior in ζ_{EdGB} railed up against the uniform prior bounds for LISA+LIGO (see Figure 7.4), however, the priors are not unreasonable since they are roughly equal to those constraints given by Ref. [128]. The the GR model MPA plots (i.e., Figure 7.1) are replicated (green), to show how the constraints vary by the extension of the model. The red mark points out where the PAR is greater than 1, indicating bias (see Figure 7.5 for more information on this). Interestingly, we see that adding the free ζ_{EdGB} parameter generally *decreases* the standard deviation on t_c ; one usually expects all parameters to become less well constrained by extending the model, but this can however be explained by intricacies of the distribution, as explained in Section 7.2.2.1.




Figure 7.4: Corner-plot for low-mass EdGB binary in LISA+LIGO, marginalised to $(\mathcal{M}_c, q, \chi_{\text{eff}}, \zeta_{\text{EdGB}})$. The ζ_{EdGB} constraints come from LISA only, as the quadratic gravity modifications are not modelled here in LIGO. The injected value of ζ_{EdGB} is zero, thus the waveform is a purely GR waveform and does not include negative PN terms in the phase. There is strong preference for positive only PN terms in the phase. Since EdGB gravity enters the waveform at -1PN , we find a sharp peak at $\zeta_{\text{EdGB}} = 0$, however, one obtains a waveform with positive only PN terms in the phase when $\beta_{\text{EdGB}} = 0$, which can be achieved either by $\zeta_{\text{EdGB}} = 0$, or by having $q = 1$ (see Figure 5.1). This is why a strong bias for $q = 1$ is evident.

where a strong bias is introduced. Due to our restriction to aligned spin systems, equation (5.11) shows that the EdGB modifications vanish identically when $q = 1$, thus at $q = 1$, the -1PN EdGB term vanishes from the waveform and the remaining (non-negative only) PN terms for the GR signal are accounted for by the remaining free parameters which thus undergo a shift from the true parameters, given the incorrect $q \approx 1$. This explains the platform in ζ_{EdGB} (when $q = 1$, ζ_{EdGB} is completely decoupled from the waveform, so the uniform prior on ζ_{EdGB} is returned).

In Figure 7.5, we plot the LISA+LIGO marginals as a function of LISA SNR, ρ_* . We have wider priors on ζ_{EdGB} here, and again we reiterate that EdGB is not modelled for the LIGO part; all the EdGB information comes from LISA. We see a few distinct transitions as ρ_* is increased. The first occurs at $\rho_* \approx 4$, where the LISA signal initially emerges in the posterior, prior to which its presence affects the posterior with similar behaviour as that we saw in Figure 7.2. The strong preference for $q = 1$ starts to weaken at $\rho_* \approx 30$, where the platform on ζ_{EdGB} starts to drop significantly. Finally, for LISA SNRs in the hundreds (where the LIGO contribution becomes insignificant) we begin to see more erratic behaviour (note in these plots the range on ζ_{EdGB} is much tighter).

This is likely due, in part at least, to sampling issues. As depicted in Figure 7.6, the heavily marginalised posterior for which $\rho_* = 300$, we can see that the ζ_{EdGB} parameter generally forms an ‘L’-shaped distribution, as viewed from the various biaxial projections. This mode is very thin with a sharp turn/corner, which may be difficult for some samplers to handle. In particular, since `Nessai` tries to learn iso-likelihood contours and generate a map to a standard Gaussian distribution (within which it is easy to propose new samples), it is possible that a difficulty occurs in ‘turning the corner’ and following the contours (especially in high dimensions) and that only one of the arms of the ‘L’ are converged upon. This is one example situation in which significant problems may occur in sampling and which have the potential to cause the return of unreliable posteriors.

Such large LISA SNRs are not reasonable for stellar order binaries at cosmological distance of course. However, there may very well exist lighter binaries in DECIGO with similar SNRs and numbers of cycles, and which thus provide posteriors with the same sort of structure, leading to these sampling problems. This issue must be addressed at the sampler level, since given the vast size of the LISA and DECIGO datasets, it is not practical to, for example, rerun the sampler many times to ensure the convergence is

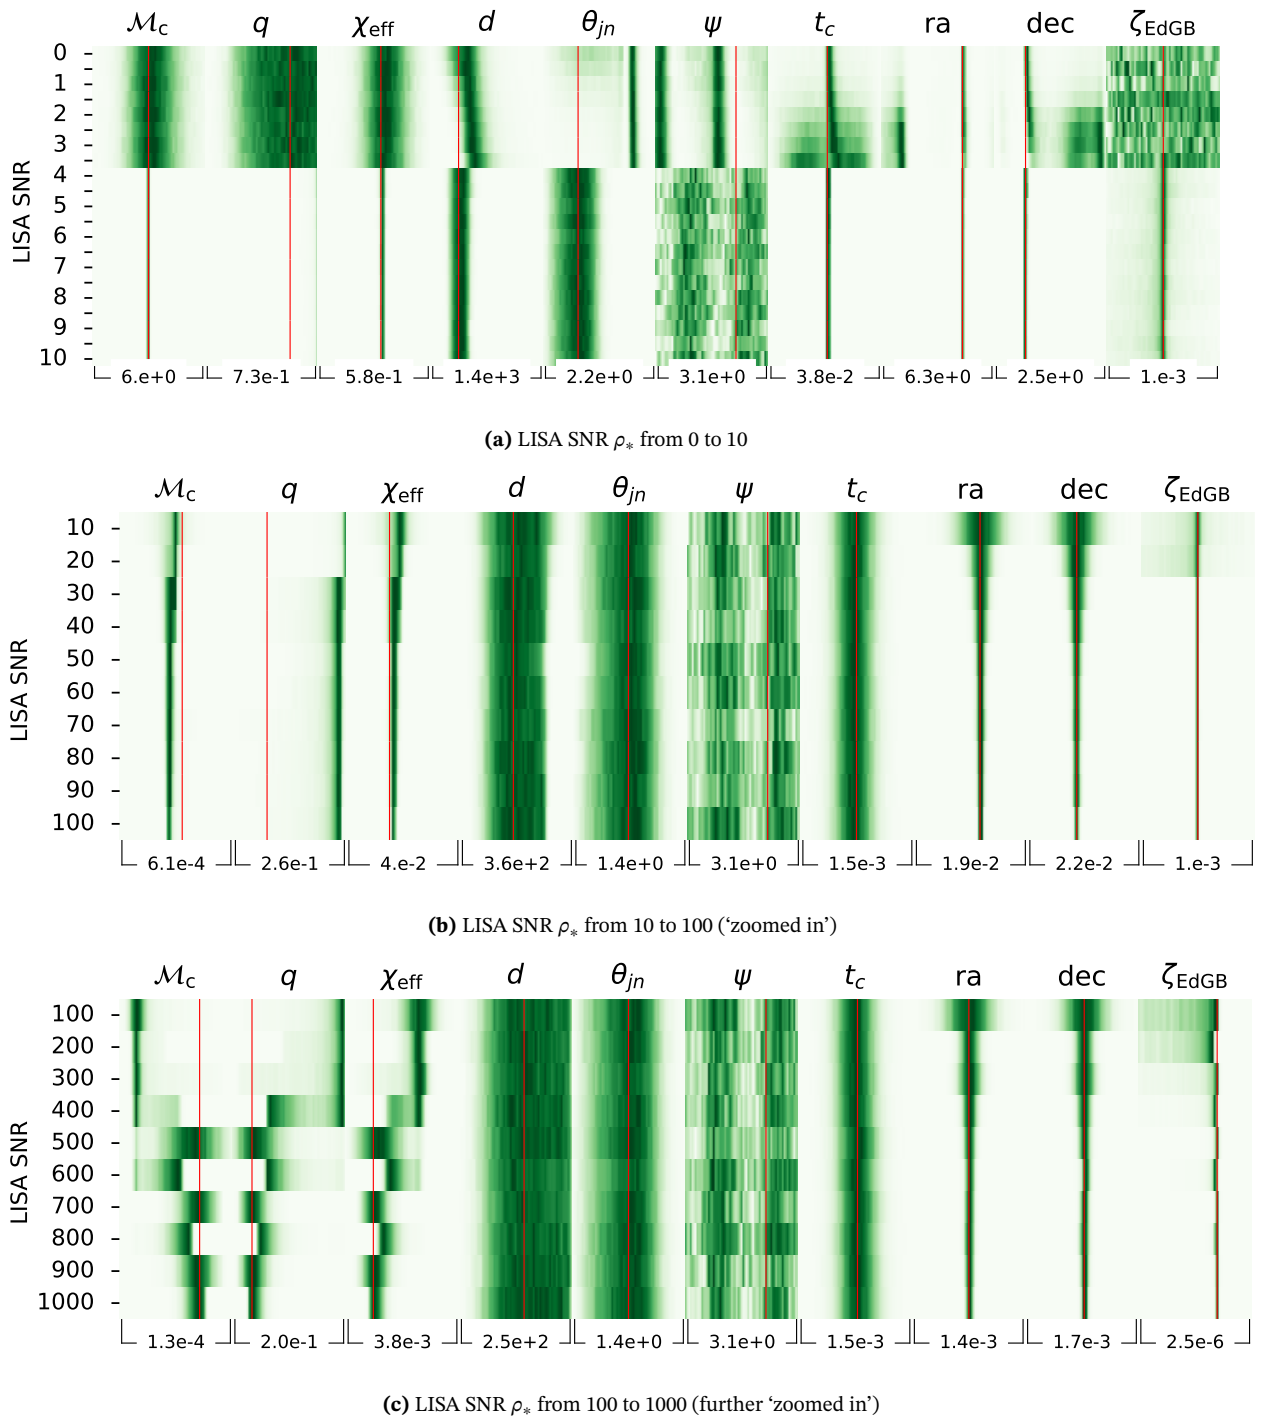


Figure 7.5: Detailed evolution of the low-mass EdGB binary in LISA(4 year)+LIGO, as a function of the LISA SNR. Firstly, notice that the extension of the model with another parameter increases the LISA SNR detection threshold to $\rho_* \approx 4$ (compare to Figure 7.1). The bias remains in the parameters (particularly in q) however, until ρ_* reaches the hundreds. At $\rho_* \approx 400$, we see the primary mode (containing the correct q) begin to appear, from which point there appears to be an issue with sampling, which is made clear in Figure 7.6, where it appears that one or the other ‘arms’ of the L-shaped modes are selected by the sampler. Around $\rho_* \approx 900$, the SNR finally appears high enough to overcome both bias and sampling issues.




Figure 7.6: As in Figure 7.4, this is the corner-plot for the low-mass EdGB binary in LISA+LIGO, marginalised onto $(\mathcal{M}_c, q, \chi_{\text{eff}}, \zeta_{\text{EdGB}})$, but with LISA SNR $\rho_* = 300$. Here we can see more detail of the bias and the probable cause of sampling issues, with a posterior mode which is both very thin and with a sharp turn. The chirp mass parameter $\mathcal{M}'_c \equiv \mathcal{M}_c - \hat{\mathcal{M}}_c$, where $\hat{\mathcal{M}}_c = 30.4 M_\odot$ is the injection parameter.

reliable. Note that it remains possible that a more complete waveform model with sub-leading order PN terms in ζ_{EdGB} will break the $q \rightarrow 1$ bias.

Finally, we point out that whilst (this single detection in) the LISA+LIGO network is not able to strongly constrain EdGB gravity, the inclusion of DECIGO in the network brings the standard deviation to $\sigma_{\zeta_{\text{EdGB}}} \approx 2 \times 10^{-8}$, corresponding to improving the current constraint on the (dimensionful) coupling parameter of $\sqrt{\alpha_{\text{EdGB}}} \lesssim 1.7 \text{ km}$ to $\sqrt{\alpha_{\text{EdGB}}} \lesssim 0.77 \text{ km}$ at 90% confidence. The multiple systems we can expect to observe (including ‘golden’ sources) will improve this further.

7.2.2.2 Dynamical Chern-Simons

The analysis of the dCS modified gravity posteriors is more straightforward than that of the EdGB gravity case; the structure of the posterior is considerably simpler, with even the standard LISA+LIGO network producing a ‘nice’, roughly Gaussian structure. The form of β_{dCS} in equation (5.10) in terms of the masses and spins of the binary components makes it less susceptible to allowing biased results in the way the EdGB parameter does, as we saw in the previous section. This is in part because $\beta_{\text{dCS}} \rightarrow 0$ as $\chi_{\text{eff}} \rightarrow 0$ (see Figure 5.1) but unlike the mass ratio, q , which is usually not so well constrained, the effective spin, χ_{eff} , can be well constrained (as in Figure 7.1, for instance), suppressing the potential biased preference of $\chi_{\text{eff}} \rightarrow 0$ for highly-spinning systems with $\beta_{\text{dCS}} \approx 0$. However, perhaps more crucially, the dCS modification enters the phase at 2PN. The 2PN GR phase coefficient is defined in terms of the masses and spins, so the 2PN dCS phase term effectively modifies this coefficient and can more freely mix with the masses and spins. One can indeed see in Figure 7.8 that including ζ_{dCS} as a free parameter generally increases the standard deviations of the masses and spins. The EdGB term, on the other hand, enters the phase at -1PN which has no GR counterpart, its strict absence from the phase (as per the injected values) is therefore both more prominent, and more limited in ways available to account for it.

In Figure 7.7, the marginals as functions of the LISA SNR again shows the same behaviour for very low SNRs, until the signal is loud enough to begin dominating the LIGO likelihood. At that point however, as the LISA constraints become apparent, we see that ζ_{dCS} itself is not strongly constrained (note the total range of ~ 400 for ζ_{dCS} in the figure). The standard deviation holds roughly constant at a value of ~ 30 (see Figure 7.8) as




Figure 7.7: Marginals of the posterior for the low-mass dCS binary in LISA+LIGO as a function of the SNR in LISA. We can see similar behaviour as that evident in Figure 7.2; the ‘correct’ posterior mode freezes out as the SNR in LISA approaches $\rho_* \approx 3$. The SNR detection threshold is less than for the EdGB gravity extension, indicating that the dCS parameter has a weaker effect on the number of waveform templates than EdGB does.

$\rho_* \rightarrow 10$. This does not comply with the small coupling approximation which demands that $\zeta < 1$ and therefore these posteriors are not usable.

This is not the case when DECIGO enters the network however; we see in Figure 7.8 that there exist posteriors for which $\sigma_{\zeta_{\text{dcs}}} < 1$ and some of these posteriors are valid. With DECIGO on its own, we have $\sigma_{\zeta_{\text{dcs}}} \approx 1$, where quite clearly there will exist samples in the posterior for which $\zeta_{\text{dcs}} > 1$, however, in conjunction with another detector, we have that $\sigma_{\zeta_{\text{dcs}}} \lesssim 0.3$ and the posterior samples mostly quite comfortably sit within the small coupling approximation limit. This is unfortunately not at all stringent enough (with this single observation), however, to improve the current constraints of $\sqrt{\alpha_{\text{dcs}}} \lesssim 8.5$ km; we instead find that $\sqrt{\alpha_{\text{dcs}}} \lesssim 47.4$ km at 90% confidence. To do better than the current best constraint, one would require (at 1-sigma) that $\sigma_{\zeta_{\text{dcs}}} \lesssim 2.27 \times 10^{-3}$ from this binary, using equation (5.6). Thus constraining the dCS parameter with low-mass binaries would require either a large number of detections with a detector network including DECIGO (around $10^4 k^2$ to reach k -sigma), or more sensitive ground-based detectors to capture their merger in fine detail, where the modification to the waveform is greatest.

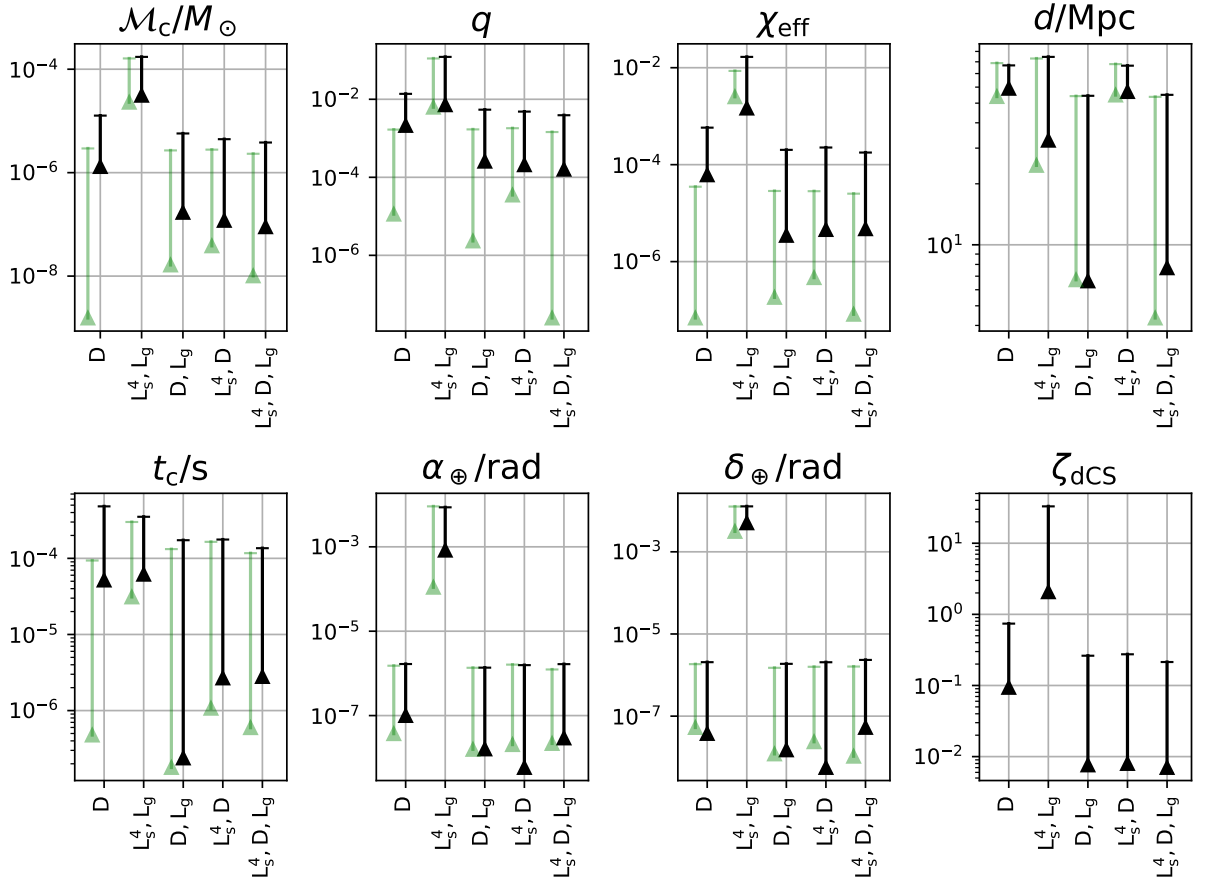


Figure 7.8: Marginal precision and accuracy plots for the low-mass dCS binary. Extending the GR model to dCS most significantly reduces the constraints one can place on the effective spin. The mass constraints are slightly weakened, and other parameters unaffected. This is to be expected by the dependence on the β_{dCS} phase modification coefficient, given by equation 5.10, which is a function of the masses and spins.

Detector Network	SNR	Total Spanned Observed Duration	f_{\min}	f_{\max}	Total Spanned Cycles
L_s^{10}	13.8	10 y	3.6 mHz	20.2 mHz	$\sim 1.7 \times 10^6$
L_s^4	10.2	4 y	5.0 mHz	20.2 mHz	$\sim 9.3 \times 10^5$
L_g	6.3	~ 5 s	~ 10 Hz	~ 20 Hz	$\mathcal{O}(10)$
L_s^{10}, L_g	15.2	10.1 y	3.6 mHz	~ 20 Hz	$\sim 1.8 \times 10^6$
D	6,579.1	0.52 days	0.1 Hz	4.39 Hz	$\sim 7.0 \times 10^3$
D, L_g	6,579.1	0.52 days	0.1 Hz	~ 20 Hz	$\sim 7.0 \times 10^3$
L_s^{10}, D	6,579.1	10.1 y	3.6 mHz	4.39 Hz	$\sim 1.8 \times 10^6$
L_s^{10}, D, L_g	6,579.1	10.1 y	3.6 mHz	~ 20 Hz	$\sim 1.8 \times 10^6$

Table 7.5: Details of the ‘high-mass’, isolated GR fiducial system signal in the various detector networks. Recall the detector shorthands in the leftmost column from Table 2.1.

7.3 ‘High-mass’ binary

We now investigate the behaviours and benefits of posteriors defined by detections of our fiducial high-mass binary, details of which can be found in Table 7.1. The merger of an inspiralling binary occurs roughly at the ISCO, given by (adapted from [146])

$$f_{\text{ISCO}} \approx 4.4 \text{ kHz} \left(\frac{M_{\odot}}{m} \right), \quad (7.4)$$

where m is the total mass, and which for our high-mass system is thus equal to $f_{\text{ISCO}} \approx 10$ Hz. This just barely enters the LIGO band, so LIGO only really captures the ringdown of this signal. Note that we cut off the DECIGO signal a few seconds before ISCO, in order to ensure that the waveform approximant used is always valid, since the time-domain inspiral approximants blow up as $\tau \rightarrow \tau_c$. Using equation (7.1) and detector properties, Table 7.5 shows the details of this signal in the various detector networks.

7.3.1 GR binary

As in the previous section, we compute the posteriors for the standard GR, high-mass binary over all the detector network combinations, and plot their parameter-wise standard deviations in Figure 7.9. We mentioned already that only the end of the merger and

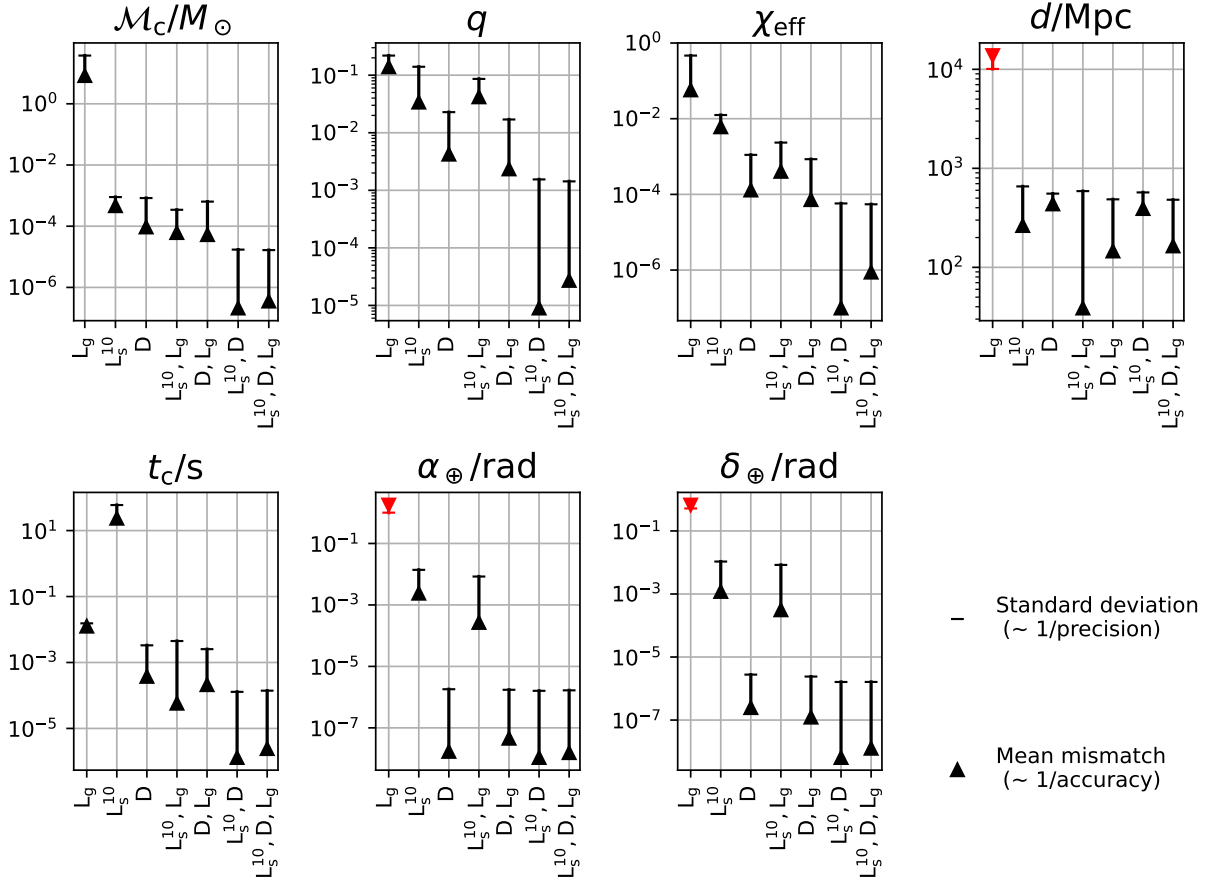


Figure 7.9: The MPA plots for the high-mass, isolated GR binary. Again, networks containing DECIGO dominate the parameter constraints. For the high-mass system with significant SNR in LISA, we now see that including LISA along with DECIGO makes a large contribution to the masses, spin and coalescence time constraints. The LIGO SNR is low (~ 6) and does not pass the detection threshold (the PARs >1 therefore are not alarming). However, there is considerable gain in the coalescence time and spin by adding LIGO to LISA.

the ringdown is caught by LIGO, due to the low-frequency at merger. The LIGO SNR is accordingly low, at around ~ 6.3 . This is why the LIGO posterior in Figure 7.9 is generally very wide, and even in the zero-noise realisation case has mean samples values offset from the true parameters by greater than 1-sigma (i.e., with PARs greater than 1). It is still a useful addition to LISA however, primarily due to the constraints it gives on the coalescence time, which selects the small region of chirp mass and effective spin that is compatible, reducing the deviation on chirp mass by a factor of a few, and on spin by around a factor of 10.

Generally speaking however, expectations are confirmed as it is evident that LIGO has a marginal role to play for very high-mass systems. The main comparison of interest

for the high-mass systems is between LISA, DECIGO, and the network of both of these detectors. Clearly, DECIGO is equal to or outperforms LISA for every parameter. Adding LISA to DECIGO does however provide significant improvements on \mathcal{M}_c , q , χ_{eff} & t_c ; by inspecting the corner plots we note that clear (anti-)correlations of $(\mathcal{M}_c, \chi_{\text{eff}})$, (\mathcal{M}_c, t_c) , (q, χ_{eff}) & (q, t_c) in LISA become correlated in the opposite sense in DECIGO.

7.3.2 Quadratic gravity

In the same form as in the previous section, we analyse the high-mass system but allowing as free parameters $\zeta_{\text{EdGB,dCS}}$. Recall the quadratic gravity modifications are not being modelled in LIGO here, however, since there is actually very little information in LIGO for the high-mass system, the posteriors obtained here shall more accurately represent the true posteriors.

7.3.2.1 Einstein-dilaton Gauss-Bonnet

The EdGB MPA plots are shown in Figure 7.10. The LIGO-only and LISA-only posteriors are omitted since the signals are not found and very little information is returned. The LISA+LIGO posterior, however, does return some useful information despite suffering from the same bias we saw in Section 7.2.2.1 for the low-mass (EdGB) case: with the injection value at $\zeta_{\text{EdGB}} = 0$, there is a vanishing -1PN term in the signal, but this can also be achieved when $q = 1$. A mass ratio of $q \rightarrow 1$ is therefore preferred, and other parameters, in particular the chirp mass, are offset to account for that preference. Besides these parameters however, the LISA+LIGO posterior is not heavily changed from the GR base model posterior (the green lines in Figure 7.10).

In terms of networks that include DECIGO, the SNR is again so high that this bias disappears. For the networks including both DECIGO and LISA, there is a very clear, large increase in the standard deviations of \mathcal{M}_c , χ_{eff} & t_c , which, by inspecting the corner plots, are found to be strongly positively correlated with ζ_{EdGB} , and q which is strongly negatively correlated. All detector networks shown in Figure 7.10 on detecting this high-mass binary produce posteriors consistent with the small coupling approximation, however, even the strongest constraint on $\sigma_{\zeta_{\text{EdGB}}}$ given by all three detectors (L_s^{10}, D, L_g) translates to a constraint on the dimensionful coupling parameter of $\sqrt{\alpha_{\text{EdGB}}} \lesssim 7.5 \text{ km}$ at 90% certainty. This is roughly $10\times$ the value we obtained from the low-mass binary, even though

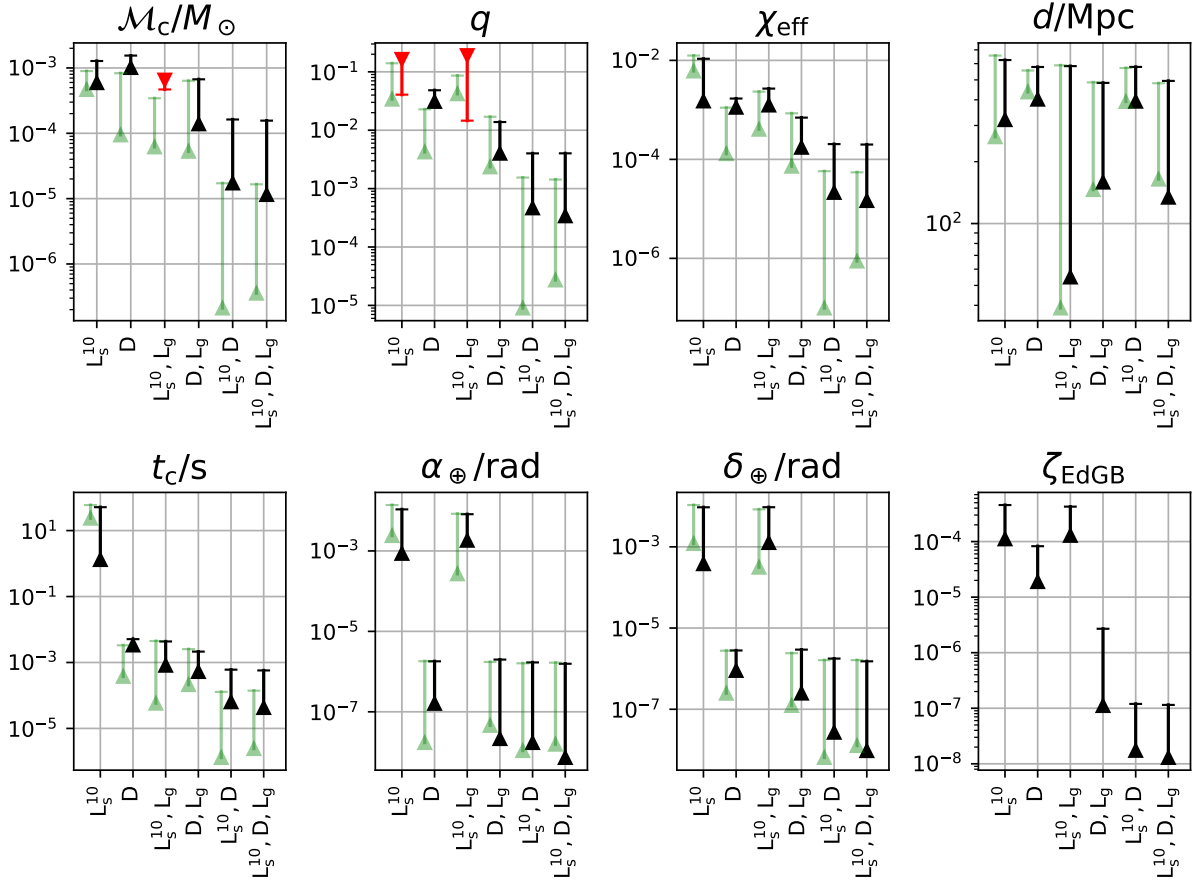


Figure 7.10: The MPA plots for the high-mass, isolated EdGB binary. The same bias we observed in the low-mass system persists here, and very high SNRs (e.g., those provided by DECIGO) are required to suppress the $q \rightarrow 1$ tendency. See Figures 7.3, 7.5 and 7.6 for comparison and more detail on the $\zeta_{\text{EdGB}}-q$ bias. The green marks are duplicates of the GR MPA plots in Figure 7.9, to see how the EdGB extension affects the uncertainties in the other parameters.

$\sigma_{\zeta_{\text{EdGB}}}$ is roughly the same between the high and low mass binaries. The difference is due to the fact that $\sqrt{\alpha_{\text{EdGB}}}$ scales with the mass of the system, as shown in equation (5.6), and this is why low mass binaries are generally far more useful for constraining quadratic gravity theories than higher mass counterparts.

7.3.2.2 Dynamical Chern-Simons

The high-mass dCS MPA plots are shown in Figure 7.11. Again, we can see in Figure 7.11 that including ζ_{dCS} as a free parameter has little effect on the location parameters but significant degeneracies with the masses and spins are clear, as one would expect by the dependencies of β_{dCS} . Of the detector networks studied here, the ζ_{dCS} constraints are not strong enough to place the vast majority of the posterior samples inside the small

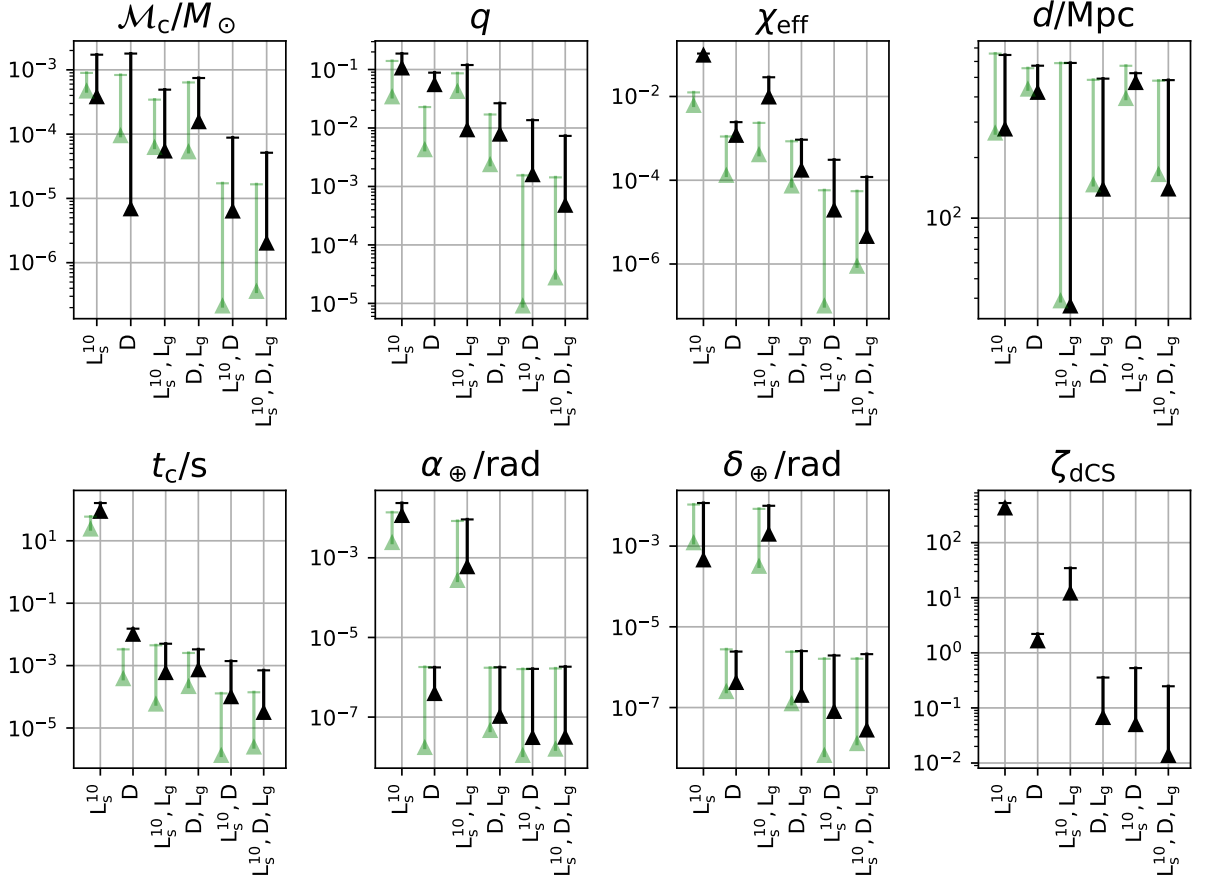


Figure 7.11: The MPA plots for the high-mass, isolated dCS binary. LISA clearly plays a far more important role for the high-mass system, and its ability to constrain the masses and spin parameters (when the signal can be detected) are more greatly affected by extending the model to dCS gravity than for the low-mass system, due to the masses and spin dependence of the dCS parameter.

coupling approximation limit, $\zeta_{\text{dcs}} < 1$, with $\sigma_{\zeta_{\text{dcs}}} \approx 0.6$.

However, supposing these posteriors were valid, this standard deviation would translate to a constraint on the dimensionful coupling parameter of $\sqrt{\alpha_{\text{dcs}}} \lesssim 371.7$ km at 90% confidence. This is around a factor of 8 worse than the bound provided by the low-mass system of 47.4 km. In order to improve on the currently existing best constraints from GW observations then, one clearly requires a vast number of favourable (low mass, highly asymmetric mass, high spin, high SNR) BHB events. The outlook for tightly constraining dCS gravity with stellar-order binary GWs therefore does not appear to be good, even for best case future scenarios in which DECIGO plays a part. This is unfortunate, since strong constraints using other methods prove difficult to obtain and using GWs has been called a ‘last resort’ [149].

7.4 Prospects for multi-band GW data analysis

7.4.1 Population models, detection rates

Evaluating an estimated detection rate involves two steps. First, assume a model of the universe that predicts a distribution of BHBs per (comoving) volume. Then, take a detector network configuration and detection statistic (i.e., a threshold SNR) and with it, integrate over the distribution those regions that constitute a positive detection. The equation describing this estimate is derived in Ref. [150] and estimations for the LISA+LIGO network computed: a 4 year (10 year) LISA mission duration should result in roughly 5–30 (35–170) multi-band detections given an SNR detection threshold of 4. Recall that in Section 7.2.1 we found that the chirp mass begun to ‘freeze out’ in LISA already at SNR $\rho_* \approx 2$. There was a significant impact on other parameters, which, we surmised, arose from strong LIGO constraints on the orientation, delaying the full freeze out until $\rho_* \approx 2.5$ (Figure 7.2).

However, these posteriors were produced using a reduced parameter model, omitting sub-leading order spin and eccentricity parameters (along with additional time-delay parameters as we shall investigate in the next chapter, which can often be necessary to model and requiring multiple parameters). Introducing this extra freedom is tantamount to requiring a higher SNR threshold, depending of course on how strongly this freedom can affect waveform templates (see Section 6.2.2). We saw this sort of effect occurring in small part as we extended our model, albeit in a different direction, when we introduced extra freedom by including the quadratic gravity parameters. By adding the dCS parameter, the signal froze out at $\rho_* \approx 3$ (see Figure 7.7). By adding the EdGB parameter, the signal appeared to freeze out (though with a strong bias) at $\rho_* \approx 4$ (Figure 7.5).

It might therefore seem likely that we would have to rely solely on very heavy BHBs for multi-band GW data analysis with a LISA+LIGO like detector network; recall that in Table 7.2 the low-mass system observed even for 10 years in LISA gives an SNR of ~ 5.7 . However, we note the recent work in Ref. [151] in which fully modelled (17 parameter, including eccentricity) BHB waveforms and the LISA detector model are given a very thorough treatment and analysis of signals present in LISA (only) data reported. The authors found that signals were generally recovered very well, even finding good poste-

riors for signals with SNRs as low as 5.7. This successful recovery of signals in LISA alone bodes well for multi-band detection of BHB signals. Moreover, the strongly constrained posteriors from future ground-based, and space-based deci-hertz detectors used as priors can bring down the SNR detection threshold for other detectors even more, bolstering the possibility of digging into the LISA noise for low-frequency information at low SNRs (see Section 8.1.2).

7.4.2 Detector network summary

In this section there will be continual reference to the marginal PAR plots, which we point to once here, in Figures 7.1, 7.3, 7.8, 7.9, 7.10 & 7.11, to avoid repetition.

We also talk about source localisation; we approximate the 90% confidence *sky localisation volume*, V , by supposing that the 90% confidence iso-contour of sky position (the distance and sky coordinates) is roughly given by an ellipsoid, with its semi-axes given by the 90% (approximately 1.65 standard deviations) bounds on each of the parameters. Then we have

$$V \approx \frac{4\pi}{3} 1.65 \sigma_d d^2 \sin(1.65 \sigma_{\alpha_{\oplus}}) \sin(1.65 \sigma_{\delta_{\oplus}}) \cos(\delta_{\oplus}), \quad (7.5)$$

where we have written the right ascension (ra) as α_{\oplus} and the declination (dec) as δ_{\oplus} .

7.4.2.1 LISA

The low-mass binary signal searched for using a 4 year LISA signal (with $\text{SNR} \sim 3.1$) did not yield a detection. The 10 year observation of the high-mass system ($\text{SNR} \sim 13.8$) provided excellent constraints with the GR analysis, which, for all parameters aside from the sky position, closely matched those provided by DECIGO ($\text{SNR} \sim 6600$).

The sky localisation volume for the high-mass binary in LISA $V \approx 1.28 \times 10^4 \text{ Mpc}^3$, corresponding to hundreds of potential host galaxies [152]. This binary is located at a distance of 3000 Mpc; for the same system at 1000 Mpc, the number of potential host galaxies from LISA alone drops down into the tens. If the host galaxy can be confidently located, one can acquire a redshift measurement of the galaxy, at once breaking the mass-redshift degeneracy and obtaining a constraint on the Hubble constant.

When searching for EdGB modifications, the constraints on all parameters are mostly unchanged beside the mass ratio, which acquires a strong bias. The constraints on ζ_{EdGB}

itself are poor since the posterior rails up against the prior, due to the presence of the platform of probability in its marginal as first discussed in detail in Section 7.2.2.1. The dCS parameter is ‘nice’ in that it is not strongly correlated with other parameters, however, its impact on the gravitational waveform is not great enough to improve the current best constraints.

7.4.2.2 DECIGO

DECIGO is by far the most sensitive detector of those considered here. The SNRs are generally extremely high, leading to tightly constrained posteriors which closely approximate Gaussian distributions. The sky localisation is excellent; with equation (7.5) we have for the low-mass system that $V \approx 10^{-5} \text{ Mpc}^3$, corresponding to 10^{-7} potential host galaxies, meaning that it is very highly likely to be able to locate the host galaxy. For the high-mass system (at 3000 Mpc) we have that $V \approx 5 \times 10^{-4}$, again giving excellent chance of host galaxy location (overlooking the technical details regarding making such a host galaxy redshift observation). Thus DECIGO on its own, in principle, appears to be sensitive enough to many GW events occurring across a large fraction of the history of the universe that it will be able to place useful constraints on the expansion history of the universe (i.e., a time dependent Hubble parameter). The EdGB parameter is well-constrained with DECIGO and bias is not noticeable when this parameter is included. The dCS parameter, on the other hand, cannot be constrained well enough to improve current best constraints, nor even to comply with the small coupling approximation.

7.4.2.3 LIGO

The vast majority of GW data analysis research to date has been focused on LIGO and its data, and its strengths and limitations are already well understood. It does not feature heavily here outside of considering it as part of a detector network and the posteriors in the above plots are mainly included for comparison between the different networks.

7.4.2.4 LISA & LIGO

We have found, albeit with significant model simplifications, that LISA is able to enhance the posterior given by LIGO for signals of very low LISA SNR ~ 3 (see Figures 7.2 & 7.7 for GR & dCS models, respectively). Even with the problematic bias that arises due to the

structure of the EdGB modification in equation (5.11), the EdGB gravity model search (with a reasonable prior) clearly finds a signal when LISA SNR ~ 4 . This is despite the fact that the EdGB parameter enters the phase at -1PN and therefore corresponds to an entirely new term in the phase; it is a good example of a newly introduced parameter that drastically increases the number of templates in the template bank which is usually expected to significantly increase the SNR detection threshold (see Section 6.2.2 for details). We saw in Figure 7.5 that in order to overcome the bias in the LISA posterior introduced by the EdGB parameter, the SNR must be increased to around $\rho_* \sim 500$ (see Figure 7.5). These SNRs at this sort of mass range are almost entirely certain not to occur, however, it is likely that with strong EdGB constraints from third generation ground-based detectors, the majority of the platform of probability we saw in the EdGB marginal, along with the associated bias, will be strongly suppressed.

7.4.2.5 LISA & DECIGO

For low-mass binaries, there is no noticeable improvement in the posteriors by adding LISA to DECIGO, aside from, interestingly, on the introduction of the quadratic gravity parameters. In the case of introducing new parameters, the variances generally increase, but LISA somewhat inhibits that increase with the information it provides on the masses and spins. For high-mass systems spanning LISA & DECIGO, the situation is different as constraints are significantly enhanced on masses and spins by adding LISA to DECIGO. The greatest benefit of this detector pair is in constraining quadratic gravity parameters themselves, with the EdGB parameter being constrained to better than 3 orders of magnitude (the EdGB-mass ratio bias is strongly present in the high-mass LISA signal, which is essentially removed by DECIGO) compared to DECIGO alone. The dCS parameter is still too poorly constrained for the small coupling approximation to be valid.

7.4.2.6 DECIGO & LIGO

If the DECIGO mission is confirmed, there will be little more help that LIGO (the particular model we have used) can offer, due mainly to the vast discrepancy in SNRs between the detectors. There is, however, a significant improvement in the standard deviation of the distance for the low-mass system by a factor ~ 3 , which can be particularly useful for localisation and host galaxy identification. Recall that our LIGO model does not include

quadratic gravity modifications, and that the modifications are greatest at merger; this is somewhat of a shortcoming of the present study. The situation will be vastly improved with third-generation ground-based detectors, too, such as the Einstein Telescope, which will have a much higher sensitivity than our LIGO model.

7.4.2.7 LISA, DECIGO & LIGO

In general, the results of the three-detector network are not surprisingly very similar to those of the LISA+DECIGO network, due to the relatively tiny SNR in LIGO with similar frequency band to DECIGO. The exception is a small but useful contribution to the constraints on the distance (for the low-mass system) one acquires from including LIGO alongside DECIGO.

7.4.3 Outlook

Clearly, DECIGO dominates in terms of providing parameter constraints for the waveform model we have considered in this chapter. It is also crucial for avoiding the bias observed when including the EdGB parameter in the model. Since the status of DECIGO is as of yet unconfirmed however, it is pragmatic to be more concerned with the LISA+LIGO network. For this network, using the reduced parameter GR model we have employed for binaries (*known* to be isolated) can dramatically improve the parameter constraints that one acquires from LIGO only, where we can see the beginnings of precision GW astronomy coming to fruition. But we must be cautious in including the very low SNR LISA signals along with the LIGO analysis; as we saw in Figure 7.1, a significant bias in multiple parameters is observed before the primary mode originating from the LISA data freezes out. Further investigation into this behaviour is required to confirm the cause (we surmised that the posterior on the orientation from LIGO suppressed the mode containing the true parameters too much, until the LISA SNR was large enough to overcome the LIGO constraints) to evaluate its impact in the fully modelled case, and to work out precisely when it is safe to trust the results of the multi-band data analysis.

We have also found that none of the detector networks used can provide usable constraints on the dCS parameter. It is conceivable that detectors such as the Einstein Telescope will observe particular ‘golden sources’ at relatively high SNR to provide better constraints for dCS gravity, and this seems to be the best hope for the foreseeable future

for dCS constraints using GWs. However, given the detector networks under consideration in this thesis, we will exclude the dCS model for the remainder of our investigations.

Multi-band GW Astronomy Part II: Binaries in Galactic Orbits

The next extensions of the GW model to be considered are the varying time-delays that modify the waveform due to the motion of the source and its background geometry. These were introduced in Chapter 4, along with *approximants* of these time-delays that require either one or two parameters to describe, rather than the full set of six Keplerian elements required in general. The ‘high-mass binary’ is used for the majority of the investigations in this chapter; we are often strictly limited to using higher mass systems, since the SNRs of the lower mass binaries can be too low to make a detection in LISA, even in the case of effectively having a strongly constrained prior (and thus lower SNR detection threshold) provided by DECIGO. However, it is reasonable to expect these sort of signals to exist in the LISA data, as discussed in Chapter 4.

In this chapter, the expectations about the effects of low-frequency waveform modifiers on multi-band GW data analysis introduced in Chapter 4, and the concepts and simplifications introduced there, are verified by experiment and shown to be useful aids for analysis. Over the course of the chapter, we will see that the structure and properties of posteriors for signals with small time-delays is rather simple. A constant or linear acceleration, for example, entering the phase at large negative PN order(s), is not easily interchangeable with (combinations of) the non-negative PN terms in the GR isolated binary waveform, and thus there is very little correlation between the parameters. However, for large (and highly non-linear) time-delays, the variation of the posterior structure

as a function of the injection parameters is highly intricate: one can already expect that predicting correlations of Keplerian and intrinsic binary parameters by, for example, trying to match effects of the time-delays (as exemplified in Figure 4.4) with PN terms in the phase of the binary’s waveform is a highly non-trivial task. We verify this through examples, showing that for those systems with significant time-delays, it is very difficult to predict how a specific system will behave, and that, in general, one must resort to running simulations for each system to know the details of the behaviour.

This chapter examines properties of posteriors of the same fiducial systems that were studied in Chapter 7 (see Table 7.1) and is structured as follows. We look at the low-mass binary in Section 8.1, and confirm the estimated detectability of (constant) acceleration, then examine the MPA plots for models of constant and linear accelerations. The same analysis for the high-mass binary is given in Section 8.2. We analyse the posteriors of the fully modelled Kepler orbiting high-mass binary in Section 8.3 generated by the LISA+LIGO network. In Section 8.4, we investigate retrieving the Keplerian elements only, using a ‘known carrier signal’. The time-delay approximants and their regions of validity as explored in Section 4.2.4 are tested in Section 8.5. We show the results of simultaneously allowing a free (constant) acceleration parameter and EdGB gravity parameter in Section 8.6.

8.1 ‘Low-mass’ binary with acceleration

In Chapter 7, we analysed ‘isolated binaries’ using a waveform model for both the injected signal and the waveform search templates, where the binary was assumed to be in a flat background geometry, i.e., with no non-linear time-delay modifications arising from the relative position of the binary. Although this may sometimes be an accurate approximation, we cannot know whether it is appropriate to use that model before having analysed the data. It is instructive, therefore, to ascertain the magnitudes of these non-linear components of the time-delays at which it becomes necessary to include the relevant parameters in the model. One might say this occurs when the posterior becomes ‘significantly disrupted’, or when the parameter becomes ‘measurable’. In Section 4.2.4, we declared this point to be when one of the mean values of the parameter marginals is offset by 1 standard deviation.

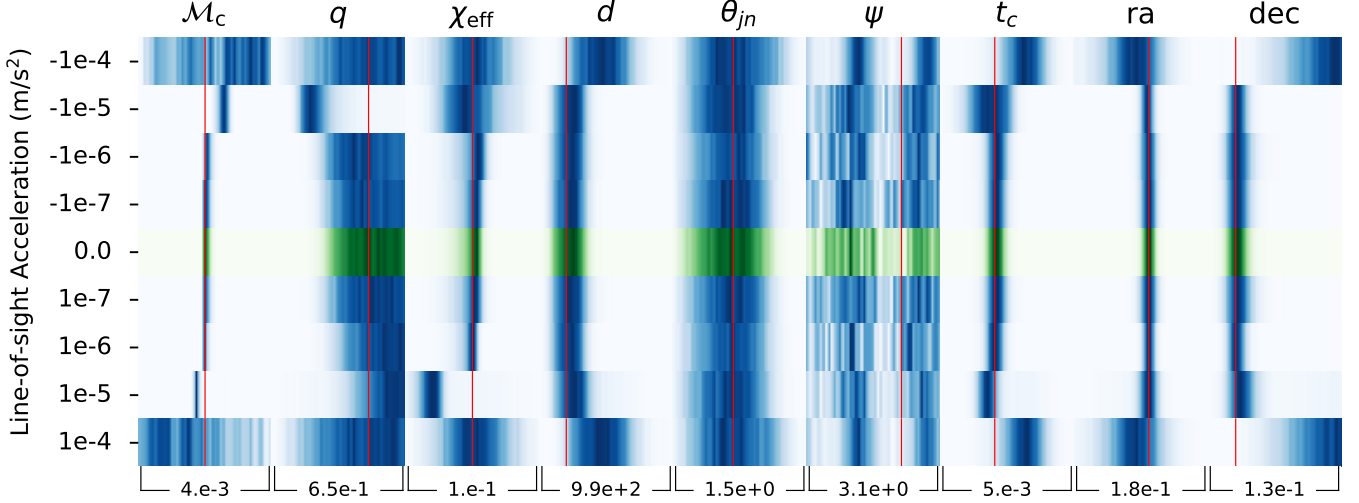


Figure 8.1: Marginals of the posterior of the low-mass system in LISA(4 year)+LIGO as a function of (constant) line-of-sight acceleration of the binary. The signal is generated by an accelerating binary, but the model used for defining the posterior does not include acceleration. We can thus observe the acceleration at which the posterior begins to break down, or, indeed, the acceleration at which (line-of-sight) acceleration must be included in the model. We observe perturbations of the posterior at acceleration magnitudes of around 10^{-6} m s^{-2} , and a strong breakdown with acceleration magnitudes of 10^{-5} m s^{-2} .

The simple, constant Newtonian line-of-sight acceleration, a , corresponds to a quadratic time-delay, which we expect to be measurable at the a that satisfies equation (4.42). For our low-mass binary in LISA+LIGO, with the signal details given in Table 7.2, the estimated detectable acceleration is

$$|a| \approx 6.8 \times 10^{-7} \text{ m s}^{-2}. \quad (8.1)$$

The posterior evolution plot showing the parameter marginals as functions of the unmodelled constant acceleration is shown in Figure 8.1. We can see that the marginals become disrupted slightly at $|a| = 10^{-6} \text{ m s}^{-2}$ and significantly at $|a| = 10^{-5} \text{ m s}^{-2}$.

This is around a factor of 10 higher than our estimate in equation (8.1). However, recall from Section 4.2.4 that (as discussed in Ref. [26]) measurable offsets in the chirp mass are around a factor of 10 greater than one might at first conclude; when spin is included as a free parameter, the degeneracy with chirp mass and spin widens the marginal by roughly a factor of 10. We omitted this extra factor of 10 from our equation (4.42) to err on the side of caution in our parameter space division scheme. With this factor included, we have a modified estimated detectable acceleration of

$$|a_{\mathcal{M}_c; \chi_{\text{eff}}}| \approx 6.8 \times 10^{-6} \text{ m s}^{-2}. \quad (8.2)$$

There is indeed excellent agreement between this value and the results in Figure 8.1. This bolsters our assertion that an IAD of $I \approx 3 \text{ m s}^{-1}$ makes a reasonable (if conservative) bound for the different regions of the Keplerian element parameter space, as given, for example, in equation (4.46).

8.1.1 Constant acceleration

We now produce and examine the MPA plots for our low-mass binary with constant acceleration as a free parameter. These are shown in Figure 8.2. The LIGO only and LISA only posteriors are omitted since clearly acceleration is not detectable in LIGO alone, and the LISA SNR is too low for a signal to be found without the help of other detectors. In comparison to the posteriors with fixed/unmodelled acceleration (green lines in Figure 8.2), notice that the parameter standard deviations for those networks including LISA hardly evolve; there is little freedom provided by other parameters to account for acceleration (recall constant acceleration enters the waveform at -4PN , see Section 5.6). This is with the exception of the coalescence time, which displays the same peculiar behaviour we observed in Section 7.2.2.1 by the inclusion of the EdGB parameter; the coalescence time constraints *improve* by including constant acceleration as a free parameter. As before, the reasons for this are presently not clear, however, the effect may be due to fine details of the structure of the posterior; see Section 7.2.2.1 for discussion and an example of a potential cause for this behaviour.

Those networks containing LISA strongly constrain the acceleration, as expected (when the signal can be located). Observe that LISA+LIGO constrains the acceleration at 1-sigma to the same degree as expected in equation (8.2). The very high SNR of DECIGO overrides the rudimentary expectations leading to equation (4.42) however. With the signal properties in DECIGO from Table 7.2, we would expect that accelerations of $|a| \gtrsim 1.8 \times 10^{-3} \text{ m s}^{-2}$ are required in order to have a measurable impact. Clearly this is not the case, and this is because our estimate of the minimum expected acceleration does not take into consideration the SNR, which strongly informs the posterior; one could indeed measure arbitrarily small accelerations if one were allowed to have arbitrarily high SNRs.

There is a strong positive correlation of the acceleration with the chirp mass and effective spin, and negative correlation with the mass ratio in DECIGO. In LISA however,

there is no such clear correlation. This is reasonable considering the durations of the signals; thinking of the acceleration as the -4PN term in the phase, the modification to the waveform takes time to accrue. Over long time periods, it is difficult to mistake a -4PN term with another PN term (recall that, perturbatively at least, a $p\text{-PN}$ term corresponds to a $\tau^{(5-2p)/8}$ term in the phase), but over relatively short intervals it becomes easier for (combinations of) other PN terms to mimic the -4PN term. The conjunction of LISA with the DECIGO posterior accordingly subdues correlations between these parameters.

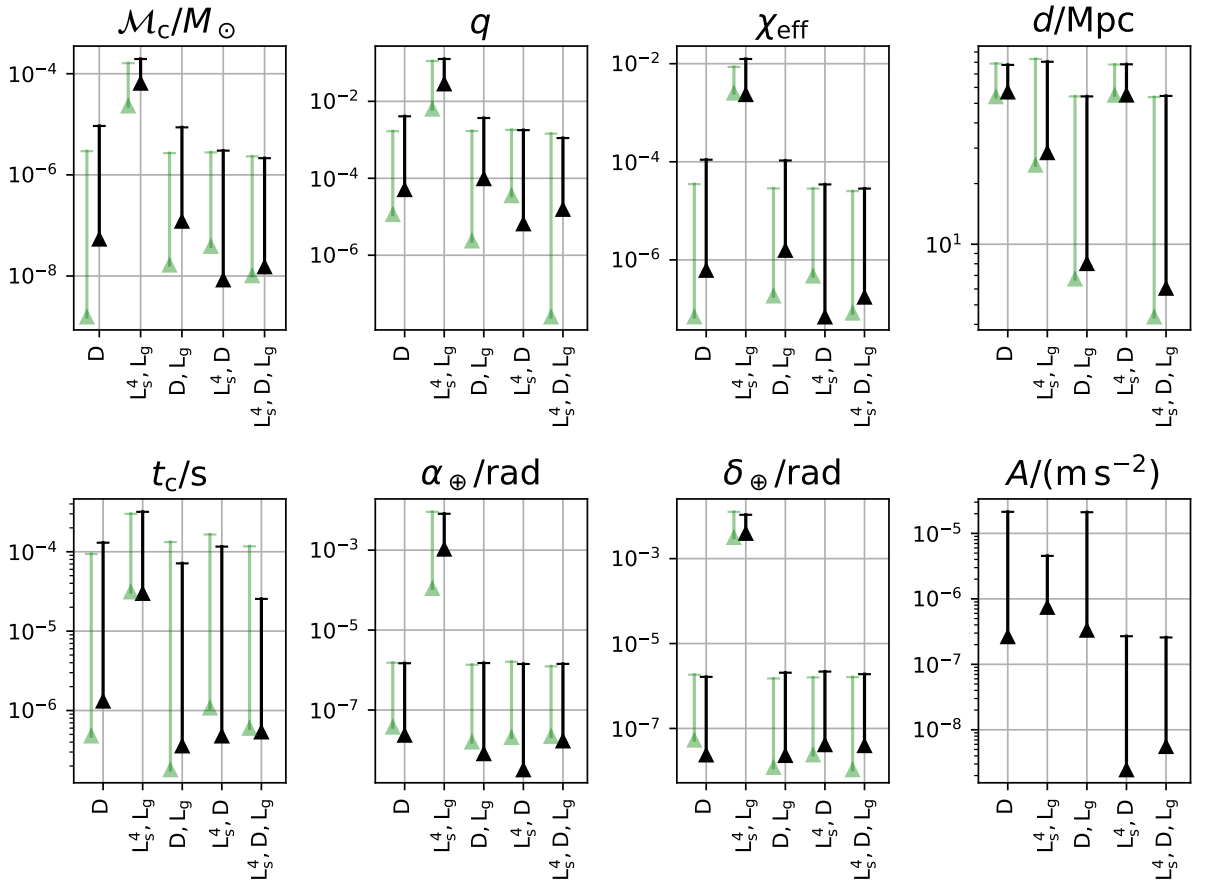


Figure 8.2: The MPA plots for the low-mass, constantly accelerated GR binary. LISA alone does not feature as the signal is not detected, and LIGO is omitted since it does not detect acceleration. DECIGO is able to constrain acceleration rather well despite the relatively short signal duration due to the high SNR. When LISA detects the signal, the acceleration constraints are strong, as to be expected, since the acceleration over long duration signals builds up significant phase offset. The green marks are a duplicate of the non-accelerating signal models, as in Figure 7.1. Interestingly, we see that on expanding the model to include acceleration, coalescence time is *improved*. We saw the same effect on coalescence time by extending the model to EdGB gravity in Chapter 7, and surmise that this could be due to similar intricacies of the distribution as described in Section 7.2.2.1.

8.1.2 Linear acceleration

In a gradual extension of the model, we now include a second acceleration parameter to model acceleration as a linear function of time. The MPA plots are given in Figure 8.4. Immediately we notice that the addition of the 4 year LISA signal to LIGO fails to locate the signal in LISA; the increase in the number of templates caused by introducing another parameter and a new term in the phase at -8PN (see Sections 5.6 & 6.2.2) increases the LISA SNR threshold for this multi-band analysis beyond the SNR of the LISA signal. For this reason, we have increased the duration of the signal to 10 years in LISA for the LISA+LIGO analysis, in order to have a posterior for comparison defined by a network that does not include DECIGO. However, one does not need to increase the duration of the LISA signal when DECIGO forms part of the network; the strong parameter constraints of DECIGO make it easier to dig into the LISA noise.

The DECIGO+LIGO posterior does not feature because the posteriors demonstrate




Figure 8.3: The corner plot of the posterior for the low-mass system given the linear acceleration model in DECIGO+LIGO, marginalised onto the parameters, (A, B) . The two acceleration parameters are strongly degenerate (within the prior ranges) due to the short signal length which does not allow for much phase offset to accrue, and thus does not allow much information on acceleration to be captured. The signals are present in DECIGO for up to a few days, so the rate of change of acceleration must be very high in order for acceleration to be measurably different between the start and the end of the DECIGO signal; DECIGO is realistically only generally capable of determining one acceleration parameter.

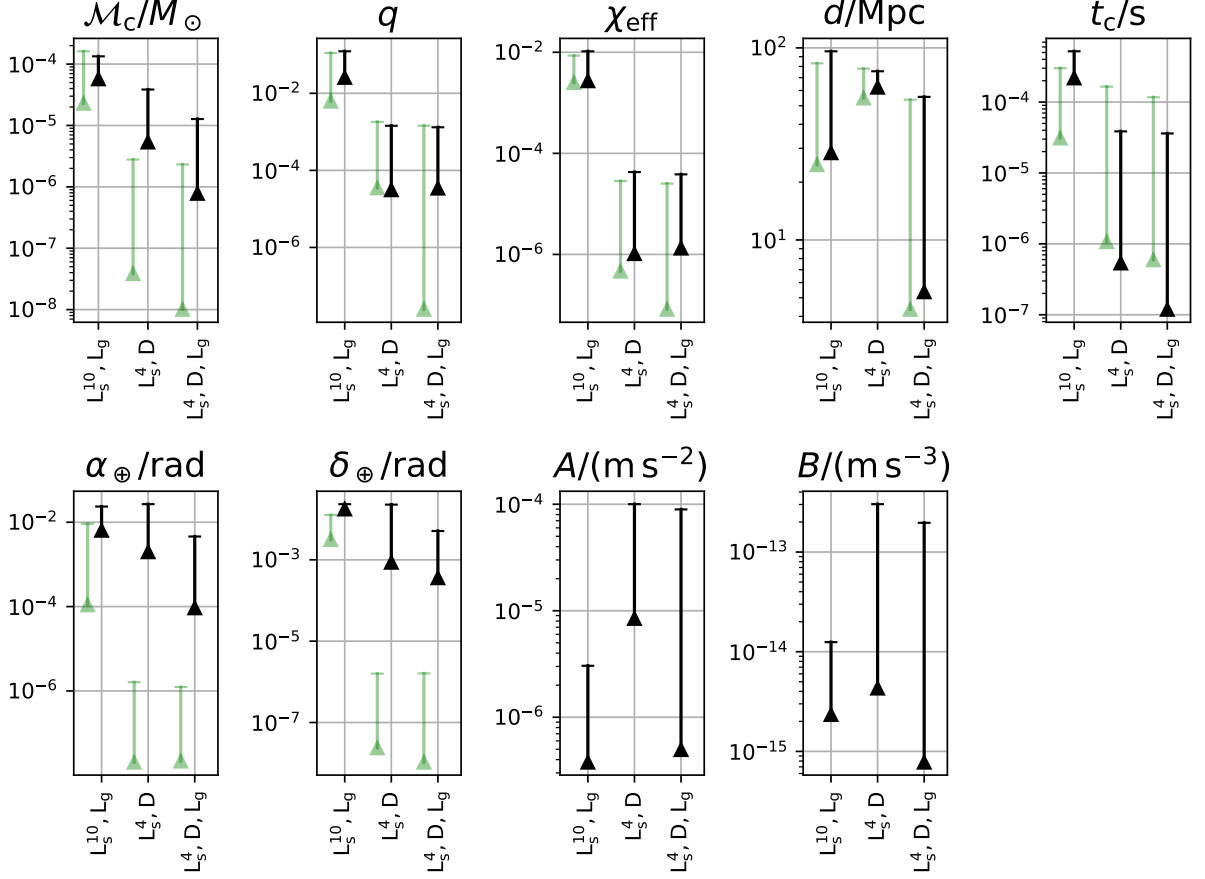


Figure 8.4: Marginal precision and accuracy plots for the low-mass, linearly accelerated GR binary. Extending the model to include a second acceleration parameter raises the LISA(4 year) SNR detection threshold such that the signal cannot be located in LISA(4 year)+LIGO. We instead display results from the LISA(10 year)+LIGO posterior, which is able to constrain acceleration extremely well, as expected when the signal is located with LISA. We also omit DECIGO since only one acceleration parameter is measurable in DECIGO alone (see Figure 8.3). The introduction of the second acceleration parameter apparently severely impacts the sky position, however, this is not realistic since the motion of the LISA and DECIGO spacecraft has not been fully modelled, and these motions are expected to help to constrain the sky position; the timescale of the induced modulations of the waveform from this motion will be not at all, or very weakly, degenerate with acceleration of the source. Again, coalescence time is *improved* by including A and B , see caption of Figure 8.2 for discussion.

the fact that only one acceleration parameter can be picked out; this is clear in Figure 8.3 showing that parameters A & B are strongly degenerate over this prior range. Note the prior ranges on A & B are large as compared to the constraints shown in Figure 8.4 provided by LISA, and already with this range the posterior takes a considerably long time to converge, likely due to this extremely thin mode. It is likely that only a constant acceleration is detectable in DECIGO for the majority of sources despite the very large SNR, due to the signal’s short duration of 11.8 days. For non-linear acceleration to be detectable, it must be of a very large (relative) amplitude to change the acceleration considerably in a few days. For many sources, it is reasonable to say that DECIGO is able to measure the approximately ‘instantaneous’ acceleration of the source in the DECIGO band. In Figure 8.4 we can see that the chirp mass constraints suffer slightly from the inclusion of the acceleration parameters, and the sky position constraints from DECIGO are almost completely lost. However, the detector modelling only includes a Roemer delay which, for a short signal duration, can be easily understood as being absorbed into the source acceleration parameters. A proper modelling of the ‘tumbling’ motion of the DECIGO spacecraft, which may happen on the order of a day, should help to break this degeneracy, restoring good sky localisation.

8.2 ‘High-mass’ binary with acceleration

We will now use the high-mass binary to verify the estimate of the expected detectable (constant) acceleration. With the signal properties from Table 7.5 and equation (4.42), we would prudently include acceleration in the model of the high-mass system with 10 years in LISA for $|\mathcal{a}| \gtrsim 5.2 \times 10^{-7} \text{ m s}^{-2}$. Again, we expect to detect accelerations a factor of 10 greater than this, at $|\mathcal{a}_{\mathcal{M}_c: \chi_{\text{eff}}}| \gtrsim 5.2 \times 10^{-6} \text{ m s}^{-2}$, due to mass-spin degeneracy [26]. This expectation again is decidedly consistent with the results we obtained, shown in Figure 8.5.

8.2.1 Constant acceleration

The same analysis of extending the model to include a constant acceleration parameter is repeated here, for the high-mass system. We display the MPA plots in Figure 8.6. LISA on its own is able to easily locate the signal and constrain the acceleration, and performs

around two times better than the low-mass system in relation to multi-detector networks (i.e., adding another detector for the high-mass system helps to constrain the acceleration half as much as adding another detector to the low-mass system), which is simply accounted for due to the high-mass signal being observed for around 2.5 times longer than the low-mass system.

DECIGO on its own constrains the low-mass acceleration around 500 times better than the high-mass system. But notice we also have that the expected detectable acceleration for the high-mass system in DECIGO is $|a_{\mathcal{M}_c: \chi_{\text{eff}}}| \gtrsim 0.95 \text{ m s}^{-2}$. The ratio of this to the expected detectable acceleration of the low-mass system (in DECIGO alone) as computed above is ~ 530 . Thus, although equation (4.42) breaks down for very high SNRs where the SNR dominates control of the variance, the *relative* expected detectabilities of the accelerations of different binary systems (at some given SNR) still appears to remain accurate, as one might have anticipated. As we saw for the low-mass system in Section 8.1.1, for the high-mass system in DECIGO there is a fairly strong positive correlation of the acceleration with the chirp mass and effective spin, and negative correlation with the mass ratio. There are no such clear correlations in LISA.

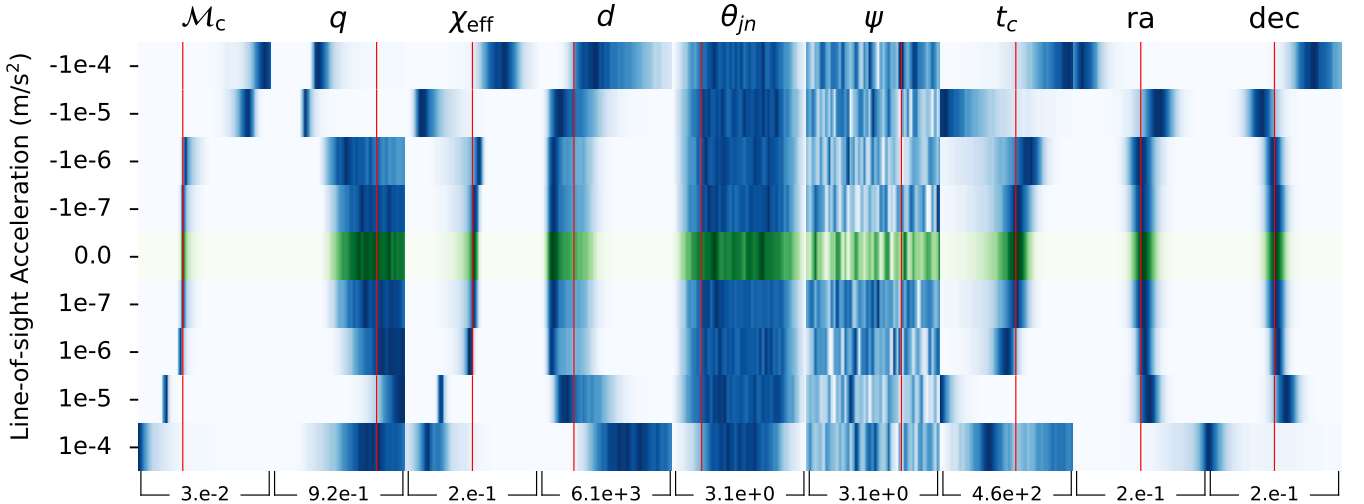


Figure 8.5: Marginals of the posterior of the high-mass system in LISA(10 year)+LIGO as a function of (constant) line-of-sight acceleration of the binary. The signal is generated by an accelerating binary, but the model used for defining the posterior does not include acceleration (see also Figure 8.1). We observe perturbations of the posterior at acceleration magnitudes of around 10^{-6} m s^{-2} , and a strong breakdown with acceleration magnitudes of 10^{-5} m s^{-2} .

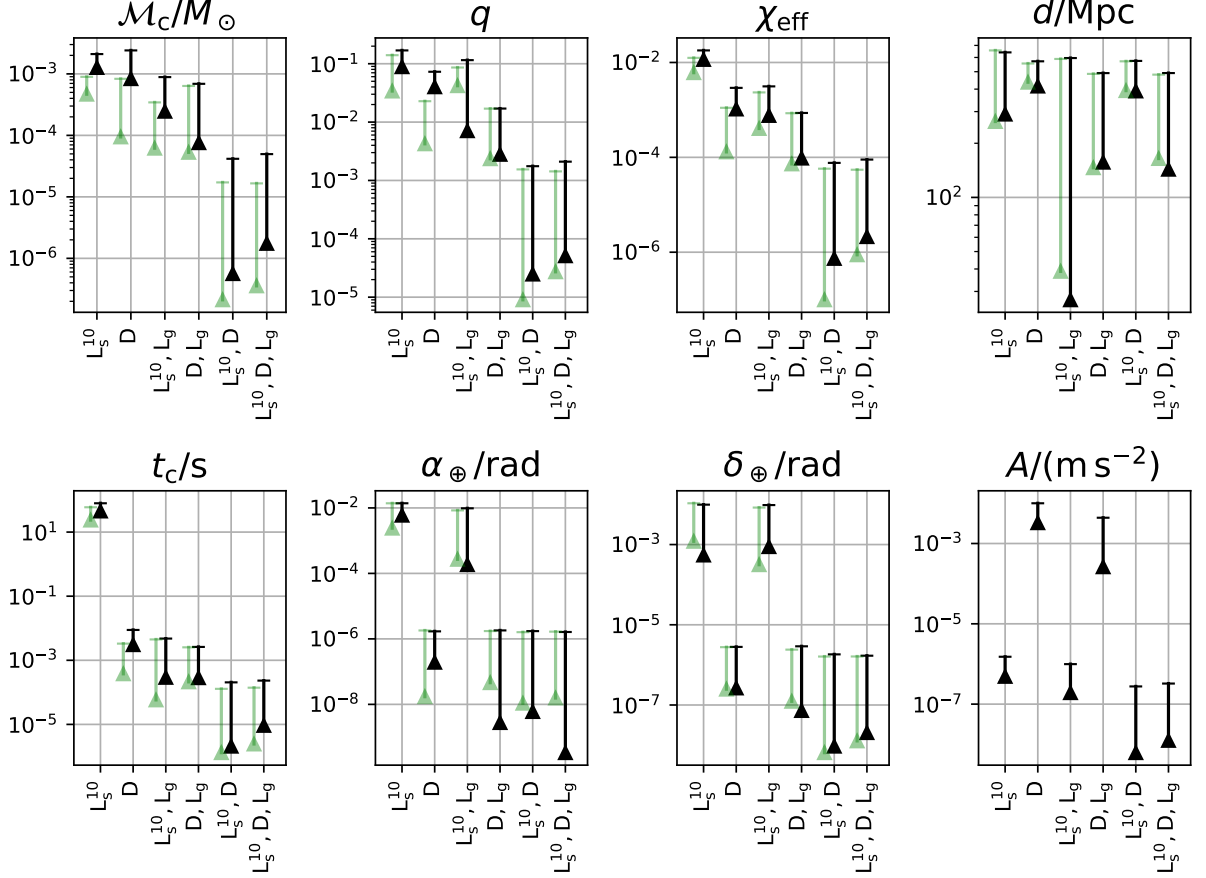


Figure 8.6: Marginal precision and accuracy plots for the high-mass, constantly accelerated GR binary. LIGO is omitted since acceleration is not measurable there. Including acceleration parameter A increases the uncertainty on masses by around a factor of a few, but with very little effect on other parameter constraints.

8.2.2 Linear acceleration

Finally, for the high-mass systems with linear acceleration, the MPA plots are shown in Figure 8.7. DECIGO alone again is clearly only sensitive to one acceleration parameter on examination of the posteriors. Marginalised onto the A & B parameters, the posteriors show precisely the same degeneracy as that demonstrated in Figure 8.3, thus DECIGO is omitted from Figure 8.7. Indeed, independent measurement of the two acceleration parameters is absolutely dependent on the presence of LISA in the network. Notice that, in all networks, there is very little impact on the parameter constraints caused by including the two acceleration parameters (besides the chirp mass, which is anyway very strongly constrained and degenerate with a constant line-of-sight velocity; these extremely precise details of the chirp mass are not so interesting). This is a fortunate state of affairs

for the purposes of other scientific pursuits such as localisation (for Hubble parameter measurement) and population studies.

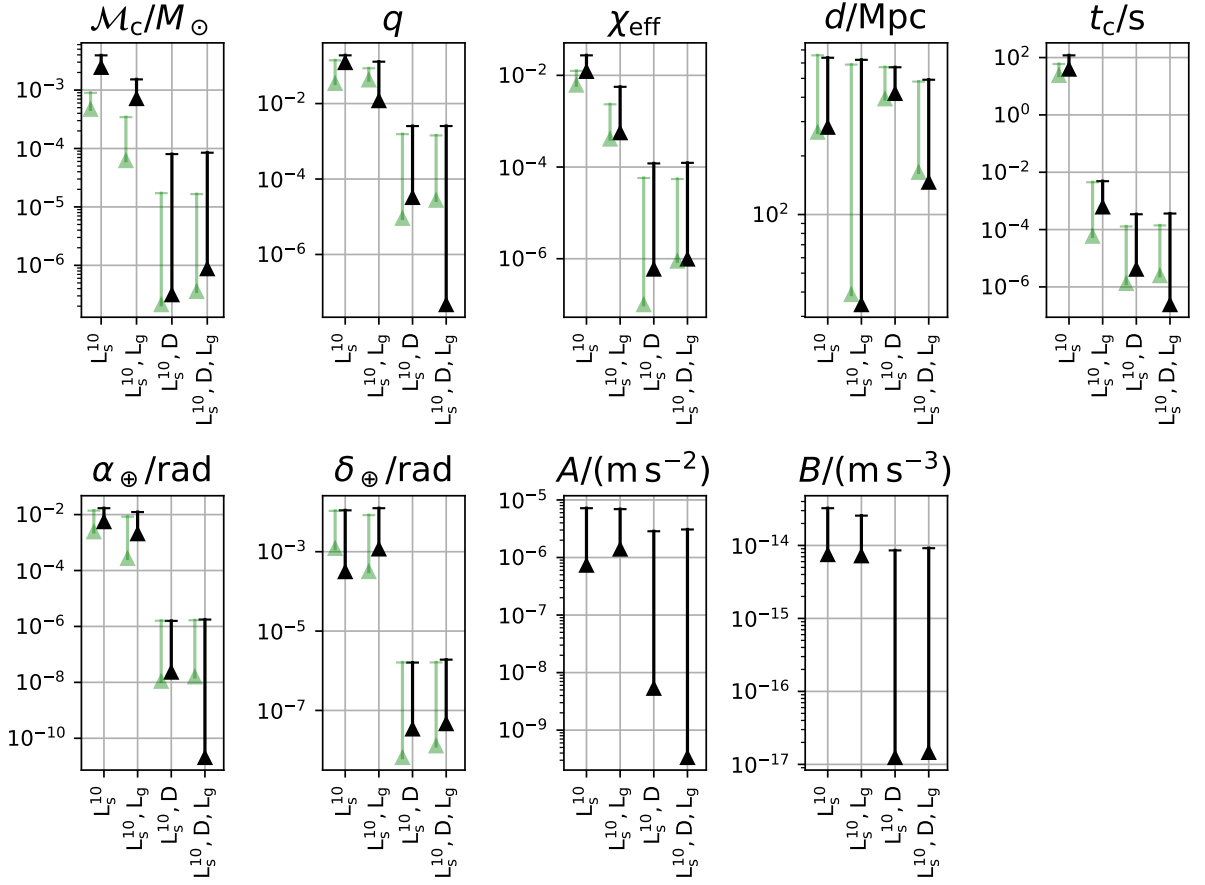


Figure 8.7: Marginal precision and accuracy plots for the high-mass, linearly accelerated GR binary. For the high-mass system, the acceleration parameters are informed by LISA with high enough SNR to prevent the posteriors from losing too much precision when introducing new parameters (compare to the low-mass example, Figure 8.4).

8.3 Complete time-delay model

Now we finally use the fully modelled waveform (as per Chapter 4 and the caveats pointed out in Chapter 6) of a binary in orbit around a supermassive, Schwarzschild black-hole. The set of results here is limited because of the lack of appropriate priors and the time required to evaluate the posteriors. With priors set too wide, one faces a great sampling difficulty as the sampler appears to struggle to locate the primary mode, and convergence

time is vastly too long¹. However, one must wait for the posterior to see if the priors had been set too narrowly. Using the Fisher information matrix to approximate the range was unfortunately not reliable (see Section 6.2.1).

We note that the PE convergence times for the non-accelerating models are fast in general due to the use of the downsampling procedure explicated in Chapter 3. On introducing the extra Keplerian elements, the PE convergence time slowed considerably. The reason for this unfortunately remains elusive; the likelihood evaluation times were only moderately affected by the extra parameters, and as we will see, the posteriors themselves are not particularly irregular. The state and progress of the sampler seemed normal throughout the runs and inspection of the likelihood values of the returned posterior samples appeared to describe a smooth function as one would expect. The dimensionality of the parameter space is certainly a factor, but the slow convergence persisted even on fixing some or all of the parameters of the binary and reducing PE to the six Kepler parameters only. We would of course like to fully understand this, but at present we cannot report further on this here.

With long PE convergence times, the range of results suffered. However, we arrived at three complete, well-formed posteriors for this section, shown in Figure 8.8 and Figure 8.9, which display posteriors for two systems we refer to as the ‘nearby’ system and the ‘distant’ system, respectively. The final posterior of the nearby system took 16.3 days to converge. This is considered slow, especially since we are using the downsampling procedure of Chapter 3. Previous attempts of the nearby system posterior took up to 36 days to complete. The final posterior of the LISA+LIGO distant system took 5.5 days to converge, with previous attempts taking up to 22 days, and this system with the inclusion of DECIGO took ~ 10 days. The convergence times could be reduced in a limited way by ‘parallelising’ the likelihood evaluations (spreading the computational work over multiple processors). It is interesting to observe that, very roughly speaking, increasing the number of time-delay parameters to n from zero ($n > 0$) increases the number of `Nessai`’s likelihood evaluations and rejected samples by a factor of $\sim n + 1$ (with any model; constant or linear acceleration, or Kepler time-delays).

¹Different samplers were tested for speed. Although `dynesty` converged relatively quickly, results were poor in that multiple runs on the same data returned inconsistent posteriors and did not locate the injection parameters. We continued using `Nessai`, which, despite being significantly slower, appears to be highly consistent. This strongly suggests that the posteriors obtained for the results presented here are accurate.

System	α (rad)	β (rad)	T (y)	$M_{\text{SMBH}} (M_{\odot})$	e	E_0 (rad)	d (Mpc)
Nearby	0.2	0.2	10.4	4×10^6	0.1435	1.14	400
Distant	$\pi/2$	0.2	10.0	4×10^6	0.1435	1.14	3000

Table 8.1: Updates to Keplerian and distance parameters for two distinct variations (*nearby* and *distant*) of the ‘Orbit 1’ example orbit of Table 4.2.

The nearby and distant systems share the same injection parameters besides the distance, d , inclination of the outer orbit, α , and a slightly different outer orbital period, T . For the BHB itself, we use our fiducial high-mass system. The parameters of the outer orbit are amendments of our ‘Orbit 1’ example, from Table 4.2. The Kepler parameters and distance are clarified in Table 8.1.

For both systems, around one full outer orbit is captured by a 10-year LISA mission. We mentioned already that signals contain more information on the Kepler parameters for certain injection parameter values: shorter orbital periods, for example. Then $T \approx 10$ y ensures the Keplerian parameters are reasonably well constrained; see Section 8.4 for further details. The SNR scales inversely with distance, so while the SNR of the distant system is ~ 15 in LISA+LIGO, the SNR of the nearby system is ~ 110 . We thus have a fair range of system properties. Unfortunately, it is not possible to use the low-mass system, since we already saw in Section 8.1.2 that more than one acceleration parameter renders the signal undetectable in LISA+LIGO. Although DECIGO is able to lower the LISA SNR detection threshold enough for multi-band analyses with two acceleration parameters to locate the low-mass system signal and constrain parameters, it is not presently known whether DECIGO would provide enough help to do the same given freedom of all six Keplerian elements. Our early attempts at doing so suggest that this is not possible.

8.3.1 Nearby system

For the nearby system in Figure 8.8, we find generally excellent constraints on all parameters. However, some very strong degeneracies remain in the Keplerian parameters, namely between all three of the α , β and M_{SMBH} parameters. This can be understood intuitively by considering the Roemer delay, whereby simultaneously increasing the SMBH




Figure 8.8: Corner plot of fully modelled high-mass system on ‘nearby’ variant of ‘example orbit 1’ (see Table 8.1), marginalised over orientation, ψ , produced with LISA(10 year)+LIGO. The primed parameters are translated injected parameters with easy to read labels defined such that the injection primed parameters, as per Table 7.1 (aside from T' defined by $T' = T - 10.4$ y) are equal to zero. The nearby system distance is $d = 400$ Mpc, giving a LISA SNR of ~ 100 and LIGO SNR ~ 40 . The resolution is very good across the parameter space; the source parameters are recovered well, despite the strong correlation between α , β and M_{SMBH} . However, the information on the Kepler parameters is strongly dependent on the source parameters, particularly the period/number of orbits observed by the detector.

mass and the inclination, and (for this particular system) decreasing the ‘yaw’ results in a similar line-of-sight acceleration as one obtains from the injection parameters. The degeneracy is also reinforced by the sinusoidal priors on the inclination, α , giving preference to values of α towards $\pi/2$. The GR effects of the gravitational and Shapiro delays, and the periastron precession help to break the degeneracies however. Recall from Section 6.3 that the modelling of GR effects is incomplete and thus often (including in this case) the constraints can be expected to be tighter for the complete model.

The sky localisation volume of the nearby system is $\sim 2.3 \text{ Mpc}^{-3}$ (90% confidence) corresponding to ~ 0.02 potential host galaxies [152] for the binary, and thus there is a high probability of a host galaxy identification. Such a GW observation with host galaxy identification would be an excellent result; if we are also able to infer constraints on the galaxy’s viewing angle [153], then we can update the priors on the Kepler orbit orientation and reanalyse the GW data to obtain an updated posterior. Moreover, with host galaxy identification and SMBH mass constraints, galaxy property (total stellar mass, luminosity) [154] and NSC property correlations [118] with the SMBH mass may be observed or tested, aiding in the pursuit of understanding galaxy structure and formation. If one is able to deduce the location of the binary in its host galaxy (i.e., the semi-major axis), these observations will also directly and strongly inform binary formation scenarios [111].

By searching over the Keplerian elements at low-frequencies, one also removes a large factor of the mass-redshift degeneracy: the peculiar velocity of the BHB in relation to the SMBH. This cannot be measured with LIGO mergers alone, but it is already encoded in the Keplerian parameters and thus can be measured by LISA in principle (and in practice, as we have seen, with our nearby and distant fiducial systems). The host galaxy’s SMBH and BHBs mass, and the BHBs outer-orbit semi-major axis are degenerate with the redshift (see Section 6.3.3).

8.3.2 Distant system

Considering the moderate SNR of the distant system in LISA+LIGO of ~ 15 and the large dimensionality of the parameter space (16 parameters in our reduced model) the posterior we obtain in Figure 8.9 is surprisingly precise. Aside from the chirp mass, which exhibits strong correlations with M_{SMBH} , α & β , there is very little correlation between




Figure 8.9: Corner plot of high-mass system on ‘distant’ variant of ‘example orbit 1’ (see Table 8.1), marginalised over ψ : with LISA(10 year)+LIGO (red) and including DECIGO (blue). LISA SNR ~ 14 , LIGO SNR ~ 6 and DECIGO SNR ~ 6600 . The main benefit of including DECIGO here is constraining q , χ_{eff} , t_c ; Kepler parameters are better constrained by factors $\lesssim 2$. Nevertheless, one still has good recovery of the source parameters. Aside from distance, the only major difference between the nearby (Figure 8.8) and distant systems is pitch, α , of the Kepler orbit. This is enough to significantly change the structure of the posterior, leaving no hint of general correlations (with the exception of \mathcal{M}_c) between intrinsic binary parameters and Keplerian parameters.

the standard binary parameters and the Keplerian parameters (correlations depend on the injected signal parameters, so this is not to be expected in general). When the effects of changing two different parameters modify the waveform in very different ways (for example, compare the effective spin which manifests in the binary’s phase as a polynomial in τ_c , in equation (5.46), to the functional form of the time-delays as shown in Figure 4.4 that encode the SMBH mass) then one can expect those parameters to be highly non-correlated. Hence, in part, the precision we observe in Figure 8.9 can be put down to the fact that the injected parameters (in particular a small semi-major axis) make for a signal model particularly well-suited for defining a precise posterior. Another way to see this is to consider Fisher information: certain injection parameters yield signals with more Fisher information/precision on certain parameters. For larger semi-major axis, where the time-delay modifications become weaker, the precision can be expected to weaken (and it does, see, for example, Figure 8.13).

The localisation volume is almost exactly the same as for the isolated high-mass system, which, as mentioned in Section 7.4.2.1 is roughly $V \approx 1.28 \times 10^4 \text{ Mpc}^3$, and corresponds to hundreds of potential host galaxies [152]. The likelihood of host galaxy identification appears poor, but as with the results discussed in Section 8.1.2, this posterior derives from an incomplete detector model that does not account for detector motion, which should help break degeneracies with the sky parameters and better locate the source. Recall this system is at a distance of 3000 Mpc ($\sim 10^{10}$ ly); there is a good possibility of observing a similar system at a smaller distance (and with correspondingly higher SNR and better sky localisation). Finally, observe that the Keplerian parameters T , e , and M_{SMBH} are individually well-constrained, so such GW systems would prove to be excellent probes of galactic structure and for testing binary formation models.

8.4 Carrier signals modified at low-frequencies

In this section we investigate the possibility of using future deci-hertz and ground-based detectors to very strongly constrain the ‘standard’ binary parameters (the masses, spins, location, orientation and eccentricity), with which we acquire precise knowledge of the earlier part of the waveform present in LISA as generated by the source, and which hence may be referred to as a ‘carrier’ signal, carrying modulations caused by the environment

of the emitting binary. These modulations can be used to probe the host environment of the GW emitter. Given the high SNRs that would be provided by a combination of DECIGO and other third-generation ground-based detectors like ET, this investigation approximates a plausible future analysis scenario. Note that while the distance is not particularly strongly constrained by DECIGO networks, as one can see in Figure 8.6 for example, the waveform itself is known very precisely; there is strong degeneracy between the distance and inclination-angle since these both strongly affect the amplitude, but the




Figure 8.10: Comparison of the ‘confirmed future detector network’, i.e., LISA+LIGO, Kepler parameter posterior (i.e., the posterior in Figure 8.9 marginalised onto the Kepler parameters, red) with a potential future network with an extremely high SNR in ground-based detectors that fix the ‘carrier wave’ (we approximate all ground-based detector measurable parameters by a delta function), measuring Kepler parameters by LISA (green). It is possible to acquire orders of magnitude improvements on some Kepler parameters by improvements in ground-based detectors alone.

amplitude itself is very well constrained: the carrier wave (including at earlier times, i.e., in the LISA band) is thus knowable, in principle, to high accuracy.




Figure 8.11: Superimposed posteriors (marginalised over period) of systems with ground-based detector measurable parameter posteriors approximated as a delta function, for orbits of various periods. For a 10 year signal, the 5, 10 and 20 years periods capture 2, 1 and 0.5 full orbits, respectively. This demonstrates how signals with shorter periods carry more information on the Kepler parameters, as well as the strong and intricate dependence of the structure of the posterior on the source parameters.

The improvements this makes to the Keplerian elements posterior is shown in Figure 8.10, in a direct comparison to the LISA+LIGO posterior marginalised over the non-Kepler parameters. Of particular importance are the SMBH mass, eccentricity and period (or semi-major axis), which are better constrained roughly by factors of 10, 2 and 2, respectively, since these parameters inform the population models and formation scenarios. If the host galaxy is located, the angular parameters α & β can be used to constrain

the viewing angle.

In fact, we can likely do significantly better than this. Recall that DECIGO is still strongly sensitive to an ‘instantaneous’ acceleration (a single acceleration parameter, see Figure 8.3), thus, as well as knowing the carrier wave, we shall also know its acceleration at a given ‘instant’. This has not been included in the model generating the known carrier wave posterior in Figure 8.10; all the acceleration is assumed only measured by LISA. We can interpret this extra constraint as essentially selecting a thin ‘slice’ through the posterior in Figure 8.10 of those parameter values consistent with the instantaneous acceleration measured by DECIGO. This could potentially lead to a much tighter constraint on the Kepler orbit parameters.

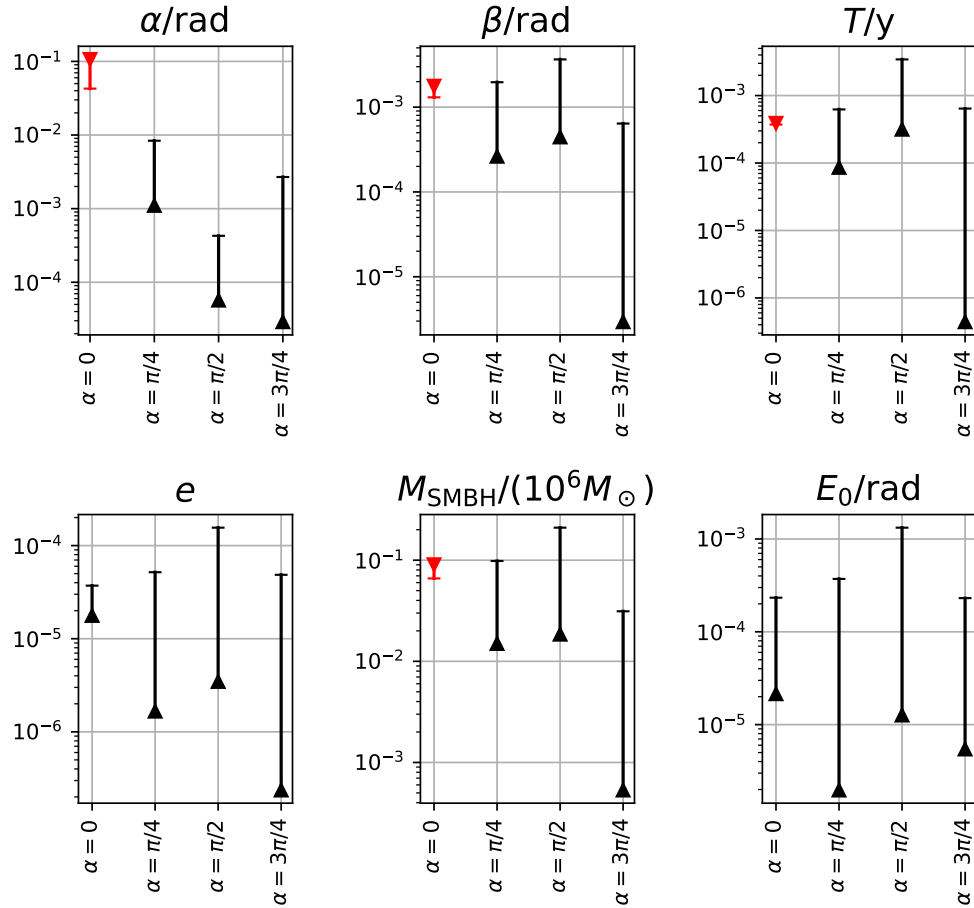


Figure 8.12: Marginal precision and accuracy plots for LISA posteriors of Kepler parameters (from known carrier wave signals) for different values of ‘pitch’, α . The red points (with $\text{PAR} > 1$) for $\alpha = 0$ are to be expected, since the prior on α is sinusoidal and excludes the injected $\alpha = 0$. The other parameters are fixed as ‘Example Orbit 1’ (see Table 4.2): these are not averages over orientations so we cannot draw general conclusions, however, this provides a modicum of evidence suggesting that the errors on the Kepler parameters are only weakly dependent on the pitch.

Figure 8.11 gives an example of posteriors as a function of the Kepler orbital period, and shows just how great a difference in parameter constraints is given by binaries in close orbits around the SMBH for a given signal length (the important factors in constraining the parameters, however, are the *number of orbits* captured by the signal and the SMBH mass, see Figure 4.7). It also demonstrates how vastly varied the (zero noise realisation) posteriors are, as functions of the source parameters.

Using our high-mass system as the carrier signal, MPA plots are shown for the same orbit (with period $T = 10$ y) for varying inclination angles α in Figure 8.12. Generally speaking, as well as smaller values of the semi-major axis, a , (or equivalently the period, T) describing a more rapidly changing time-delay, so does a smaller outer orbit inclination angle, α (see Section 4.2.2). However, whilst not much is yet known about the prior on a (or T) since this depends on population models, the prior on α is $\sin(\alpha)$, following the assumption of isotropy. Thus, the parameter constraints one acquires from the posterior do not simply reflect the degree to which the waveform depends on α , as the prior also contributes. For our fiducial system and its variations shown in Figure 8.12 then, there is indeed not as great a variability in parameter constraints as one might have initially supposed.

8.5 Approximating the time-delays

In Section 4.2.4 we discussed the possibility of describing the Kepler orbit and associated GR time-delays by a reduced parameter model (RPM). In particular, we divided the Kepler parameter space \mathcal{K} into four regions, each of which consisting of those points of the parameter space where the true time-delays (over a fixed signal duration) either must be fully modelled (\mathcal{K}_6), or for which a polynomial approximation of the time-delay is indistinguishable from the true time-delay (\mathcal{K}_{P_1} , \mathcal{K}_{P_0} , $\mathcal{K}_{P_{\text{zero}}}$). We will put these time-delay approximants to the test here, by using LISA datasets defined using the full time-delay signal model, but searching for the signal using the various approximants we have defined. We will use our high-mass fiducial BHB with the ‘Orbit 1’ Kepler orbit of Table 4.2 (with varied semi-major axis) observed by LISA+LIGO. In Figure 4.7 we can see that the boundaries between the regions of validity of the various approximants are intersected by a line of decreasing semi-major axis. We therefore produce posteriors defined by our




Figure 8.13: Marginals of the posteriors produced with fully modelled time-delay data and model (corresponding to region \mathcal{K}_6 ; see Section 4.2.4), marginalised over semi-major axis a and other parameters for clarity. We use the high-mass system and Example Orbit 1 (see Table 4.2). Notice that the variance of the chirp mass increases as the semi-major axis decreases already by 100 pc, despite unreasonably tight uniform priors on the Keplerian elements and chirp mass (the tight β and M_{SMBH} constraints at small a are likely due to having tight priors on \mathcal{M}_c and are not realistic). The unreasonable priors were required in order to acquire results in a practical time. More appropriate (wider) priors would only further decrease parameter constraints however, thus it is reasonable to state that the constraints on the intrinsic binary parameters start to degrade at around 100 pc.

different model approximants as functions of the semi-major axis to find the point at which they break down, and compare this to predictions.

We first produce the posteriors and plot the marginals of the fully modelled search in Figure 8.13. One might say this is the ‘correct’ or exact approach, since the signal is generated by the same model as the templates that are used in the search. It is very important to point out that all posteriors that rail up against the bounds in the following figures are railing up against the priors. This is especially important in Figure 8.13, since the priors on the Keplerian elements are unreasonably tight, artificially making the posteriors far too constrained. This was necessary due to the long PE convergence times (see previous section for more details). The posteriors cover all 16 parameters of our fully modelled signal and were marginalised over some parameters for clarity. Despite the very tight priors (especially on the Kepler parameters) we can still see behaviour pertinent to our investigation; the chirp mass constraints begin to degrade as the semi-major axis decreases. This occurs at around $a \sim 100$ pc. With more appropriate priors, we would see greater degradation. Note that the improvement in the constraints on β and M_{SMBH} as a decreases is in large part a product of the tight priors on \mathcal{M}_c : the posteriors are not reliable. To summarise, we have lower-bounds in Figure 8.13 on, for example,

the standard deviations of the marginals. The true marginals should be wider, yet we still see that chirp mass constraints degrade at $a \sim 100$ pc. Therefore we can also expect that with more appropriate priors the chirp mass constraints degrade at $a \gtrsim 100$ pc.

We now generate the data with the fully modelled time-delays, but search for signals using an unmodified signal model (i.e., not modelling time-delays). The result is shown in Figure 8.14, and we can see immediately that the posteriors maintain their structure down to $a \sim 10$ pc, and indeed only showing the beginnings of variation at $a \sim 1$ pc; the unmodelled time-delay analysis in this respect appears to significantly *outperform* the fully modelled analysis, even with its tight priors. This highlights the potential effectiveness, for the purposes of PE, of devising and using an optimal model.




Figure 8.14: Marginals of the posteriors produced with fully modelled time-delay data but with no time-delays present in the template model for parameter estimation (corresponding to region $\mathcal{K}_{P_{\text{zero}}}$; see Section 4.2.4). We use the high-mass system and Example Orbit 1 (see Table 4.2). There is noticeable variation appearing as the semi-major axis decreases to around $a = 1$ pc. One can acquire better precision and accuracy on the intrinsic binary parameters than using the fully modelled time-delay model (as in Figure 8.13), thus one could use the simpler model for more intrinsic binary parameter information until the semi-major axis is approximately 1–10 pc. This is broadly consistent with Figure 4.7, where we see that $\partial\mathcal{K}_{P_0}(M_{\text{SMBH}} = 4 \times 10^6 M_{\odot}) \approx 5$ pc.

Ultimately, one requires a model selection analysis to compare the performance of the different models. Unfortunately, since a proper treatment of model selection requires posteriors we do not currently possess, this falls beyond the scope of this thesis². However, whilst the non-accelerated model demonstrably performs well (to minimum values

²Although *Nessai* approximates the *evidence* of the data given a model (required to compute a Bayes factor) as part of the process of nested sampling (see Section 2.4.2), it only evaluates the integral within the prior ranges supplied, and not over the whole parameter space as is required in the definition of the evidence. The evidences returned are thus not suitable for model selection, especially between models with different parameters.




Figure 8.15: Marginals of the posteriors produced with fully modelled time-delay data, with time-delays in parameter estimation modelled as constant accelerations (corresponding to region $\mathcal{K}_{P_0^{bf}}$). The breakdown in the posterior occurs at around 0.04 pc, again consistent with Figure 4.7 which has $\partial\mathcal{K}_{P_1}(M_{\text{SMBH}} = 4 \times 10^6 M_{\odot}) \approx 0.25$ pc.




Figure 8.16: Marginals of the posteriors produced with fully modelled time-delay data, with time-delays in parameter estimation modelled as linear accelerations (corresponding to region $\mathcal{K}_{P_1^{bf}}$). The breakdown in the posterior occurs at around 0.01 pc, again consistent with Figure 4.7 which has $\partial\mathcal{K}_{P_6}(M_{\text{SMBH}} = 4 \times 10^6 M_{\odot}) \approx 0.06$ pc.

of a), we can only know that it performs well due to having full knowledge of the injected signal. Further, clearly we already know what the correct model is, so for computing the Bayes factor, say, the prior odds are absolutely dominated, yielding a probability of 1 that the full Kepler time-delay model is correct. But we are interested in models that can reproduce the waveform indistinguishably from the true model in as few parameters as possible in order to minimise the posterior evaluation effort and maximise the information on the binary parameters (by removing degeneracies with the Kepler parameters). In practice, by examining Figure 8.14, this appears to be feasible. We note that models that cannot accurately reproduce the waveform will have poor fits to the data and correspondingly lower maximum likelihood values. This provides an easily accessible measure of goodness of fit of a model; maximum log likelihood values for the evaluated posteriors are shown in Table 8.2.

The marginals of posteriors generated by searching with a constant acceleration model are plotted in Figure 8.15, where the posterior appears to start breaking down at around $a \sim 0.06$ pc, and the marginals of posteriors generated by searching with a linear acceleration model are plotted in Figure 8.16, with posteriors breaking down at around $a \sim 0.01$ pc. The maximum log likelihoods found by the sampler for these posteriors are shown in Table 8.2. Therefore in practice we have been able to dig a little deeper into

a (pc)	Number of time-delay model parameters			
	6 (control)	0	1	2
1×10^{-2}	-4.32	-24.10	-	-32.17
2×10^{-2}	-	-	-46.54	-5.02
3×10^{-2}	-	-	-15.46	-4.50
4×10^{-2}	-	-	-8.73	-4.68
5×10^{-2}	-	-	-4.79	-4.29
6×10^{-2}	-	-	-4.24	-4.27
7×10^{-2}	-	-	-4.25	-4.44
8×10^{-2}	-	-	-3.88	-4.44
9×10^{-2}	-	-	-4.22	-4.36
1×10^{-1}	-4.19	-76.29	-4.40	-4.22
1×10^0	-4.16	-4.30	-	-
1×10^1	-4.21	-4.18	-	-
1×10^2	-4.16	-4.38	-	-
1×10^3	-4.10	-	-	-
1×10^4	-4.16	-	-	-

Table 8.2: Maximum log likelihoods sampled by *Nessai* for all time-delay models. The ‘control’, or fully modelled time-delays described by the six Keplerian parameters have very similar maximum log likelihood values (recall we are using zero noise realisations) for any semi-major axis, a , as the sampler is always able to closely match the injected signal with a model template. The approximants, on the other hand, are able to match the signal accurately over a limited region of the parameter space. As one decreases a , there are quite clear jumps in the maximum log likelihoods to lower values, where the model cannot very well match the injected signal (red numbers). The jumps correlate strongly with the visual break-downs in the posterior marginals in Figures 8.14–8.16. Some peculiar values are highlighted (light blue background); these log likelihood values are unexpectedly high, the reasons for which are not clear, but may be an artefact of the downsampling procedure being employed.

smaller semi-major axes than those suggested in Figure 4.5 (we can allow IADs up to 30 m s^{-1} rather than the 3 m s^{-1} plotted, due the mass-spin degeneracy as mentioned in Ref. [26]). We cannot make a direct comparison to Figure 4.7 since the alignments of the SMBH, binary and observer are slightly different between the systems analysed in that figure and the results evaluated in this section, but there is a good order of magnitude agreement in the range of validity of the various approximants. Thus it appears, from the results of this preliminary investigation, that one can acquire better binary parameter constraints by using the time-delay approximants for the corresponding regions \mathcal{K}_6 , $\mathcal{K}_{P_1^{\text{bf}}}$, $\mathcal{K}_{P_0^{\text{bf}}}$ & $\mathcal{K}_{P_{\text{zero}}}$.

8.6 Modified gravity with acceleration

For our final posterior analysis, we include both a single (constant) acceleration parameter and the EdGB modified gravity parameter as free parameters of the model. The high-mass fiducial system is used as the base binary and the signal is observed by the




Figure 8.17: Corner plot of the high-mass, isolated GR binary posterior, generated with model searching simultaneously over a constant acceleration, A , and EdGB gravity, ζ_{EdGB} , marginalised onto the A - ζ_{EdGB} plane, using LISA+LIGO+DECIGO (recall that DECIGO is required to provide useful constraints on ζ_{EdGB} ; see, for example, Section 7.2.2.1). There is strong negative correlation for small values of A and ζ_{EdGB} , highlighting the importance of correctly accounting for acceleration in modelling; one may otherwise arrive at incorrect conclusions about EdGB gravity (or other modified gravity theories: see Section 8.6).

LISA+DECIGO+LIGO network with its very high SNR ~ 6600 , noted in Table 7.5. We mentioned the potential degeneracy between the acceleration and modified gravity parameters in Section 5.6 and this is borne out clearly in Figure 8.17: the marginalisation of the posterior onto the $A-\zeta_{\text{EdGB}}$ parameter axes; we observe a strong negative correlation between the constant acceleration and EdGB parameters. Note that for this posterior, the sky position was overconstrained in the priors and is therefore roughly consistent with a signal originating from a binary with the host galaxy having been identified. Thus, if one neglects modelling acceleration when it is not known that acceleration is negligible, one risks obtaining a false positive result for EdGB gravity, and vice versa.

8.7 Discussion

The majority of results presented in this chapter derive from the high-mass fiducial system; the high-mass GW sources are potentially rich sources of information for both probing galactic environments and testing alternative theories of gravity. We have seen in Figure 8.9 that even for relatively distant sources, one is able to obtain quite impressive constraints on the Keplerian elements using planned future detector networks (i.e., LISA+LIGO, where LIGO shall be superseded by more sensitive ground-based GW observatories). However, those constraints are heavily dependent on the actual orbital parameters themselves, so an element of good fortune is required to observe such ideal or ‘golden’ signals carrying abundant information. One can expect the Keplerian elements to be generally well-constrained when the orbital period is of order the signal duration as demonstrated in Figure 8.11 (aside from some special cases, such as when the axis of rotation of a circular orbit coincides with our line-of-sight to the galaxy) since small changes in the parameters are not as easily absorbed into other parameters as when one only sees a small fraction of the orbit, that is, when the period $T \lesssim T_{\text{obs}}$, where T_{obs} is the signal duration. By equation (4.14), we are then limited, roughly, to acquiring strong constraints on the Keplerian elements by

$$a \lesssim \left(\frac{GM_{\text{SMBH}} T_{\text{obs}}^2}{4\pi^2} \right)^{1/3}. \quad (8.3)$$

so that for a given M_{SMBH} , there is an approximate upper limit on the semi-major axis of the orbit, which may lead to SMBH mass dependent time-delay model selection ef-

fects which should be taken into consideration when using the constraints on the Kepler parameters to inform galaxy structure models. One also obtains an approximate maximum SMBH mass for which Kepler parameters of an orbiting binary can be measured for a given signal duration, by taking the shortest orbital period to occur at the ISCO radius of the SMBH, defined by $R_{\text{ISCO}} = 6GM/c^2$ [26]. Substituting this for the semi-major axis, we require

$$M_{\text{SMBH}} \lesssim \left(\frac{c}{\sqrt{6}} \right)^3 \frac{T_{\text{obs}}}{2\pi G} \approx 6.9 \times 10^{11} M_{\odot},$$

for measuring Kepler parameters of 10 year signals. No known SMBH is as massive, so this limit does not rule out potential ability to measure BHB time-delays from any galaxy. Since binaries at small semi-major axes may reside in accretion disks, binaries on eccentric orbits of SMBHs may also provide potential to probe the accretion disk density profile, allowing one to test various models via model selection analyses. For orbital periods $T \gtrsim T_{\text{obs}}$, the time-delay parameters tend towards becoming nuisance parameters that somewhat spoil our PE analyses; we shall be limited in the conclusions we can draw from these detections about the environments of galaxies with lower mass SMBHs, or from BHBs with wider orbits. However, we saw that the time-delay approximants introduced in Section 4.2.4 can minimise the degree to which posteriors are ‘degraded’ (whilst simultaneously speeding up the PE process due to the lower dimensionality) as a result of removing degeneracies between the Kepler and the binary parameters.

For the high-mass systems, the low SNR in LIGO does not pass the detection threshold for LIGO alone of $\text{SNR} \sim 8$. However, future ground-based detectors, in particular ET, scheduled to commence operations in 2035 with a 50 year lifetime, and thus operating at an appropriate time to capture mergers of the inspirals in LISA, should bring the noise floor LIGO currently provides down by a factor of a few, and with more sensitivity to lower frequencies. This will greatly increase the SNR of our high-mass fiducial system, making it easily detectable in ET alone. This is important since the parameter space for time-delayed systems is extremely large, particularly considering searching for coalescence times within a few decades, making it difficult to find signals in LISA alone. With a detection in ET, one will know that the signal exists in LISA and searches (using the ET posterior as a prior for LISA) using the various time-delay models will be tractable. Without an ET detection, one may have a troublesome time locating a signal in LISA, since the priors on the coalescence time are strictly unbounded.

Analysing signals of low-mass systems is clearly more difficult due to the very low SNRs one can expect in LISA. We saw that, for a 4 year LISA signal in a detector network with LIGO, only a constant acceleration model yielded a detection. If a second acceleration parameter is modelled, the LISA SNR detection threshold increases such that the signal can no longer be found in LISA (although a 10 year LISA mission successfully locates the signal). Supposing a low-mass signal's time-delays *require* a linear acceleration (that a constant acceleration model is not suitable), then it would not be detected at all in a 4 year LISA signal using any model. However, with a merger detected in LIGO, we know that the signal must be present in LISA. Then a non-detection itself can be helpful; for a detection in LIGO, one has two possibilities: a detection in LISA, in which case both the eccentricity and acceleration are likely low, thus preferring field formation of the binary, or non-detection in LISA, where the eccentricity [144] and/or acceleration are likely non-negligible, thus preferring dynamical formation.

Including DECIGO in the network obviously significantly strengthens all parameter constraints, either directly, or by lowering SNR detection thresholds. The effects of DECIGO on posteriors for time-delayed signals is similar to that we saw for the posteriors of isolated binaries in the previous chapter, with good sky localisation and host galaxy identification potential for binaries, given they are close enough. However, in the fully modelled time-delay posterior, although the Kepler parameters are significantly improved, there is a trade-off in that the sky localisation constraints are degraded. Being that the signal only exists in DECIGO for ~ 0.52 days (see Table 7.5), the detector's motion, modelled as being stationary in Earth's reference frame, can easily be degenerate with Kepler parameters that replicate the time-delays on the source side, rather than the detector side. Properly modelling DECIGO and its base station's motions may help break those degeneracies however (although the details of the DECIGO spacecraft, and the mission itself, are not yet confirmed), and without much of an effect on the other parameter constraints, due to the very high SNR.

Final Remarks and Outlook

Multi-band GW data analysis has far-reaching potential applications and benefits for astronomy, cosmology and fundamental physics. A primary aim of this thesis has been to make progress in revealing features of the posterior probability distributions of multi-band signals, and the expectations one can reasonably hold about future multi-band detector network data analysis.

As preparation for our envisaged analyses of long-duration LISA signals (i.e., those with very large datasets), we first had to bypass the bottleneck of computing time and resources at the parameter estimation stage. We proposed and tested a procedure applicable to *simulations* of the real-world data analysis setting to drastically decrease the computing time required to produce posteriors. This was achieved in Chapter 3 and found to be effective, reducing the likelihood evaluation times and thus PE convergence times by factors of hundreds to thousands (as compared to our estimates of the frequency domain likelihood evaluation times) whilst reproducing the likelihood function to high accuracy. This relatively fast approach to estimating posteriors allowed for large-scale analyses, such as high dimensional PE and detailed inquiry of *parameterised families* of posteriors, for example, defining and evaluating posteriors as functions of the detector noise amplitude (equivalently, signal SNR), or of one of the model's injection parameters, as we did throughout Chapters 7 & 8.

Although this downsampling procedure is fast and accurate, the scope of its validity is not very well understood; it was developed for and tested on signals that are *slowly evolving*, a property that was only ever loosely defined. Moreover, the reasons determining

the optimal downsampling rate (or optimum number of samples) are unclear. Further investigations to clear up these issues may prove to be enlightening and reassuring, especially in terms of allowing for more general applications. However, for our BHB systems, the performance of the results of applying the downsampling procedure were extremely consistent throughout all of the analyses presented in this thesis, including for the high dimensional (15 dimensional) LISA posteriors. As far as the study of slowly evolving waveforms in LISA goes (for data containing a signal from a single source), there does not presently appear to be any pressing requirements or practical advantages for further investigations of these features and dependencies of the downsampling procedure.

We have investigated modelling expectations and requirements for long-duration (the lower-frequency part of low to intermediate mass BHBs) binary signals, detectable, in principle, by LISA. As soon as one includes the inspiral part of a CBC observed by LISA in the analysis of the merger observed by LIGO, the number of potential detectable environmental effects jumps from zero to many; the mergers in LIGO last seconds and arise from a small spacetime region with an effectively constant background geometry (as if emitted from a single point), whereas the inspiral signals in LISA lasting a few years are emitted along a curve through spacetime. This extra dimension provides space for significant waveform modifications to enter the waveform; we gave a short introduction to the variety of such modifiers in Section 4.1.

Although it is important to develop a comprehensive understanding of the interplay of these effects, we are rather limited in scope in the present work by the vast scale of the fully modelled problem. We focused on the standard (first-order) GR time-delays that arise from a binary set in a background Schwarzschild geometry (with important caveats pointed out in Section 6.3) as these are the most significant waveform modifiers for some of the GW sources we expect to observe both with LISA and ground-based observatories. Since multi-band GW astronomy is touted as being a particularly useful tool for testing gravity theories, we also included the EdGB and dCS modified gravity parameters to examine their posteriors and test this assertion.

Our initial set of investigations looked into the *isolated binary* or ‘first-order’ applications of multi-band GW data analysis to CBCs, which exclude the GR and orbital motion time-delays. We confirmed that GW150914-like binaries in (4-year) LISA data, with SNRs of around 4, sit on or near the detection threshold. If the SNR is high enough,

one can achieve orders of magnitude improvement on the masses, spins and sky localisation, which, with knowledge of host galaxy redshift, could lead to improvements in the estimate of the Hubble parameter (although many sources need to be observed at a given distance to alleviate the problem of peculiar motion of the source). However, many sources are likely to be too quiet for multi-band detection, especially if there is appreciable eccentricity, since this decreases the SNR significantly. We should hope for an extended LISA mission duration of 10 years to maximise detections.

Unexpected behaviour of the posteriors at very low LISA SNR (SNR below the detection threshold) was discovered in the LISA+LIGO network. The posteriors seem to ‘become worse’ by including the low SNR LISA signal. It appears that this is due to the ‘incorrect’ posterior on the binary’s orientation from LIGO, with strong bias for values away from the truth. This overpowers and suppresses the primary posterior mode of LISA, forcing a secondary mode of LISA to become more dominant, preferring values of time, distance and sky position away from the truth. This is an interesting result that requires a further investigation into when, or even if, it is reasonable to rely on results that derive from LISA signals with SNRs at and below (or even near) the detection threshold.

We found that modified gravity parameters are particularly troublesome for networks comprising the currently confirmed future detectors, including LISA and LIGO(-like) detectors. The leading order PN term of the EdGB parameter in the phase depends on the properties of the binary components in such a way that the modification vanishes when the component masses and spins are equal. Thus a GR waveform analysed with an EdGB model gives a strong bias to $q = 1$. Vastly increasing the SNR, or, equivalently, including DECIGO in the network, strongly suppresses this tendency of $q \rightarrow 1$ however, simultaneously suppressing the ζ_{EdGB} prior that appears in the posterior when $q \rightarrow 1$ (the prior permeates through to the posterior when $q = 1$ since ζ_{EdGB} is uncoupled from the signal there). Sub-leading order PN terms in ζ_{EdGB} might alleviate this bias as they should likely retain a presence in the waveform when $q = 1$. The outlook for constraining EdGB gravity therefore depends far more strongly on SNR than detections in multiple frequency bands, so for the foreseeable future, unless a detector such as DECIGO is approved, these constraints shall likely rest with the future planned ground-based detectors. Unfortunately, dCS is even harder to constrain, as we found no improvements to the currently existing constraints even using DECIGO. However, it is very likely that

sub-leading order PN terms in the dCS waveform will help matters when they are known.

Introducing additional modifications in the form of time-delays to the model, with the simplest (detectable) time-delay being (equivalent to) a constant acceleration, we saw that the LISA+LIGO network was able to recover the parameters well. However, with a linear acceleration, the LISA SNR threshold for the LISA+LIGO network rose above the LISA signal SNR for our low-mass fiducial system, and the signal became unrecoverable. One must either use a 10 year LISA signal to recover linear acceleration, or else one can include DECIGO in the network to reduce the SNR detection threshold. For the low-mass binary undergoing small line-of-sight accelerations with magnitudes less than around 10^{-6} m s^{-2} with a 4-year observation period in LISA, the acceleration can be safely neglected. It then follows that, for the LISA+LIGO network, given a detection in LIGO and inferring the LISA SNR for ranges of eccentricities and accelerations, the detection itself (or non-detection) of the signal in LISA can be used to inform formation scenarios. One can think of this as mitigation of selection effects in LISA via detections of the mergers in LIGO.

The full extension of the model to the six-parameter (six Keplerian elements) time-delay clearly also requires significant SNR in LISA, since introducing further parameters generally increases the SNR detection threshold. We should therefore not expect to be able to find signals in LISA of a low-mass system with the LISA+LIGO network unless at very close range. However, the higher mass systems (we primarily studied a 10 year LISA signal, but at large distance) of course far more often have SNRs which clearly sit comfortably over the detection threshold. They therefore constrain the constant and linear acceleration, and the fully extended time-delay models easily, although with the important caveat that the ability to constrain the Keplerian parameters strongly depends on the source parameters of the system.

In terms of testing binary formation scenarios and galactic structure, the overall situation is fortunate in the sense that the proposed scenarios, favouring low-mass binaries forming both in the field and dynamically, and high-mass binaries only forming through hierarchical mergers in dynamical environments, generally align with the detectability expectations of the time-delay parameters across the mass range. Thus, prospects for discriminating between scenarios are bolstered *if* the currently expected scenarios are roughly correct. Moreover, high-mass BHBs are thought to form via BHs migrating to

AGN discs, at which place the time-delays are considerable and Kepler parameters more readily measurable. If this picture is correct, the prospects of observing GWs that can be used as probes of the Kepler parameters, and thus the AGN disc region are improved.

We also found that, for certain regions of the Kepler parameter space, one can significantly enhance the intrinsic binary parameter constraints by using low-order polynomial approximants to describe the time-delays. The true time-delay imprinted upon a GW may be very well approximated by, say, a quadratic or cubic in time, requiring only one or two parameters to describe (since linear modifications in time are not detectable). By substituting the Kepler model for an approximant, one already benefits from the smaller dimensionality of the parameter space, but degeneracies of the intrinsic binary parameters with the Kepler parameters are also removed. Although the approximant parameters are essentially nuisance parameters, one obtains better constraints on the binary's parameters whose useful information is still embedded in the signal. Further work on model comparison and selection is required to determine how well the performance of these approximants can be *known*; it is clear to us which models perform well since we know the injection parameters, but we must of course be able to determine the model and its parameters of an unknown signal. There are potentially more natural and better basis functions with which to define a (truncated) series for expanding the time-delays, for example, Chebyshev polynomials, which may more readily describe the arcs of the orbits, and some investigation of this may be useful.

The combination of efficient sampling methods and vast reduction in likelihood evaluation time (using the downsampling procedure explicated in Chapter 3) drastically decreases the posterior convergence time. Even so, some of the posteriors produced took up to weeks to converge. Despite our efforts, we were unable to solve this problem; the posteriors appeared as 'nice', smooth functions, as one would expect and which can be seen by inspecting the likelihood values at the returned posterior sample locations (as well as by 'manually' inspecting the likelihood function), and there does not appear to be any modelling errors in the code. However, the number of likelihood evaluations required by the sampler increased far more than might be reasonably expected when increasing the number of free parameters in the model (especially on adding multiple Kepler parameters). Posteriors produced using 9 intrinsic binary parameters completed in a few hours, with around 5×10^5 likelihood evaluations, whereas the 15 parameter model required up

to as much as ~ 40 days and 2×10^8 likelihood evaluations. This could be simply due to the difficult shape of posteriors and high dimensionality, however, this unresolved issue, likely to do with sampling, requires further investigation.

A shortcoming of the results presented here is that the merger of our fiducial systems always occurs soon after the signal has left the LISA band. This is of course generally not the case, and there could be decades between a detectable signal in LISA and merger observed on the ground. Such multi-band detections would essentially yield decades long observation periods, able (in principle) to provide stronger constraints on both the intrinsic binary parameters and the Kepler parameters; small changes in the parameters will accumulate, leading to greater modifications over longer periods of time. However, for a given signal, a greater period between observed inspiral and merger significantly weakens the possibility of correctly identifying the inspiral's corresponding merger signal, since both the average instantaneous SNR and the coalescence time information decreases as one goes to earlier times in the inspiral.

The BHB science case for the space-based detectors has been somewhat clarified by the results presented here. If one is interested in galactic nucleus structure and dynamics then LISA is preferred, as we have seen one can successfully recover the Kepler parameters of high-mass binaries in NSCs. If one is instead interested in very precise measurements of the 'isolated binary' parameters, then DECIGO should be preferred, although it can still measure a constant acceleration and thus constrain formation scenarios to some degree. Of course BHBs are not the only GW source of interest; a primary target of LISA, for example, is merging SMBHs, which on its own is highly compelling. For BHB science, ideally we would have a detector network consisting of all possible detectors, with LISA launched first, and DECIGO and other ground-based detectors set to be in operation as the LISA mission comes to an end, optimised to capture as many full IMR sequences as possible.

From the seemingly endless set of possibilities for both astrophysical GW sources and future detector network scenarios that can be dreamt up and investigated, we have named and tested a few. It has been demonstrated that there is great potential in multi-band gravitational wave astronomy if the right sources are present, potentially simultaneously providing insights into galactic structure, cosmology and fundamental physics. There is still a great deal of work to go in developing and testing the GW models, and

then whittling down those models, essentially by placing priors on them, derived from studying galactic structure, binary formation scenarios and dense star cluster dynamics.

We have also seen that parameter estimation can be aided with ‘optimised modelling’, but that there remain aspects of the analysis and modelling, and subtle properties of the posteriors that require careful examination and better understanding. The LISA data analysis (and hence multi-band analysis) task that remains is immense, but the overall picture is becoming clearer, and we still have over a decade until the LISA spacecraft launch and a further few years of data collection before multi-band GW astronomy becomes a reality. During this period we can be sure to expect new advances in computing, approaches to analysis, and an extensive knowledge of gravitational waveforms to be developed, as we prepare to enter another new phase in our understanding of the Universe.

Closed-form approximations of IADs

We shall compute very rough approximations of IADs, as defined in equation (4.39). We ignore all GR contributions to the time-delay and restrict attention to circular orbits only, so that $\Delta t_{\text{tot}} = \Delta t_{\text{Roemer}} = d_{\text{Roemer}}/c$, where $\Delta t_{\text{Roemer}} = \Delta t_{\text{Roemer}}(t, \alpha, a, M_{\text{SMBH}}, E'_0)$. The circular orbit defines a sinusoidal line-of-sight distance with a period given by Kepler's third law of

$$T = \sqrt{\frac{4\pi^2 a^3}{GM_{\text{SMBH}}}}, \quad (\text{A.1})$$

with amplitude (maximal line-of-sight distance from the SMBH) a , multiplied by a factor of $\cos(\alpha)$ due to the orbit's inclination α . The line-of-sight distance (from the SMBH distance) is thus given by

$$d(t) = -a \cos(\alpha) \sin(E'_0 + 2\pi t/T), \quad (\text{A.2})$$

and line-of-sight acceleration is thus

$$a(t) = \frac{d^2 d}{dt^2} = \frac{4\pi^2}{T^2} a \cos(\alpha) \sin(E'_0 + 2\pi t/T). \quad (\text{A.3})$$

The section of the orbit we observe of which acceleration we approximate as a first order polynomial begins at time $t = 0$ and lasts 10 years. Now, rather than computing $P_1^{\text{bf}}(a)$ precisely, we will make a rudimentary approximation of it, $P_1^{\text{app}}(a)$, defined to be the line tangent to a at the half-way point of the signal (i.e., 5 years through a 10 year signal) as illustrated in Figure A.1; the blue line shows this linear approximation of the acceleration over the signal. We have further defined $E'_{1/2} = E'_0 + 2\pi \frac{5y}{T}$ as the 'mid-signal' eccentric




Figure A.1: An example of a full period of line-of-sight acceleration of a circular orbit (black curve), and an observed section between E_0 and E_1 , where acceleration is approximated by the blue line.

anomaly, and $E'_1 = E'_0 + 2\pi \frac{10y}{T}$ as the ‘final’ eccentric anomaly. The equation for our linear approximation, $P_1^{\text{app}}(a)$ is

$$\begin{aligned} P_1^{\text{app}}(a) &= \left. \frac{d \cdot a}{dt} \right|_{t=5y} (t - 5y) + a(5y) \\ &= \frac{4\pi^2}{T^2} a \cos(\alpha) \left[\frac{2\pi}{T} \cos(E'_{1/2}) \cdot (t - 5y) + \sin(E'_{1/2}) \right]. \end{aligned} \quad (\text{A.4})$$

Since it would be very useful to dispose of the requirement for the absolute value operation in the integral, we demand further restrictions. Firstly, in order for the integrated difference between the sinusoidal acceleration and linear approximation not to always vanish when $E'_{1/2} = 0$ (difference between 2 odd functions), we shall compute the difference for one half of the signal only, for example, in Figure A.1, just the solid blue line and not the dashed blue line, and compare this to half of the IAD required, that is, for example, taking roughly half of each side of equation (4.46) to produce the new criterion

$$I_1^{\text{app, half sig}} = \int_{t(E'_{1/2})}^{t(E'_1)} dt |\ddot{\alpha}(t) - P_1^{\text{app}}(\ddot{\alpha}(t))| \approx 1.5 \text{ m s}^{-1}. \quad (\text{A.5})$$

However, one can then imagine situations in which the signal length is of order T ; the integrated difference (without taking the absolute value) of the curve and the solid blue line in Figure A.2a, for example, may well evaluate to around zero, which would clearly be a poor estimate of the IAD. Taking advantage of the symmetry properties of the sine function and the problem at hand, we can just restrict the allowed values of $E'_{1/2}$ to being between $0 < E'_{1/2} < \pi/2$, as per Figure A.2b and simultaneously restricting the




Figure A.2: An example of a situation in which the integral of the difference between the acceleration and a linear approximation (without taking the absolute value) fails to accurately approximate the IAD (left) and an example of a linear approximation to acceleration given the restricted range of parameters (right): $0 \leq E'_{1/2} < \pi/2$, and $E'_1 - E'_{1/2} < \pi/2$ (right).

signal length such that $E'_1 - E'_{1/2} < \pi/2$. We can then write the integral as

$$I_1^{\text{app, half sig}} = \int_{t(E'_{1/2})}^{t(E'_1)} dt P_1^{\text{app}}(\ddot{a}(t)) - \ddot{a}(t) \approx 1.5 \text{ m s}^{-1}, \quad (\text{A.6})$$

which, finally, can easily be computed analytically.

We can expect that the many assumptions and approximations made here shall render this approach rather inaccurate as compared to the fully modelled numerical analysis. However, since the linear approximations are not true lines of best fit, the IAD will be greater in general and the criteria thus stronger; there is less danger of concluding that only two parameters are required for points that should in fact be within \mathcal{K}_6 . Most importantly however, we are now able to derive some relationships that provide, in a comprehensible way, an approximation of the surface $\partial\mathcal{K}_6$, so that we have some idea of when we must use the full six orbital elements in our parameter estimation of a signal. The approximate form of equation (4.46) shall be written

$$I_1^{\text{app}} \equiv \int_{t(E'_{1/2})}^{t(E'_1)} dt P_1^{\text{app}}(a(t)) - a(t) = 1.5 \text{ m s}^{-1}. \quad (\text{A.7})$$

Substituting equations (A.3) and (A.4) and integrating, the condition for which using a




Figure A.3: The shaded regions show, for various values of $E'_{1/2}$, the combinations of semi-major axis and SMBH mass that result in an approximate IAD (for half of a 10 year signal) of less than 1.5 m s^{-1} for $\alpha = \pi/4$. There is little dependence on $E'_{1/2}$ due to the fact that, in this range, the periods are on the order of hundreds to thousands of years; a 10 year signal represents a very small section of that orbit, the observed acceleration over which thus being approximately linear irrespective of initial position. There is good agreement with the suggestion from the example systems in Figure 4.5, where linear acceleration models can be used when the observed fraction of orbit is between 10^{-2} to 10^{-3} ; for 10 year signals, this corresponds to periods of 10^3 to 10^4 years.

linear approximation to describe the acceleration for a 10 year signal may be written

$$\begin{aligned}
1.5 \text{ m s}^{-1} &\geq \frac{4\pi^2}{T^2} a \cos(\alpha) \int_{t(E'_{1/2})}^{t(E'_1)} dt C(t - 5 \text{ y}) + \sin(E'_{1/2}) - \sin(E'_0 + 2\pi t/T) \\
&= \frac{4\pi^2}{T^2} a \cos(\alpha) \left[C\left(\frac{1}{2}t^2 - 5t \text{ y}\right) + \sin(E'_{1/2})t + \frac{T}{2\pi} \cos(E'_0 + 2\pi t/T) \right]_{5 \text{ y}}^{10 \text{ y}} \\
&= \frac{4\pi^2}{T^2} a \cos(\alpha) \left[\frac{25\pi}{T} \cos(E'_{1/2})y^2 + \sin(E'_{1/2})5 \text{ y} + \frac{T}{2\pi} \left(\cos(E'_1) - \cos(E'_{1/2}) \right) \right],
\end{aligned} \tag{A.8}$$

where $C = \frac{2\pi}{T} \cos(E'_{1/2})$. Figure A.3 gives a depiction of regions (shaded) in which the linear model of acceleration is safe to use. For orbits with parameters outside of these regions, the accelerations must be described using the six Keplerian elements. Note, however, that we have erred on the side of caution at every step (apart from, of course, ignoring the GR elements of the time-delays); the true \mathcal{K}_6 is likely to be a significantly smaller region.

The estimates we have derived and shown in Figure A.3 are well into the ‘safe’ region




Figure A.4: An example of a constant approximation to acceleration; the integrated difference over the dashed blue line gives a poor representation of the IAD, so use the first half of the signal over the given restricted range, $0 \leq E'_{1/2} < \pi/2$ (solid blue line).

(only as far as ignoring GR time-delay contributions) and perhaps too conservative if one compares those results to the example orbits considered in Figure 4.5. There, one can see that linear approximations of the (ascribed) accelerations for all three systems are appropriate when the observed fraction is anything less than 10^{-3} , which, for a 10 year signal, corresponds to a period of $T = 10^4$ y; the red dashed line in Figure A.3.

It is of course much easier to compute the IAD for zeroth order polynomials ($P_0 = \text{const.}$) and for the zero polynomial ($P_{\text{zero}} = 0$), and to evaluate them precisely or estimate them with a higher degree of accuracy. The best-fit constant for a function over an interval is its average value, which can be computed by dividing its integral by the interval range. Thus, using equation (A.3), we have

$$P_0^{\text{bf}} = \frac{1}{10 \text{ y}} \int_{\text{signal}} dt \, a(t) = \frac{1}{10 \text{ y}} \int_{\text{signal}} dt \, \frac{4\pi^2}{T^2} a \cos(\alpha) \sin(E'_0 + 2\pi t/T). \quad (\text{A.9})$$

However, we still have the problem of finding general analytic solutions of the IAD, since the integrand contains an absolute value. We again therefore roughly approximate P_0 , setting this equal the value of a at the mid-point, $E'_{1/2}$ of the signal

$$P_0^{\text{app}}(a) = a(t(E'_{1/2})), \quad (\text{A.10})$$

then evaluate the difference over the first half of the signal subject to the constraints $0 \leq E'_{1/2} < \pi/2$, and $E'_1 - E'_{1/2} < \pi/2$. That is, similar to equation (A.5), we take the limit

to be

$$\begin{aligned}
1.5 \text{ m s}^{-1} &\approx I_0^{\text{app,half sig}} \equiv \int_{t(E'_0)}^{t(E'_{1/2})} dt P_0^{\text{app}}(\ddot{\alpha}) - \ddot{\alpha}(t) \\
&= \frac{4\pi^2}{T^2} a \cos(\alpha) \left[\sin(E'_{1/2})t + \frac{T}{2\pi} \cos(E'_0 + 2\pi t/T) \right]_{0y}^{5y} \\
&= \frac{4\pi^2}{T^2} a \cos(\alpha) \left(\sin(E'_{1/2}) \cdot 5y + \frac{T}{2\pi} \left[\cos(E'_{1/2}) - \cos(E'_0) \right] \right). \quad (\text{A.11})
\end{aligned}$$

See Figure A.4 for an example of this approximation.

The approximation of the zero acceleration IAD condition, given the same range constraints on $E'_1 - E'_0$ and T as above, should use the second half of the signal (so that the integrand never changes sign) and is given by

$$\begin{aligned}
1.5 \text{ m s}^{-1} &\approx I_{\text{zero}}^{\text{app,half sig}} \equiv \int_{t(E'_{1/2})}^{t(E'_1)} dt \ddot{\alpha}(t) \\
&= \frac{2\pi}{T} a \cos(\alpha) \left[\cos(E'_{1/2}) - \cos(E'_1) \right]. \quad (\text{A.12})
\end{aligned}$$

The ranges of application of the different acceleration models with boundaries given by the constraint equations (A.8), (A.11) and (A.12) are shown in Figure 4.7. Of course if one uses a different set of acceleration models (rather than standard polynomials) then these regions would be enclosed by boundaries defined by different isosurface constraint equations.

Shapiro Delay

Shapiro delay, the relativistic time-delay arising due from the extra propagation distance due to curvature of spacetime, is formulated in approximate form by considering the situation pictured in Figure B.1, for null rays propagating between emitter and observer.



Figure B.1: Shapiro’s original problem.

The (one-way) Shapiro delay was computed by Shapiro to be [124] the quantity given in equation (4.30):

$$\Delta t \approx \frac{2GM_{\text{SMBH}}}{c^3} \left(\ln \left[\frac{z_p + (z_p^2 + r_0^2)^{1/2}}{-z_e + (z_e^2 + r_0^2)^{1/2}} \right] - \frac{1}{2} \left[\frac{z_p}{(z_p^2 + r_0^2)^{1/2}} + \frac{z_e}{(z_e^2 + r_0^2)^{1/2}} \right] \right), \quad (\text{B.1})$$

to first order in the Schwarzschild radius $r_s = 2GM_{\text{SMBH}}/c^2$, where z_p , z_e and r_0 are, respectively, the distance between the orbiting body and periastron, the distance “along the line of flight” between observer and periastron, and the impact parameter.




Figure B.2: Time-delay due to situation of emitter on ‘near-side’ of SMBH.

However, we will also be interested in the travel time-delay caused by proximity to mass for full orbits of the emitter, including the situation in which the emitter is between the massive object and the observer, as in Fig.B.2

We can use the framework already derived by Shapiro by considering breaking up the path described in the original problem, Fig. B.1, into two parts: the Fig.B.2 part of the trajectory, and the difference between the two, illustrated in Fig.B.3. That is, Fig. B.1 is equal to Fig. B.3 plus B.2.




Figure B.3: Path difference between near-side and far-side position.

The path of the original problem and that of the difference are both described by Shapiro's framework, which we can use to calculate the time-delay from the path when emitter is on the near-side (NS) of the massive object, by subtracting the 'difference' from the 'original'.

We use three symbols to keep track of the different delays. For Shapiro's original setup, with emitter on the far-side (FS) as in Fig.B.1, we call the time-delay Δt_{FS} , for the case with emitter on the near-side as in Fig.B.2, we call the time-delay Δt_{NS} , and for the difference between these delays as in Fig.B.3, we write the Shapiro delay difference between these paths as

$$\Delta t_{\text{diff}} = \Delta t_{\text{FS}} - \Delta t_{\text{NS}}. \quad (\text{B.2})$$

Using (4.33) and the fact that, for Δt_{diff} we have $r_p = r_e \Rightarrow z_p = z_e$, then

$$\begin{aligned} \Delta t_{\text{diff}} &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{z_p + (z_p^2 + r_0^2)^{1/2}}{-z_p + (z_p^2 + r_0^2)^{1/2}} \right] - \frac{z_p}{(z_p^2 + r_0^2)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{z_p^2 + 2z_p(z_p^2 + r_0^2)^{1/2}}{z_p^2 + r_0^2} + 1 \right] - \frac{z_p}{(z_p^2 + r_0^2)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + 2(1 + \epsilon)^{1/2}}{1 + \epsilon} + 1 \right] - \frac{1}{(1 + \epsilon)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{1 + \epsilon} + \frac{2}{(1 + \epsilon)^{1/2}} + 1 \right] - \frac{1}{(1 + \epsilon)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{(1 + \epsilon)^{1/2}} + 1 \right]^2 - \frac{1}{(1 + \epsilon)^{1/2}} \right), \end{aligned} \quad (\text{B.3})$$

where $\epsilon \equiv r_0^2/z_p^2$.

Now we should also take the Δt_{FS} in the limit as $z_e \rightarrow \infty$, a commonly made approximation assuming Earth is far from the source, which simplifies the equation. In that

case we have

$$(z_e^2 + r_0^2)^{1/2} = z_e + r_0^2/2z_e$$

so that

$$\begin{aligned} \Delta t_{\text{FS}} &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{z_p + (z_p^2 + r_0^2)^{1/2}}{-z_e + z_e + r_0^2/2z_e} \right] - \frac{1}{2} \left[\frac{z_p}{(z_p^2 + r_0^2)^{1/2}} + \frac{z_e}{z_e + r_0^2/2z_e} \right] \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[2z_e \frac{z_p + (z_p^2 + r_0^2)^{1/2}}{r_0^2} \right] - \frac{1}{2} \left[\frac{z_p}{(z_p^2 + r_0^2)^{1/2}} + 1 \right] \right). \end{aligned}$$

We can see that the logarithm diverges as $z_e \rightarrow \infty$. This is due to the nature of gravity's long range and occurs from the mathematical procedure of taking the limit (additional small time-delays accrue as one moves to infinity, still, an infinite sum of increasingly smaller time-delays may be infinite). Fortunately we can easily get around this technicality, since we are only sensitive to changing Shapiro delays, thus we can simply omit the constants from the above relation, which then reads

$$\begin{aligned} \Delta t_{\text{FS}} &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{z_p + (z_p^2 + r_0^2)^{1/2}}{r_0^2} \right] - \frac{z_p}{2(z_p^2 + r_0^2)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + (1 + r_0^2/z_p^2)^{1/2}}{r_0^2/z_p} \right] - \frac{1}{2(1 + r_0^2/z_p^2)^{1/2}} \right) \\ &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + (1 + \epsilon)^{1/2}}{z_p \epsilon} \right] - \frac{1}{2(1 + \epsilon)^{1/2}} \right). \end{aligned} \tag{B.4}$$

We can use these expressions to find the Δt_{NS} delay, which is

$$\begin{aligned}
 \Delta t_{\text{NS}} &= \Delta t_{\text{FS}} - \Delta t_{\text{diff}} \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + (1 + \epsilon)^{1/2}}{z_p \epsilon} \right] - \frac{1}{2(1 + \epsilon)^{1/2}} - \ln \left[\frac{1}{(1 + \epsilon)^{1/2}} + 1 \right]^2 + \frac{1}{(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1 + (1 + \epsilon)^{1/2}}{z_p \epsilon} \right] + \ln \left[\frac{1 + \epsilon}{2 + 2(1 + \epsilon)^{1/2} + \epsilon} \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{2z_p \epsilon} \frac{1 + (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon + \frac{1}{2}\epsilon + \epsilon(1 + \epsilon)^{1/2}}{1 + (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon} \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{2z_p \epsilon} \left(1 + \frac{\frac{1}{2}\epsilon + \epsilon(1 + \epsilon)^{1/2}}{1 + (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon} \right) \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{2z_p \epsilon} \left(1 + \frac{(\frac{1}{2}\epsilon + \epsilon(1 + \epsilon)^{1/2})(1 - (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon)}{(1 + (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon)(1 - (1 + \epsilon)^{1/2} + \frac{1}{2}\epsilon)} \right) \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{2z_p \epsilon} \left(1 + \frac{\frac{1}{2}\epsilon(-1 - \frac{3}{2}\epsilon + (1 + \epsilon)^{3/2})}{\frac{1}{4}\epsilon^2} \right) \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{1}{2z_p \epsilon} \left(1 + \frac{-2 - 3\epsilon + 2(1 + \epsilon)^{3/2}}{\epsilon} \right) \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right) \\
 &\approx \frac{2GM}{c^3} \left(\ln \left[\frac{(1 + \epsilon)^{3/2} - (1 + \epsilon)}{z_p \epsilon^2} \right] + \frac{1}{2(1 + \epsilon)^{1/2}} \right). \tag{B.5}
 \end{aligned}$$

The z_p and r_0 are functions of time and determined by the position of the emitter in relation to the central mass and the observer. We already found these in fact in (4.10) and (4.27) where we have

$$z_p = |z_{\text{obs}}| \quad \text{and} \quad r_0 = \sqrt{r_{\text{src}}^2 - z_p^2}. \tag{B.6}$$

We need to know when to use Δt_{NS} or Δt_{FS} ; the sign of z_{obs} indicates whether the emitter is on the near-side or far-side of the central mass. Also note that as $z_p \rightarrow 0$, $\epsilon \rightarrow \infty$, where both Δt_{NS} and Δt_{FS} tend to

$$-\frac{2GM}{c^3} \ln r_0.$$

Thus we can write

$$\Delta t_{\text{Shap}} = \begin{cases} \Delta t_{\text{FS}} & \text{for } z_{\text{obs}} \geq 0, \\ \Delta t_{\text{NS}} & \text{for } z_{\text{obs}} < 0. \end{cases} \quad (\text{B.7})$$

Finally, we define the ‘delayed time’ as the Shapiro delayed coordinate time,

$$t_{\text{Shap}}(t) \equiv t + \Delta t_{\text{Shap}}.$$

Recall from the discussion above (B.4) that we discarded an infinite constant from taking $z_e \rightarrow \infty$. Since the observer is not at an infinite distance, note that this Shapiro time is an accurate approximation up to an additive constant. The time dilated evolution of the binary itself is dealt with in the modelling of the gravitational redshift; the Shapiro delay accounts for the signal propagation time only and not the source evolution.

Bibliography

- [1] B. P. Abbott et al. “Tests of General Relativity with GW150914”. In: *Phys. Rev. Lett.* 116 (22 May 2016), p. 221101. DOI: [10.1103/PhysRevLett.116.221101](https://doi.org/10.1103/PhysRevLett.116.221101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.221101>.
- [2] R. A. Hulse and J. H. Taylor. “Discovery of a pulsar in a binary system.” In: 195 (Jan. 1975), pp. L51–L53. DOI: [10.1086/181708](https://doi.org/10.1086/181708).
- [3] R. Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run”. In: (Nov. 2021). arXiv: [2111.03606](https://arxiv.org/abs/2111.03606) [gr-qc].
- [4] John Baker et al. *The Laser Interferometer Space Antenna: Unveiling the Millihertz Gravitational Wave Sky*. 2019. DOI: [10.48550/ARXIV.1907.06482](https://doi.org/10.48550/ARXIV.1907.06482). URL: <https://arxiv.org/abs/1907.06482>.
- [5] Charles Hailey et al. “A density cusp of quiescent X-ray binaries in the central parsec of the Galaxy”. In: *Nature* 556 (Apr. 2018), pp. 70–73. DOI: [10.1038/nature25029](https://doi.org/10.1038/nature25029).
- [6] Albert Einstein. “Näherungsweise Integration der Feldgleichungen der Gravitation”. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)* (Jan. 1916), pp. 688–696.
- [7] Albert Einstein. “Über Gravitationswellen”. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* (Jan. 1918), pp. 154–167.
- [8] Jorge Cervantes-Cota, Salvador Galindo-Uribarri, and George Smoot. “A Brief History of Gravitational Waves”. In: *Universe* 2.3 (Sept. 2016), p. 22. ISSN: 2218-1997. DOI: [10.3390/universe2030022](https://doi.org/10.3390/universe2030022). URL: <http://dx.doi.org/10.3390/universe2030022>.
- [9] Robert M Wald. *General relativity*. Chicago, IL: Chicago Univ. Press, 1984.

- [10] Y. Fourès-Bruhat. “Théorème d’existence pour certains systèmes d’équations aux dérivées partielles non linéaires”. In: *Acta Mathematica* 88.none (1952), pp. 141–225. DOI: [10.1007/BF02392131](https://doi.org/10.1007/BF02392131). URL: <https://doi.org/10.1007/BF02392131>.
- [11] Hans Ringström. “Origins and development of the Cauchy problem in general relativity”. In: *Class. Quantum Grav.* 32.12 (2015). DOI: <https://doi.org/10.1088/0264-9381/32/12/124003>.
- [12] Yvonne Choquet-Bruhat and Robert Geroch. “Global aspects of the Cauchy problem in general relativity”. In: *Communications in Mathematical Physics* 14.4 (1969), pp. 329–335. DOI: [cmp/1103841822](https://doi.org/cmp/1103841822). URL: <https://doi.org/>.
- [13] Mihalis Dafermos. *Singularities and cosmic censorship in general relativity - Part 1*. DAMTP. 2013. URL: <http://web.math.princeton.edu/~dafermos/online-talks.html>.
- [14] F. A. E. Pirani. “On the Physical significance of the Riemann tensor”. In: *Acta Physica Polonica* 15 (Jan. 1956), pp. 389–405.
- [15] A. Einstein and N. Rosen. “On gravitational waves”. In: *Journal of the Franklin Institute* 223.1 (1937), pp. 43–54. ISSN: 0016-0032. DOI: [https://doi.org/10.1016/S0016-0032\(37\)90583-0](https://doi.org/10.1016/S0016-0032(37)90583-0). URL: <https://www.sciencedirect.com/science/article/pii/S0016003237905830>.
- [16] Emanuele Berti. “The First Sounds of Merging Black Holes”. In: *Physics* 9 (Feb. 2016). ISSN: 1943-2879. DOI: [10.1103/physics.9.17](https://doi.org/10.1103/physics.9.17). URL: <http://dx.doi.org/10.1103/Physics.9.17>.
- [17] Sean M. Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Cambridge University Press, 2019. DOI: [10.1017/9781108770385](https://doi.org/10.1017/9781108770385).
- [18] Demetrios Christodoulou. *The Formation of Black Holes in General Relativity*. 2008. arXiv: [0805.3880](https://arxiv.org/abs/0805.3880) [gr-qc].
- [19] S. Klainerman and I. Rodnianski. *On the formation of trapped surfaces*. 2009. arXiv: [0912.5097](https://arxiv.org/abs/0912.5097) [gr-qc].
- [20] Roger Penrose. “Gravitational Collapse and Space-Time Singularities”. In: *Physical Review Letters* 14.3 (1965), pp. 57–59. ISSN: 0031-9007. DOI: [10.1103/physrevlett.14.57](https://doi.org/10.1103/physrevlett.14.57).

- [21] Maximiliano Isi et al. “Testing the No-Hair Theorem with GW150914”. In: *Physical Review Letters* 123.11 (Sept. 2019). ISSN: 1079-7114. DOI: [10.1103/physrevlett.123.111102](https://doi.org/10.1103/physrevlett.123.111102). URL: <http://dx.doi.org/10.1103/PhysRevLett.123.111102>.
- [22] Éanna É Flanagan and Scott A Hughes. “The basics of gravitational wave theory”. In: *New Journal of Physics* 7 (Sept. 2005), pp. 204–204. ISSN: 1367-2630. DOI: [10.1088/1367-2630/7/1/204](https://doi.org/10.1088/1367-2630/7/1/204). URL: <http://dx.doi.org/10.1088/1367-2630/7/1/204>.
- [23] I. Bombaci et al. “Was GW190814 a Black Hole–Strange Quark Star System?”. In: *Physical Review Letters* 126.16 (Apr. 2021). ISSN: 1079-7114. DOI: [10.1103/physrevlett.126.162702](https://doi.org/10.1103/physrevlett.126.162702). URL: <http://dx.doi.org/10.1103/PhysRevLett.126.162702>.
- [24] The LIGO Scientific Collaboration et al. *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run*. 2021. arXiv: [2111.03606](https://arxiv.org/abs/2111.03606) [gr-qc].
- [25] Clifford M. Will. “On the unreasonable effectiveness of the post-Newtonian approximation in gravitational physics”. In: *Proceedings of the National Academy of Sciences* 108.15 (Mar. 2011), pp. 5938–5945. DOI: [10.1073/pnas.1103127108](https://doi.org/10.1073/pnas.1103127108). URL: <https://doi.org/10.1073/pnas.1103127108>.
- [26] Michele Maggiore. *Gravitational waves*. Oxford: Oxford Univ. Press, 2008. DOI: [10.1093/acprof:oso/9780198570745.001.0001](https://doi.org/10.1093/acprof:oso/9780198570745.001.0001). URL: <https://cds.cern.ch/record/1080850>.
- [27] Héctor Estellés et al. “Phenomenological time domain model for dominant quadrupole gravitational wave signal of coalescing binary black holes”. In: *Physical Review D* 103.12 (June 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.103.124060](https://doi.org/10.1103/physrevd.103.124060). URL: <http://dx.doi.org/10.1103/PhysRevD.103.124060>.
- [28] Matthew Pitkin. “Prospects of observing continuous gravitational waves from known pulsars”. In: *Mon. Not. Roy. Astron. Soc.* 415 (2011), pp. 1849–1863. DOI: [10.1111/j.1365-2966.2011.18818.x](https://doi.org/10.1111/j.1365-2966.2011.18818.x). arXiv: [1103.5867](https://arxiv.org/abs/1103.5867) [astro-ph.HE].
- [29] Paul D. Lasky. “Gravitational Waves from Neutron Stars: A Review”. In: *Publications of the Astronomical Society of Australia* 32 (2015), e034. DOI: [10.1017/pasa.2015.35](https://doi.org/10.1017/pasa.2015.35).

- [30] R. Abbott et al. “Search for Gravitational Waves Associated with Gamma-Ray Bursts Detected by Fermi and Swift during the LIGO–Virgo Run O3a”. In: *The Astrophysical Journal* 915.2 (July 2021), p. 86. DOI: [10.3847/1538-4357/abee15](https://doi.org/10.3847/1538-4357/abee15). URL: <https://doi.org/10.3847/1538-4357/abee15>.
- [31] Ernazar Abdikamalov, Giulia Pagliaroli, and David Radice. “Gravitational Waves from Core-Collapse Supernovae”. In: (2020). DOI: [10.48550/ARXIV.2010.04356](https://arxiv.org/abs/2010.04356). URL: <https://arxiv.org/abs/2010.04356>.
- [32] Chia-Feng Chang and Yanou Cui. “Gravitational Waves from Global Cosmic Strings and Cosmic Archaeology”. In: (2021). DOI: [10.48550/ARXIV.2106.09746](https://arxiv.org/abs/2106.09746). URL: <https://arxiv.org/abs/2106.09746>.
- [33] Sajal Mukherjee, Sanjit Mitra, and Sourav Chatterjee. “Gravitational wave observatories may be able to detect hyperbolic encounters of black holes”. In: *Monthly Notices of the Royal Astronomical Society* 508.4 (Sept. 2021), pp. 5064–5073. DOI: [10.1093/mnras/stab2721](https://doi.org/10.1093/mnras/stab2721). URL: <https://doi.org/10.1093/mnras/stab2721>.
- [34] Hengrui Xing et al. “Spinning black holes as cosmic string factories”. In: *Physical Review D* 103.8 (Apr. 2021). ISSN: 2470-0029. DOI: [10.1103/PhysRevD.103.083019](https://doi.org/10.1103/PhysRevD.103.083019). URL: <http://dx.doi.org/10.1103/PhysRevD.103.083019>.
- [35] Neil J Cornish and Tyson B Littenberg. “Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches”. In: *Classical and Quantum Gravity* 32.13 (June 2015), p. 135012. DOI: [10.1088/0264-9381/32/13/135012](https://doi.org/10.1088/0264-9381/32/13/135012). URL: <https://doi.org/10.1088/0264-9381/32/13/135012>.
- [36] Nelson Christensen. “Stochastic gravitational wave backgrounds”. In: *Reports on Progress in Physics* 82.1 (Nov. 2018), p. 016903. DOI: [10.1088/1361-6633/aae6b5](https://doi.org/10.1088/1361-6633/aae6b5). URL: <https://doi.org/10.1088/1361-6633/aae6b5>.
- [37] Maria Chiara Guzzetti et al. In: *La Rivista del Nuovo Cimento* 39.9 (Aug. 2016), pp. 399–495. ISSN: 0393697X, 0393697X. DOI: [10.1393/ncr/i2016-10127-1](https://doi.org/10.1393/ncr/i2016-10127-1). URL: <https://doi.org/10.1393/ncr/i2016-10127-1>.
- [38] David J. Weir. “Gravitational waves from a first-order electroweak phase transition: a brief review”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376.2114 (Jan. 2018), p. 20170126.

- DOI: [10.1098/rsta.2017.0126](https://doi.org/10.1098/rsta.2017.0126). URL: <https://doi.org/10.1098/rsta.2017.0126>.
- [39] Martti Raidal, Ville Vaskonen, and Hardi Veermäe. “Gravitational waves from primordial black hole mergers”. In: *Journal of Cosmology and Astroparticle Physics* 2017.09 (Sept. 2017), pp. 037–037. DOI: [10.1088/1475-7516/2017/09/037](https://doi.org/10.1088/1475-7516/2017/09/037). URL: <https://doi.org/10.1088/1475-7516/2017/09/037>.
- [40] Timothy Clifton et al. “Modified gravity and cosmology”. In: *Physics Reports* 513.1-3 (Mar. 2012), pp. 1–189. DOI: [10.1016/j.physrep.2012.01.001](https://doi.org/10.1016/j.physrep.2012.01.001). URL: <https://doi.org/10.1016/j.physrep.2012.01.001>.
- [41] Nicolás Yunes and Frans Pretorius. “Fundamental theoretical bias in gravitational wave astrophysics and the parametrized post-Einsteinian framework”. In: *Physical Review D* 80.12 (Dec. 2009). DOI: [10.1103/physrevd.80.122003](https://doi.org/10.1103/physrevd.80.122003). URL: <https://doi.org/10.1103/physrevd.80.122003>.
- [42] Kent Yagi, Nicolás Yunes, and Takahiro Tanaka. “Gravitational Waves from Quasicircular Black-Hole Binaries in Dynamical Chern-Simons Gravity”. In: *Physical Review Letters* 109.25 (Dec. 2012). DOI: [10.1103/physrevlett.109.251105](https://doi.org/10.1103/physrevlett.109.251105). URL: <https://doi.org/10.1103/physrevlett.109.251105>.
- [43] Luc Blanchet. “Gravitational Radiation from Post-Newtonian Sources and Inspiral Compact Binaries”. In: *Living Reviews in Relativity* 17.1 (Feb. 2014). DOI: [10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2). URL: <https://doi.org/10.12942/lrr-2014-2>.
- [44] J. Lange et al. “Parameter estimation method that directly compares gravitational wave observations to numerical relativity”. In: *Physical Review D* 96.10 (Nov. 2017). DOI: [10.1103/physrevd.96.104041](https://doi.org/10.1103/physrevd.96.104041). URL: <https://doi.org/10.1103/physrevd.96.104041>.
- [45] Yi Pan et al. “Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism”. In: *Physical Review D* 89.8 (Apr. 2014). DOI: [10.1103/physrevd.89.084006](https://doi.org/10.1103/physrevd.89.084006). URL: <https://doi.org/10.1103/physrevd.89.084006>.
- [46] A. Buonanno and T. Damour. “Effective one-body approach to general relativistic two-body dynamics”. In: *Physical Review D* 59.8 (Mar. 1999). DOI: [10.1103/physrevd.59.084006](https://doi.org/10.1103/physrevd.59.084006).

- physrevd . 59 . 084006. URL: <https://doi.org/10.1103/physrevd.59.084006>.
- [47] Mark Hannam et al. “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms”. In: *Physical Review Letters* 113.15 (Oct. 2014). DOI: [10.1103/physrevlett.113.151101](https://doi.org/10.1103/physrevlett.113.151101). URL: <https://doi.org/10.1103/physrevlett.113.151101>.
- [48] Laura Sberna et al. “Nonlinear effects in the black hole ringdown: Absorption-induced mode excitation”. In: *Physical Review D* 105.6 (Mar. 2022). DOI: [10.1103/physrevd.105.064046](https://doi.org/10.1103/physrevd.105.064046). URL: <https://doi.org/10.1103/physrevd.105.064046>.
- [49] B. P. Abbott et al. “Directly comparing GW150914 with numerical solutions of Einstein’s equations for binary black hole coalescence”. In: *Physical Review D* 94.6 (Sept. 2016). DOI: [10.1103/physrevd.94.064035](https://doi.org/10.1103/physrevd.94.064035). URL: <https://doi.org/10.1103/physrevd.94.064035>.
- [50] B P Abbott et al. “Effects of waveform model systematics on the interpretation of GW150914”. In: *Classical and Quantum Gravity* 34.10 (Apr. 2017), p. 104002. DOI: [10.1088/1361-6382/aa6854](https://doi.org/10.1088/1361-6382/aa6854). URL: <https://doi.org/10.1088/1361-6382/aa6854>.
- [51] I M Romero-Shaw et al. “Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue”. In: *Monthly Notices of the Royal Astronomical Society* 499.3 (Sept. 2020), pp. 3295–3319. ISSN: 1365-2966. DOI: [10.1093/mnras/staa2850](https://doi.org/10.1093/mnras/staa2850). URL: <http://dx.doi.org/10.1093/mnras/staa2850>.
- [52] Davide Gerosa et al. “A generalized precession parameter χ_p to interpret gravitational-wave data”. In: *Physical Review D* 103.6 (Mar. 2021). DOI: [10.1103/physrevd.103.064067](https://doi.org/10.1103/physrevd.103.064067). URL: <https://doi.org/10.1103/physrevd.103.064067>.
- [53] G Hobbs et al. “The International Pulsar Timing Array project: using pulsars as a gravitational wave detector”. In: *Classical and Quantum Gravity* 27.8 (Apr. 2010), p. 084013. DOI: [10.1088/0264-9381/27/8/084013](https://doi.org/10.1088/0264-9381/27/8/084013). URL: <https://doi.org/10.1088/0264-9381/27/8/084013>.

- [54] Jorge Cervantes-Cota, Salvador Galindo-Uribarri, and George Smoot. “A Brief History of Gravitational Waves”. In: *Universe* 2.3 (Sept. 2016), p. 22. DOI: [10.3390/universe2030022](https://doi.org/10.3390/universe2030022). URL: <https://doi.org/10.3390/universe2030022>.
- [55] Ajit Kembhavi and Pushpa Khare. “Future Gravitational Wave Detectors”. In: *Gravitational Waves: A New Window to the Universe*. Singapore: Springer Singapore, 2020, pp. 133–148. DOI: [10.1007/978-981-15-5709-5_9](https://doi.org/10.1007/978-981-15-5709-5_9). URL: https://doi.org/10.1007/978-981-15-5709-5_9.
- [56] J Aasi et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). URL: <https://doi.org/10.1088/0264-9381/32/7/074001>.
- [57] Peter R. Saulson. “Gravitational wave detection: Principles and practice”. In: *Comptes Rendus Physique* 14.4 (2013). Gravitational waves / Ondes gravitationnelles, pp. 288–305. ISSN: 1631-0705. DOI: <https://doi.org/10.1016/j.crhy.2013.01.007>. URL: <https://www.sciencedirect.com/science/article/pii/S163107051300008X>.
- [58] Massimo Tinto and Sanjeev V. Dhurandhar. “Time-delay interferometry”. In: *Living Rev. Rel.* 24.1 (2021), p. 1. DOI: [10.1007/s41114-020-00029-6](https://doi.org/10.1007/s41114-020-00029-6).
- [59] Jessica Page and Tyson B. Littenberg. “Bayesian time delay interferometry”. In: *Physical Review D* 104.8 (Oct. 2021). DOI: [10.1103/physrevd.104.084037](https://doi.org/10.1103/physrevd.104.084037). URL: <https://doi.org/10.1103/physrevd.104.084037>.
- [60] Massimo Tinto, Michele Vallisneri, and J. W. Armstrong. “Time-delay interferometric ranging for space-borne gravitational-wave detectors”. In: *Physical Review D* 71.4 (Feb. 2005). DOI: [10.1103/physrevd.71.041101](https://doi.org/10.1103/physrevd.71.041101). URL: <https://doi.org/10.1103/physrevd.71.041101>.
- [61] Bernard F Schutz. “Networks of gravitational wave detectors and three figures of merit”. In: *Classical and Quantum Gravity* 28.12 (May 2011), p. 125023. DOI: [10.1088/0264-9381/28/12/125023](https://doi.org/10.1088/0264-9381/28/12/125023). URL: <https://doi.org/10.1088/0264-9381/28/12/125023>.
- [62] Travis Robson, Neil J Cornish, and Chang Liu. “The construction and use of LISA sensitivity curves”. In: *Classical and Quantum Gravity* 36.10 (Apr. 2019),

- p. 105011. ISSN: 1361-6382. DOI: [10.1088/1361-6382/ab1101](https://doi.org/10.1088/1361-6382/ab1101). URL: <http://dx.doi.org/10.1088/1361-6382/ab1101>.
- [63] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. free software (GPL). 2018. DOI: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16).
- [64] Kent Yagi and Naoki Seto. “Detector configuration of DECIGO/BBO and identification of cosmological neutron-star binaries”. In: *Physical Review D* 83.4 (Feb. 2011). DOI: [10.1103/physrevd.83.044011](https://doi.org/10.1103/physrevd.83.044011). URL: <https://doi.org/10.1103/physrevd.83.044011>.
- [65] Stanislav Babak, Martin Hewitson, and Antoine Petiteau. *LISA Sensitivity and SNR Calculations*. 2021. DOI: [10.48550/ARXIV.2108.01167](https://arxiv.org/abs/2108.01167). URL: <https://arxiv.org/abs/2108.01167>.
- [66] D. V. Martynov et al. “Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy”. In: *Physical Review D* 93.11 (June 2016). DOI: [10.1103/physrevd.93.112004](https://doi.org/10.1103/physrevd.93.112004). URL: <https://doi.org/10.1103/physrevd.93.112004>.
- [67] Neil Cornish and Travis Robson. “Galactic binary science with the new LISA design”. In: *Journal of Physics: Conference Series* 840 (May 2017), p. 012024. DOI: [10.1088/1742-6596/840/1/012024](https://doi.org/10.1088/1742-6596/840/1/012024). URL: <https://doi.org/10.1088/1742-6596/840/1/012024>.
- [68] C J Moore, R H Cole, and C P L Berry. “Gravitational-wave sensitivity curves”. In: *Classical and Quantum Gravity* 32.1 (Dec. 2014), p. 015014. DOI: [10.1088/0264-9381/32/1/015014](https://doi.org/10.1088/0264-9381/32/1/015014). URL: <https://doi.org/10.1088/0264-9381/32/1/015014>.
- [69] J. Veitch and A. Vecchio. “Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network”. In: *Physical Review D* 81.6 (Mar. 2010). DOI: [10.1103/physrevd.81.062003](https://doi.org/10.1103/physrevd.81.062003). URL: <https://doi.org/10.1103/physrevd.81.062003>.
- [70] Fuzhen Zhang. *The Schur Complement and its Applications*. Vol. 4. Numerical Methods and Algorithms. New York: Springer, 2005. DOI: [10.1007/b105056](https://doi.org/10.1007/b105056).

- [71] George Labahn and Tamir Shalom. “Inversion of Toeplitz matrices with only two standard equations”. In: *Linear Algebra and its Applications* 175 (1992), pp. 143–158. ISSN: 0024-3795. DOI: [https://doi.org/10.1016/0024-3795\(92\)90306-U](https://doi.org/10.1016/0024-3795(92)90306-U). URL: <https://www.sciencedirect.com/science/article/pii/002437959290306U>.
- [72] William Trench. “An Algorithm for the Inversion of Finite Toeplitz Matrices”. In: *Siam Journal on Applied Mathematics - SIAMAM* 12 (Sept. 1964). DOI: [10.1137/0112045](https://doi.org/10.1137/0112045).
- [73] Daniel Foreman-Mackey et al. “emcee: The MCMC Hammer”. In: *Publications of the Astronomical Society of the Pacific* 125.925 (Mar. 2013), pp. 306–312. DOI: [10.1086/670067](https://doi.org/10.1086/670067). URL: <https://doi.org/10.1086/670067>.
- [74] G Ashton and C Talbot. “Bilby_MCMC: an MCMC sampler for gravitational-wave inference”. In: *Monthly Notices of the Royal Astronomical Society* 507.2 (Aug. 2021), pp. 2037–2051. DOI: [10.1093/mnras/stab2236](https://doi.org/10.1093/mnras/stab2236). URL: <https://doi.org/10.1093/mnras/stab2236>.
- [75] John Skilling. “Nested sampling for general Bayesian computation”. In: *Bayesian Analysis* 1.4 (2006), pp. 833–859. DOI: [10.1214/06-BA127](https://doi.org/10.1214/06-BA127). URL: <https://doi.org/10.1214/06-BA127>.
- [76] Michael J. Williams, John Veitch, and Chris Messenger. “Nested sampling with normalizing flows for gravitational-wave inference”. In: *Physical Review D* 103.10 (May 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.103.103006](https://doi.org/10.1103/physrevd.103.103006). URL: <http://dx.doi.org/10.1103/PhysRevD.103.103006>.
- [77] Walter Del Pozzo and John Veitch. *CPNest: Parallel nested sampling*. Astrophysics Source Code Library, record ascl:2205.021. May 2022. ascl: [2205.021](https://ascl.net/2205.021).
- [78] Joshua S. Speagle. “DYNesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences”. In: *mnras* 493.3 (Apr. 2020), pp. 3132–3158. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278). arXiv: [1904.02180](https://arxiv.org/abs/1904.02180) [astro-ph.IM].
- [79] Gregory Ashton et al. “Bilby: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. In: *The Astrophysical Journal Supplement Series* 241.2 (Apr. 2019), p. 27. ISSN: 1538-4365. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc). URL: <http://dx.doi.org/10.3847/1538-4365/ab06fc>.

- [80] John Baker et al. *The Laser Interferometer Space Antenna: Unveiling the Millihertz Gravitational Wave Sky*. 2019. arXiv: [1907.06482](https://arxiv.org/abs/1907.06482) [astro-ph.IM].
- [81] Scott A. Hughes. “Gravitational Waves from Merging Compact Binaries”. In: *Annual Review of Astronomy and Astrophysics* 47.1 (Sept. 2009), pp. 107–157. ISSN: 1545-4282. DOI: [10.1146/annurev-astro-082708-101711](https://doi.org/10.1146/annurev-astro-082708-101711). URL: <http://dx.doi.org/10.1146/annurev-astro-082708-101711>.
- [82] Stefan Grimm and Jan Harms. “Multiband gravitational-wave parameter estimation: A study of future detectors”. In: *Physical Review D* 102.2 (July 2020). ISSN: 2470-0029. DOI: [10.1103/physrevd.102.022007](https://doi.org/10.1103/physrevd.102.022007). URL: <http://dx.doi.org/10.1103/PhysRevD.102.022007>.
- [83] Chang Liu et al. “Multiband observation of LIGO/Virgo binary black hole mergers in the gravitational-wave transient catalog GWTC-1”. In: *Monthly Notices of the Royal Astronomical Society* 496.1 (June 2020), pp. 182–196. ISSN: 1365-2966. DOI: [10.1093/mnras/staa1512](https://doi.org/10.1093/mnras/staa1512). URL: <http://dx.doi.org/10.1093/mnras/staa1512>.
- [84] Scott E. Perkins, Nicolás Yunes, and Emanuele Berti. “Probing fundamental physics with gravitational waves: The next generation”. In: *Physical Review D* 103.4 (Feb. 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.103.044024](https://doi.org/10.1103/physrevd.103.044024). URL: <http://dx.doi.org/10.1103/PhysRevD.103.044024>.
- [85] Kaze W. K. Wong et al. “Expanding the LISA Horizon from the Ground”. In: *Physical Review Letters* 121.25 (Dec. 2018). ISSN: 1079-7114. DOI: [10.1103/physrevlett.121.251102](https://doi.org/10.1103/physrevlett.121.251102). URL: <http://dx.doi.org/10.1103/PhysRevLett.121.251102>.
- [86] Leor Barack and Curt Cutler. “Using LISA extreme-mass-ratio inspiral sources to test off-Kerr deviations in the geometry of massive black holes”. In: *Physical Review D* 75.4 (Feb. 2007). ISSN: 1550-2368. DOI: [10.1103/physrevd.75.042003](https://doi.org/10.1103/physrevd.75.042003). URL: <http://dx.doi.org/10.1103/PhysRevD.75.042003>.
- [87] Lorenzo Speri and Jonathan R. Gair. “Assessing the impact of transient orbital resonances”. In: *Physical Review D* 103.12 (June 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.103.124032](https://doi.org/10.1103/physrevd.103.124032). URL: <http://dx.doi.org/10.1103/PhysRevD.103.124032>.

- [88] Michele Vallisneri. “Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects”. In: *Physical Review D* 77.4 (Feb. 2008). ISSN: 1550-2368. DOI: [10.1103/physrevd.77.042001](https://doi.org/10.1103/physrevd.77.042001). URL: <http://dx.doi.org/10.1103/PhysRevD.77.042001>.
- [89] Carl L. Rodriguez et al. “Inadequacies of the Fisher information matrix in gravitational-wave parameter estimation”. In: *Physical Review D* 88.8 (Oct. 2013). ISSN: 1550-2368. DOI: [10.1103/physrevd.88.084013](https://doi.org/10.1103/physrevd.88.084013). URL: <http://dx.doi.org/10.1103/PhysRevD.88.084013>.
- [90] Rory Smith et al. “Fast and accurate inference on gravitational waves from precessing compact binaries”. In: *Physical Review D* 94.4 (Aug. 2016). ISSN: 2470-0029. DOI: [10.1103/physrevd.94.044031](https://doi.org/10.1103/physrevd.94.044031). URL: <http://dx.doi.org/10.1103/PhysRevD.94.044031>.
- [91] Barak Zackay, Liang Dai, and Tejaswi Venumadhav. *Relative Binning and Fast Likelihood Evaluation for Gravitational Wave Parameter Estimation*. 2018. arXiv: [1806.08792](https://arxiv.org/abs/1806.08792) [astro-ph.IM].
- [92] Soichiro Morisaki. “Accelerating parameter estimation of gravitational waves from compact binary coalescence using adaptive frequency resolutions”. In: *Physical Review D* 104.4 (Aug. 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.104.044062](https://doi.org/10.1103/physrevd.104.044062). URL: <http://dx.doi.org/10.1103/PhysRevD.104.044062>.
- [93] Serena Vinciguerra, John Veitch, and Ilya Mandel. “Accelerating gravitational wave parameter estimation with multi-band template interpolation”. In: *Classical and Quantum Gravity* 34.11 (May 2017), p. 115006. ISSN: 1361-6382. DOI: [10.1088/1361-6382/aa6d44](https://doi.org/10.1088/1361-6382/aa6d44). URL: <http://dx.doi.org/10.1088/1361-6382/aa6d44>.
- [94] G Larry Bretthorst. *Bayesian spectrum Analysis and parameter estimation*. Springer-Verlag Berlin Heidelberg, 1988. URL: <http://bayes.wustl.edu/glb/book.pdf>.
- [95] Michael L. Katz et al. “GPU-accelerated massive black hole binary parameter estimation with LISA”. In: *Physical Review D* 102.2 (July 2020). ISSN: 2470-0029. DOI: [10.1103/physrevd.102.023033](https://doi.org/10.1103/physrevd.102.023033). URL: <http://dx.doi.org/10.1103/PhysRevD.102.023033>.

- [96] André Klein and Guy Mélard. “Computation of the Fisher information matrix for time series models”. In: *Journal of Computational and Applied Mathematics* 64 (1995), pp. 57–68.
- [97] Alan Heavens. “Generalisations of Fisher Matrices”. In: *Entropy* 18.6 (June 2016), p. 236. ISSN: 1099-4300. DOI: [10.3390/e18060236](https://doi.org/10.3390/e18060236). URL: <http://dx.doi.org/10.3390/e18060236>.
- [98] Michele Vallisneri. *Beyond Fisher: exact sampling distributions of the maximum-likelihood estimator in gravitational-wave parameter estimation*. 2011. arXiv: [1108.1158 \[gr-qc\]](https://arxiv.org/abs/1108.1158).
- [99] Laura Sberna et al. “Observing GW190521-like binary black holes and their environment with LISA”. In: (2022). DOI: [10.48550/ARXIV.2205.08550](https://arxiv.org/abs/2205.08550). URL: <https://arxiv.org/abs/2205.08550>.
- [100] Nicola Tamanini et al. “Peculiar acceleration of stellar-origin black hole binaries: Measurement and biases with LISA”. In: *Physical Review D* 101.6 (Mar. 2020). DOI: [10.1103/physrevd.101.063002](https://doi.org/10.1103/physrevd.101.063002). URL: <https://doi.org/10.1103/physrevd.101.063002>.
- [101] Alejandro Torres-Orjuela et al. “Detecting the beaming effect of gravitational waves”. In: *Physical Review D* 100.6 (Sept. 2019). DOI: [10.1103/physrevd.100.063012](https://doi.org/10.1103/physrevd.100.063012). URL: <https://doi.org/10.1103/physrevd.100.063012>.
- [102] Oleg Bulashenko and Helena Ubach. *Lensing of gravitational waves: universal signatures in the beating pattern*. 2021. DOI: [10.48550/ARXIV.2112.10773](https://arxiv.org/abs/2112.10773). URL: <https://arxiv.org/abs/2112.10773>.
- [103] Bin Liu, Dong Lai, and Ye-Fei Yuan. “Merging compact binaries in hierarchical triple systems: Resonant excitation of binary eccentricity”. In: *Physical Review D* 92.12 (Dec. 2015). DOI: [10.1103/physrevd.92.124048](https://doi.org/10.1103/physrevd.92.124048). URL: <https://doi.org/10.1103/physrevd.92.124048>.
- [104] Barnabás Deme et al. “Detecting Kozai–Lidov Imprints on the Gravitational Waves of Intermediate-mass Black Holes in Galactic Nuclei”. In: *The Astrophysical Journal* 901.2 (Sept. 2020), p. 125. DOI: [10.3847/1538-4357/abafa3](https://doi.org/10.3847/1538-4357/abafa3). URL: <https://doi.org/10.3847/1538-4357/abafa3>.

- [105] Enrico Barausse, Vitor Cardoso, and Paolo Pani. “Can environmental effects spoil precision gravitational-wave astrophysics?” In: *Physical Review D* 89.10 (May 2014). DOI: [10.1103/physrevd.89.104059](https://doi.org/10.1103/physrevd.89.104059). URL: <https://doi.org/10.1103/physrevd.89.104059>.
- [106] Andrea Caputo et al. “Gravitational-wave Detection and Parameter Estimation for Accreting Black-hole Binaries and Their Electromagnetic Counterpart”. In: *The Astrophysical Journal* 892.2 (Apr. 2020), p. 90. DOI: [10.3847/1538-4357/ab7b66](https://doi.org/10.3847/1538-4357/ab7b66). URL: <https://doi.org/10.3847/1538-4357/ab7b66>.
- [107] Mark H Y Cheung et al. “Stellar-mass microlensing of gravitational waves”. In: *Monthly Notices of the Royal Astronomical Society* 503.3 (Feb. 2021), pp. 3326–3336. DOI: [10.1093/mnras/stab579](https://doi.org/10.1093/mnras/stab579). URL: <https://doi.org/10.1093/mnras/stab579>.
- [108] Eduardo Vitral and Gary A. Mamon. “Does NGC 6397 contain an intermediate-mass black hole or a more diffuse inner subcluster?” In: *Astronomy and Astrophysics* 646 (Feb. 2021), A63. DOI: [10.1051/0004-6361/202039650](https://doi.org/10.1051/0004-6361/202039650). URL: <https://doi.org/10.1051/0004-6361/202039650>.
- [109] Marco Gorghetto, Edward Hardy, and Horia Nicolaescu. “Observing invisible axions with gravitational waves”. In: *Journal of Cosmology and Astroparticle Physics* 2021.06 (June 2021), p. 034. DOI: [10.1088/1475-7516/2021/06/034](https://doi.org/10.1088/1475-7516/2021/06/034). URL: <https://doi.org/10.1088/1475-7516/2021/06/034>.
- [110] Scott E. Perkins, Nicolás Yunes, and Emanuele Berti. “Probing fundamental physics with gravitational waves: The next generation”. In: *Physical Review D* 103.4 (Feb. 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.103.044024](https://doi.org/10.1103/physrevd.103.044024). URL: <http://dx.doi.org/10.1103/PhysRevD.103.044024>.
- [111] Michael Zevin et al. “One Channel to Rule Them All? Constraining the Origins of Binary Black Holes Using Multiple Formation Pathways”. In: *The Astrophysical Journal* 910.2 (Apr. 2021), p. 152. DOI: [10.3847/1538-4357/abe40e](https://doi.org/10.3847/1538-4357/abe40e). URL: <https://doi.org/10.3847/1538-4357/abe40e>.
- [112] Volker Beckmann and Chris R. Shrader. *Active Galactic Nuclei*. 2012.

- [113] Marc S Seigar. *Spiral Structure in Galaxies*. 2053-2571. Morgan & Claypool Publishers, 2017. ISBN: 978-1-6817-4609-8. DOI: [10.1088/978-1-6817-4609-8](https://doi.org/10.1088/978-1-6817-4609-8). URL: <https://dx.doi.org/10.1088/978-1-6817-4609-8>.
- [114] Fabio Antonini, Mark Gieles, and Alessia Gualandris. “Black hole growth through hierarchical black hole mergers in dense star clusters: implications for gravitational wave detections”. In: *Monthly Notices of the Royal Astronomical Society* 486.4 (May 2019), pp. 5008–5021. DOI: [10.1093/mnras/stz1149](https://doi.org/10.1093/mnras/stz1149). URL: <https://doi.org/10.1093/mnras/stz1149>.
- [115] Michela Mapelli. “Binary Black Hole Mergers: Formation and Populations”. In: *Frontiers in Astronomy and Space Sciences* 7 (July 2020). DOI: [10.3389/fspas.2020.00038](https://doi.org/10.3389/fspas.2020.00038). URL: <https://doi.org/10.3389/fspas.2020.00038>.
- [116] Michael Zevin et al. “Implications of Eccentric Observations on Binary Black Hole Formation Channels”. In: *The Astrophysical Journal Letters* 921.2 (Nov. 2021), p. L43. DOI: [10.3847/2041-8213/ac32dc](https://doi.org/10.3847/2041-8213/ac32dc). URL: <https://doi.org/10.3847/2041-8213/ac32dc>.
- [117] R. Abbott et al. “Properties and Astrophysical Implications of the $150M_{\odot}$ Binary Black Hole Merger GW190521”. In: *The Astrophysical Journal* 900.1 (Sept. 2020), p. L13. DOI: [10.3847/2041-8213/aba493](https://doi.org/10.3847/2041-8213/aba493). URL: <https://doi.org/10.3847/2041-8213/aba493>.
- [118] Nadine Neumayer, Anil Seth, and Torsten Böker. “Nuclear star clusters”. In: *The Astronomy and Astrophysics Review* 28.1 (July 2020). DOI: [10.1007/s00159-020-00125-0](https://doi.org/10.1007/s00159-020-00125-0). URL: <https://doi.org/10.1007/s00159-020-00125-0>.
- [119] Howard D. Curtis. “Preface”. In: *Orbital Mechanics for Engineering Students (Third Edition)*. Ed. by Howard D. Curtis. Third Edition. Boston: Butterworth-Heinemann, 2014, pp. xi–xiii. ISBN: 978-0-08-097747-8. DOI: <https://doi.org/10.1016/B978-0-08-097747-8.05001-5>. URL: <https://www.sciencedirect.com/science/article/pii/B9780080977478050015>.
- [120] James Diebel. “Representing Attitude: Euler Angles, Unit Quaternions, and Rotation Vectors”. In: *Matrix* 58 (Jan. 2006).
- [121] Reza N. Jazar. *Advanced dynamics: rigid body, multibody, and aerospace applications*. eng. Hoboken, N.J: Wiley, 2011. ISBN: 1-5231-2339-7.

- [122] John M. A. Danby. *Fundamentals of celestial mechanics*. 1992.
- [123] Steven Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. New York: John Wiley and Sons, 1972. ISBN: 978-0-471-92567-5.
- [124] Irwin I. Shapiro. “Fourth Test of General Relativity”. In: *Phys. Rev. Lett.* 13 (26 Dec. 1964), pp. 789–791. DOI: [10.1103/PhysRevLett.13.789](https://doi.org/10.1103/PhysRevLett.13.789). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.789>.
- [125] Folkert Müller-Hoissen. “Gauss-Bonnet actions and their dimensionally reduced descendants”. In: *Fisica de Particularas e Campos* 94.8 (1989), pp. 116–124.
- [126] F. Mueller-Hoissen. “Gauss-Bonnet actions and their dimensionally reduced descendants”. In: 23 (Sept. 1989). DOI: [10.1103/PhysRevD.85.064022](https://doi.org/10.1103/PhysRevD.85.064022). URL: http://inis.iaea.org/search/search.aspx?orig_q=RN:23080054.
- [127] Ethan Dyer and Kurt Hinterbichler. “Boundary terms, variational principles, and higher derivative modified gravity”. In: *Physical Review D* 79.2 (Jan. 2009). DOI: [10.1103/physrevd.79.024028](https://doi.org/10.1103/physrevd.79.024028). URL: <https://doi.org/10.1103/physrevd.79.024028>.
- [128] Scott E. Perkins et al. “Improved gravitational-wave constraints on higher-order curvature theories of gravity”. In: *Phys. Rev. D* 104 (2 July 2021), p. 024060. DOI: [10.1103/PhysRevD.104.024060](https://doi.org/10.1103/PhysRevD.104.024060). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.024060>.
- [129] Edward Witten. “Quantum field theory and the Jones polynomial”. In: *Communications in Mathematical Physics* 121.3 (1989), pp. 351–399. DOI: [cmp/1104178138](https://doi.org/10.1007/BF01200506). URL: <https://doi.org/>.
- [130] J. E. Nelson. “2+1 dimensional gravity”. In: (2018). DOI: [10.48550/ARXIV.1804.08456](https://doi.org/10.48550/ARXIV.1804.08456). URL: <https://arxiv.org/abs/1804.08456>.
- [131] Edward Witten. “Three-Dimensional Gravity Revisited”. In: (2007). DOI: [10.48550/ARXIV.0706.3359](https://doi.org/10.48550/ARXIV.0706.3359). URL: <https://arxiv.org/abs/0706.3359>.
- [132] S. Deser, R. Jackiw, and S. Templeton. “Three-Dimensional Massive Gauge Theories”. In: *Phys. Rev. Lett.* 48 (15 Apr. 1982), pp. 975–978. DOI: [10.1103/PhysRevLett.48.975](https://doi.org/10.1103/PhysRevLett.48.975). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.48.975>.

- [133] R. Jackiw and S.-Y. Pi. “Chern-Simons modification of general relativity”. In: *Physical Review D* 68.10 (Nov. 2003). DOI: [10.1103/physrevd.68.104012](https://doi.org/10.1103/physrevd.68.104012). URL: <https://doi.org/10.1103/physrevd.68.104012>.
- [134] Stephon Alexander and Nicolás Yunes. “Chern–Simons modified general relativity”. In: *Physics Reports* 480.1-2 (Aug. 2009), pp. 1–55. DOI: [10.1016/j.physrep.2009.07.002](https://doi.org/10.1016/j.physrep.2009.07.002). URL: <https://doi.org/10.1016/j.physrep.2009.07.002>.
- [135] Nicolás Yunes and Frans Pretorius. “Dynamical Chern-Simons modified gravity: Spinning black holes in the slow-rotation approximation”. In: *Physical Review D* 79.8 (Apr. 2009). DOI: [10.1103/physrevd.79.084043](https://doi.org/10.1103/physrevd.79.084043). URL: <https://doi.org/10.1103/physrevd.79.084043>.
- [136] Kent Yagi et al. “Post-Newtonian, Quasi-Circular Binary Inspirals in Quadratic Modified Gravity”. In: *Phys. Rev. D* 85 (Mar. 2012). DOI: [10.1103/PhysRevD.85.064022](https://doi.org/10.1103/PhysRevD.85.064022).
- [137] Nicolás Yunes, Kent Yagi, and Frans Pretorius. “Theoretical physics implications of the binary black-hole mergers GW150914 and GW151226”. In: *Physical Review D* 94.8 (Oct. 2016). DOI: [10.1103/physrevd.94.084002](https://doi.org/10.1103/physrevd.94.084002). URL: <https://doi.org/10.1103/physrevd.94.084002>.
- [138] Frank W. J. Olver. *Asymptotics and special functions*. eng. Computer science and applied mathematics. New York: Academic Press, 1974. ISBN: 012525850X.
- [139] Serge Droz et al. “Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform”. In: *Physical Review D* 59.12 (May 1999). DOI: [10.1103/physrevd.59.124016](https://doi.org/10.1103/physrevd.59.124016). URL: <https://doi.org/10.1103/physrevd.59.124016>.
- [140] Jolien D. E. Creighton and Warren G. Anderson. *Gravitational-wave physics and astronomy: An introduction to theory, experiment and data analysis*. Sept. 2011. DOI: [10.1002/9783527636037](https://doi.org/10.1002/9783527636037).
- [141] Curt Cutler et al. “What can we learn from multi-band observations of black hole binaries?” In: 51.3, 109 (May 2019), p. 109. arXiv: [1903.04069](https://arxiv.org/abs/1903.04069) [[astro-ph.HE](https://arxiv.org/abs/1903.04069)].
- [142] Alberto Sesana. “Prospects for Multiband Gravitational-Wave Astronomy after GW150914”. In: *Physical Review Letters* 116.23 (June 2016). DOI: [10.1103/physrevlett.116.231102](https://doi.org/10.1103/physrevlett.116.231102). URL: <https://doi.org/10.1103/physrevlett.116.231102>.

- [143] Christopher J Moore, Davide Gerosa, and Antoine Klein. “Are stellar-mass black-hole binaries too quiet for LISA?” In: *Monthly Notices of the Royal Astronomical Society: Letters* 488.1 (July 2019), pp. L94–L98. DOI: [10.1093/mnrasl/slz104](https://doi.org/10.1093/mnrasl/slz104). URL: <https://doi.org/10.1093/mnrasl/slz104>.
- [144] Xian Chen and Pau Amaro-Seoane. “Revealing the Formation of Stellar-mass Black Hole Binaries: The Need for Deci-Hertz Gravitational-wave Observatories”. In: *The Astrophysical Journal* 842.1 (June 2017), p. L2. ISSN: 2041-8213. DOI: [10.3847/2041-8213/aa74ce](https://doi.org/10.3847/2041-8213/aa74ce). URL: <http://dx.doi.org/10.3847/2041-8213/aa74ce>.
- [145] Bin Liu and Dong Lai. “Probing the Spins of Supermassive Black Holes with Gravitational Waves from Surrounding Compact Binaries”. In: *The Astrophysical Journal* 924.2 (Jan. 2022), p. 127. DOI: [10.3847/1538-4357/ac3aef](https://doi.org/10.3847/1538-4357/ac3aef). URL: <https://doi.org/10.3847/1538-4357/ac3aef>.
- [146] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Gravitational Waves. OUP Oxford, 2008. ISBN: 9780198570745. URL: <https://books.google.co.uk/books?id=AqVpQgAACAAJ>.
- [147] R. Abbott et al. “Search for intermediate-mass black hole binaries in the third observing run of Advanced LIGO and Advanced Virgo”. In: *Astronomy & Astrophysics* 659 (Mar. 2022), A84. DOI: [10.1051/0004-6361/202141452](https://doi.org/10.1051/0004-6361/202141452). URL: <https://doi.org/10.1051/0004-6361/202141452>.
- [148] Muxin Liu et al. *Dark-siren Cosmology with Decihertz Gravitational-wave Detectors*. 2022. arXiv: [2205.06991](https://arxiv.org/abs/2205.06991) [[astro-ph.CO](https://arxiv.org/abs/2205.06991)].
- [149] Remya Nair et al. “Fundamental Physics Implications for Higher-Curvature Theories from Binary Black Hole Signals in the LIGO-Virgo Catalog GWTC-1”. In: *Physical Review Letters* 123.19 (Nov. 2019). DOI: [10.1103/physrevlett.123.191101](https://doi.org/10.1103/physrevlett.123.191101). URL: <https://doi.org/10.1103/physrevlett.123.191101>.
- [150] Davide Gerosa et al. “Multiband gravitational-wave event rates and stellar physics”. In: *Physical Review D* 99.10 (May 2019). DOI: [10.1103/physrevd.99.103004](https://doi.org/10.1103/physrevd.99.103004). URL: <https://doi.org/10.1103/physrevd.99.103004>.

- [151] Riccardo Buscicchio et al. “Bayesian parameter estimation of stellar-mass black-hole binaries with LISA”. In: *Physical Review D* 104.4 (Aug. 2021). DOI: [10.1103/physrevd.104.044065](https://doi.org/10.1103/physrevd.104.044065). URL: <https://doi.org/10.1103/physrevd.104.044065>.
- [152] Hsin-Yu Chen and Daniel E. Holz. *Finding the One: Identifying the Host Galaxies of Gravitational-Wave Sources*. 2016. DOI: [10.48550/ARXIV.1612.01471](https://arxiv.org/abs/1612.01471). URL: <https://arxiv.org/abs/1612.01471>.
- [153] Jeremy S. Heyl, Lars Hernquist, and David N. Spergel. “Inferring Galaxy Viewing Angles”. In: *Astrophysical Journal* 448 (July 1995), p. 64. DOI: [10.1086/175941](https://doi.org/10.1086/175941).
- [154] Xuheng Ding et al. “The Mass Relations between Supermassive Black Holes and Their Host Galaxies at $1 < z < 2$ with HST-WFC3”. In: *The Astrophysical Journal* 888.1 (Jan. 2020), p. 37. DOI: [10.3847/1538-4357/ab5b90](https://doi.org/10.3847/1538-4357/ab5b90). URL: <https://doi.org/10.3847/1538-4357/ab5b90>.