

**GRAVITATIONAL-WAVE
ASTRONOMY WITH
COALESCING COMPACT
BINARIES: DETECTION AND
PARAMETER ESTIMATION
WITH ADVANCED DETECTORS**

RORY JAMES EDWIN SMITH

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

DOCTOR OF PHILOSOPHY

**Astrophysics and Space Research
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
January 2013**

Acknowledgements

It is my sincere pleasure to thank the following people: My supervisors Ilya Mandel and Alberto Vecchio for their mentorship, support and patience; Chad Hanna for his mentorship, and for always keeping an open door despite his busy schedule; The members of the Bayesian and CBC groups for all their help; All my friends in the ASR group at Birmingham, especially Charlotte Bond, Mark Burke, Ludovico Carbone, Chris Collins, Paul Fulda, Kat Grover, Carl-Johan Haster, Konstantinos Kolokythas, Ilya Mandel, Chiara Mingarelli, Trevor Sidery and Will Vousden for years of fun and adventures; Andy Pearsall, Brigid Jones, Amelia Perry, Charlotte Kimbly, Stacey McGowan, Jim Phillips, Peter Rothery, Patrick Scott, Jim Smyth and Mat Turnbull for their friendship over the years; The Dog & Dolphin, refuge from real world problems; And my family, especially my parents, for their love and support.

Several of the chapters in this thesis were the result of collaborations and benefited from discussions with several people. Chapter 2 was based on work done in collaboration with Ilya Mandel and Alberto Vecchio, and benefited from discussions with Eliu Huerta. Chapter 3 was based on work done in collaboration with Kip Cannon, Chad Hanna, Drew Keppel and Ilya Mandel, and benefited from discussions with Tyson Littenberg. Chapter 4 was based on work done in collaboration with Chad Hanna, Ilya Mandel and Alberto Vecchio, and benefited from discussions with several people, in particular: Kipp Cannon, Scott Field, Drew Keppel, Vivien Raymond, Manuel Tiglio and Alan Weinstein. Finally, I am grateful for the many helpful comments on this thesis from Stephen Fairhurst and Andreas Freise.

My work has been supported by the STFC and the Perimeter Institute for Theoretical Physics.

Abstract

The current generation of interferometric gravitational-wave detectors, LIGO and Virgo, are undergoing upgrade to their so-called advanced phase. These instruments, together with new instruments in Japan and India, KAGRA and LIGO India, will form a network of advanced gravitational-wave detectors with which detections are expected to become routine. Amongst the prime sources for gravitational-wave astronomy are coalescing compact binaries consisting of neutron stars and/or black holes. Filtering detector data to detect these sources relies on precise templates of the expected gravitational-wave signals. In addition, estimating the parameters encoded in the signals (masses, spins etc...) requires sophisticated Bayesian inference techniques. Templates are typically computationally expensive to generate and can be a bottle-neck in data analysis.

Here we focus on two aspects of gravitational-wave astronomy using coalescing compact binaries. The first part of this thesis focuses on studying the requirements of template waveforms to detect intermediate-mass black holes through the coalescence of a stellar-mass companion into an intermediate-mass black hole in Advanced LIGO. The second part of this thesis focuses on numerical and analytic techniques to improve the efficiency of (Bayesian) parameter estimation on coalescing binaries when parameter estimation is dominated by template waveform generation. Such efficiency improvements to parameter estimation are crucial for gravitational-wave astronomy using advanced detectors.

Contents

Acknowledgements	i
Contents	iii
List of Figures	v
List of Tables	vii
1 INTRODUCTION	1
2 DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS	21
2.1 Introduction	21
2.2 Matched filtering of gravitational-wave signals	25
2.3 Waveforms	28
2.3.1 EOBNR	29
2.3.1.1 Inspiral-plunge waveform	29
2.3.1.2 Merger-ringdown waveform	30
2.3.2 HG waveforms	30
2.3.3 TaylorT4	32
2.4 SNR from inspiral, merger and ringdown	32
2.5 Effectiveness of inspiral-only templates for IMRAC searches	38
2.6 Comparison of inspiral-only waveforms	44
2.7 Discussion and conclusion	47

CONTENTS

3	TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS	49
3.1	Introduction	49
3.2	CBC parameter estimation	51
3.2.1	MCMC	54
3.3	Singular value decomposition of gravitational-wave template banks for CBC sources	56
3.4	Interpolating template waveforms using the SVD	59
3.5	Parameter estimation using interpolated waveforms	62
3.6	Results: Comparison of parameter estimates using interpolated and non-interpolated waveforms	67
3.7	Discussion and conclusion	71
4	PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS	73
4.1	Introduction	73
4.2	Interpolating inspiral-only time-domain waveforms	75
4.2.1	Parameter estimation using interpolated TaylorT4 waveforms: 5-dimensional case	78
4.2.1.1	Prior distributions	78
4.2.1.2	SVD set up	79
4.2.1.3	Results	79
4.3	Interpolating inspiral-merger-ringdown waveforms	81
4.4	Future work and Discussion	84
5	RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION	85
5.1	Introduction	85
5.1.1	Directly interpolating the likelihood function	88
5.1.2	Likelihood interpolation: Examples	89
5.1.2.1	Example 1: High-mass binary black holes	90
5.1.2.2	Example 2: Binary neutron stars	92
5.1.2.3	Practical considerations	92

5.1.2.4 Discussion and conclusion	94
6 SUMMARY	97
A FREQUENTLY USED QUANTITIES	101
A.1 Definitions of discrete representations of continuous functions . . .	101
A.2 Post-Newtonian phasing formulae	103
A.2.1 TaylorF2	103
A.2.2 TaylorT4	103
B SINGULAR VALUE DECOMPOSITION OF $M \times N$ MATRICES	105
B.1 Definitions of the singular value decomposition of $M \times N$ matrices	105
B.1.1 Step-by-step example	106
B.1.2 Geometric picture of the SVD	108
REFERENCES	111

CONTENTS

List of Figures

1.1	Inspiral phase of a gravitational wave from compact binary coalescence.	6
1.2	Anticipated Advanced LIGO noise sources.	8
1.3	Noise amplitude spectral densities of various Earth based gravitational-wave detectors.	9
2.1	Advanced LIGO high-power, zero-detuned noise PSD	33
2.2	RMS strain of IMRAC sources in Advanced LIGO.	35
2.3	SNR of IMRAC sources in Advanced LIGO as a function of total mass and mass ratio.	37
2.4	Effectiveness of inspiral-only templates to filter inspiral, merger and ringdown signals.	41
2.5	Effectiveness of inspiral-only EOBNR templates at filtering HG signals (left) and TaylorT4 signals (right) as a function of the source component masses encoded in the signal.	45
3.1	Coalescing compact binary system relative to a ground based gravitational-wave detector.	52
3.2	Initial LIGO noise power spectral density.	64
3.3	Patching of parameter space.	67
3.4	Marginalized one-dimensional posterior probability density functions produced using interpolated and non-interpolated template waveforms.	68

LIST OF FIGURES

4.1	Histogram of mismatches between (normalized) interpolated and non-interpolated TaylorT4 waveforms.	77
4.2	Marginalized one-dimensional posterior probability density functions produced using interpolated and non-interpolated TaylorT4 waveforms.	80
4.3	Histogram of mismatches between (normalized) interpolated and non-interpolated EOBNR waveforms.	83
5.1	Interpolated and non-interpolated log likelihoods (top), and percentage error (bottom) for a data set containing a gravitational-wave signal from the coalescence of binary black holes.	91
B.1	Action of a matrix on the unit circle.	108

List of Tables

2.1	Effectiveness of inspiral-only searches.	43
2.2	Effectiveness of <i>inspiral-only</i> EOBNR template waveforms in re-covering signals modelled using different waveform families.	46
3.1	Mean posterior parameter estimates (and standard deviations) of the marginalized PDFs using interpolated and non-interpolated SPA waveforms (Fig. 3.4).	68
3.2	Computational time of template waveform generation in units of computational time of interpolated waveforms.	70
4.1	Mean posterior value parameter estimates (and standard deviations) of the marginalized PDFs using interpolated and non-interpolated TaylorT4 waveforms (Fig. 4.2).	81

LIST OF TABLES

Chapter 1

INTRODUCTION

Gravitational-wave astronomy will open an entirely new observational window to the Universe with the direct detections of gravitational waves from high energy astrophysical phenomena. A network of ground-based interferometric gravitational-wave observatories (LIGO-Hanford and LIGO-Livingston (USA) and Virgo (Italy) [1]) are undergoing upgrades to their so-called “advanced” configurations which will be an order of magnitude more sensitive than their initial configurations and are scheduled to begin collecting data around 2015 [2]. In addition, two new ground-based interferometric detectors in Japan and India, KAGRA [3] and LIGO India [4] respectively, are expected to be online around 2020. This will usher in the “advanced detector era” in which gravitational-wave detections are expected to become routine [5].

The advanced detector network is expected to detect gravitational waves from a variety of fascinating astrophysical and cosmological events; from the late stages of the coalescence of compact binaries [5; 6; 7; 8] to a stochastic background of gravitational waves of astrophysical and cosmological origin [9; 10]. Of all the target sources, coalescing compact binaries are perhaps the the most likely candidates for detection [5]. Compact binary systems, consisting of neutron stars and/or black holes are believed to be in abundance throughout the Universe. Indeed, the first indirect evidence for the existence of gravitational radiation came from the observation of the binary pulsar system PSR B1913+16, known as the “Hulse-Taylor” binary pulsar after its discoverers, which is undergoing orbital decay through the emission of gravitational waves [11]. For their discovery, Hulse

1. INTRODUCTION

and Taylor were awarded the Nobel Prize in 1993. The Hulse-Taylor binary pulsar is undergoing orbital decay very slowly and will not merge for around three hundred million years. In contrast the advanced-detector network will directly observe the gravitational waves from the final minutes to seconds of the coalescence of compact binaries before they merge to form a single object. This makes compact binary coalescence (CBC) an ideal laboratory for probing dynamical, strong-field gravity.

With observations of gravitational waves from coalescing compact binaries we will be able to make precise measurements of their masses and spins as well as of their demographics [12; 13]. These measurements are essential for astrophysics/astronomy and fundamental physics using gravitational waves from compact binaries. Determination of the masses and spins will distinguish neutron stars and black holes and can be used to test the null-hypothesis that black holes are described by the Kerr metric [14; 15]. In addition, determination of source demographics, together with merger rates, places constraints on the formation mechanisms of binary systems [16]. While there are around ten known binary neutron star (NS-NS) systems [17] there are currently no known binary black hole (BH-BH) or black-hole-neutron-star (BH-NS) systems [18]. Despite this, BH-BH and BH-NS systems are believed to be fairly common throughout the Universe: population synthesis models suggest that realistic detection rates of BH-BH and BH-NS systems in Advanced LIGO/Virgo (aLIGO/Virgo) could be as high as 40 yr^{-1} [5]. Unsurprisingly, binary black holes are difficult to observe in the electromagnetic spectrum and gravitational waves are likely the *only* means by which they can be detected.

There are numerous open issues relating to black hole systems which gravitational -wave detections will be able to make important contributions to. In particular gravitational-wave detectors will be able to search for binary systems with total masses up to around $500 M_{\odot}$ [19] and will therefore be able to provide the first unambiguous observations of “intermediate-mass” black holes (IMBHs). Intermediate-mass is a blanket term which refers to black holes with masses (roughly) in the range $50 - 10^3 M_{\odot}$. Believed to form in globular clusters [e.g., 20; 21; 22], IMBHs are extremely difficult to observe in the electromagnetic spectrum and there are currently only tentative results suggesting their existence.

Detecting IMBHs will thus have important implications for globular cluster dynamics. In addition, detections of compact binary coalescence will be important in addressing the so called black hole “mass-gap” problem [23]: The stellar mass-distribution is continuous in the range $0.1 - 100 M_{\odot}$ and one would expect the same to be true of ensuing compact remnants. However, while neutron stars have been observed to have masses up to around $2 M_{\odot}$, the lightest black holes are observed to have masses around $5 M_{\odot}$, with no compact objects with masses in-between [24]. Gravitational-wave observations of BH-BH or BH-NS systems will therefore be able to shed light on whether there is a genuine gap between the heaviest neutron stars and the lightest black holes.

The central theme of the work presented here is the detection of, and parameter estimation on gravitational waves from CBC sources. Gravitational-wave astronomy using compact binaries will ultimately be limited by our ability to robustly detect and estimate the parameters of the sources. Here we focus on two challenges for detection and parameter estimation in aLIGO. Firstly, we study the potential to detect IMBHs using aLIGO through the coalescence of a stellar-mass companion into an IMBH. We refer to this kind of binary as an “intermediate mass-ratio coalescence” (IMRAC) source. Secondly, we address the computational efficiency of performing parameter estimation on gravitational waves from CBC sources. Parameter estimation on gravitational waves from CBC sources is currently highly computationally intensive [25], and the computational cost of performing parameter estimation on aLIGO data will be exacerbated if not directly addressed.

Detecting gravitational-wave signals from coalescing compact binaries and estimating the source parameters

Despite the scientific potential of gravitational-wave detectors, detection and parameter estimation are extremely difficult in practice. A passing gravitational wave creates a transverse strain in spacetime which we aim to detect using interferometers. An interferometric gravitational-wave detector in the path of such a gravitational-wave thus experiences a relative change in the length of its arms. This strain is the observable quantity upon which gravitational-wave astronomy

1. INTRODUCTION

relies. The exact form of the strain from a gravitational-wave from binary coalescence depends sensitively on the parameters of the binary, e.g. its masses and spins. However, in general a gravitational-wave signal from a coalescing binary has a characteristic signature. The coalescence of a binary system can be approximately split into three stages; *(i)* the so-called “inspiral” phase during which the components of the binary evolve along a quasi-circular orbit, “spiralling” into one another, *(ii)* a rapid merger where the two objects form a single object, and *(iii)* a quasi-normal ringing phase where the (merged) object behaves like a damped oscillator [26]. While it is possible to have inspiralling binaries on elliptical orbits, binaries observable over the bandwidth of aLIGO/Virgo are expected to have circularized, see e.g. [27]. These three stages of coalescence imprint a distinct signature to the gravitational-wave signal. Schematically it is useful to consider a gravitational-wave signal $h(t)$, as seen in a detector, to have the form

$$\begin{aligned}
 h(t) &= h^{\text{inspiral}}(t)\Theta(\tau - t) \\
 &+ h^{\text{merger-RD}}(t)\Theta(t - \tau),
 \end{aligned}
 \tag{1.1}$$

where $h^{\text{inspiral}}(t)$ is the part of the waveform from the inspiral phase of coalescence and $h^{\text{merger-RD}}(t)$ is the part of the waveform from the merger and ringdown (RD) phases of coalescence. Here $\Theta(\cdot)$ is the usual Heaviside step function. The time $t = 0$ is taken to be the time at a fiducial gravitational-wave frequency, f_{gw} . The time τ corresponds to the time from f_{gw} until the binary reaches a point close to the “inner-most stable circular orbit” (ISCO), after which the binary undergoes a plunge before the components merge [28]. Typically, τ is referred to as the “chirp time”. To leading order the chirp time is given by [29]

$$\tau = 2.18 \text{ s} \left(\frac{1.12 M_{\odot}}{M_c} \right)^{5/3} \left(\frac{100 \text{ Hz}}{f_{\text{gw}}} \right)^{8/3},
 \tag{1.2}$$

where $M_c = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ is known as the “chirp mass” of the binary. For a binary consisting of two neutron stars, each with mass $m_1 = m_2 = 1.4 M_{\odot}$, at 10 Hz, roughly the low frequency cut off of aLIGO, one has access to around $t \approx 20$ min of the gravitational-wave signal from the inspiral phase of coalescence.

At leading order, the gravitational-wave frequency at which the binary reaches the ISCO is

$$f_{\text{ISCO}} = 4.4 \text{ kHz} \left(\frac{M_{\odot}}{M} \right), \quad (1.3)$$

where M is the total mass of the binary. Thus a binary consisting of two neutron stars with total mass of $2.8 M_{\odot}$ will reach ISCO at around 1.5 kHz after which the components begin to merge.

The inspiral phase of coalescence

To leading order a coalescing binary is a mass-quadrupole radiator and the gravitational-wave strain from the inspiral phase (at leading order) is controlled by the second time-derivative of the binary's quadrupole moment, with a characteristic $1/r$ fall-off where r is the distance from the source to an observer. For a binary with chirp mass M_c , angular frequency Ω , and orbital phase $\phi(t)$, the inspiral phase of the strain signal of a face-on, overhead binary is to leading order given by [29]

$$\begin{aligned} h^{\text{inspiral}}(t) &= h_0^{\text{inspiral}}(t) \cos \phi(t), \\ h_0^{\text{inspiral}}(t) &\approx 10^{-23} \frac{1\text{Mpc}}{r} \left(\frac{M_c}{M_{\odot}} \right)^{5/3} \Omega^{2/3}, \end{aligned} \quad (1.4)$$

It is simple to estimate the typical strain from the inspiral of a compact binary source using Eq. (1.4). The gravitational-wave frequency is twice the orbital frequency, hence $\Omega = \pi f_{\text{gw}}$ where f_{gw} is the gravitational-wave frequency. Thus for a binary neutron star system with component masses $m_1 = m_2 = 1.4 M_{\odot}$ located at a fiducial distance of 15 Mpc (roughly the distance to the Virgo galaxy cluster) at a gravitational-wave frequency of 10 Hz, around the low frequency cut off of Advanced LIGO, the strain amplitude is approximately 10^{-23} . For an interferometric gravitational-wave detector with arms of length $l = 4 \text{ Km}$, the strain is $\delta l/l \approx 10^{-23}$ where δl is the change in the arm length. Thus the change in length of the arms which one aims to measure is approximately $\delta l \approx 10^{-20} m$.

As the binary evolves in time the frequency evolution is driven by the energy radiated through gravitational waves. The gravitational-wave frequency evolution

1. INTRODUCTION

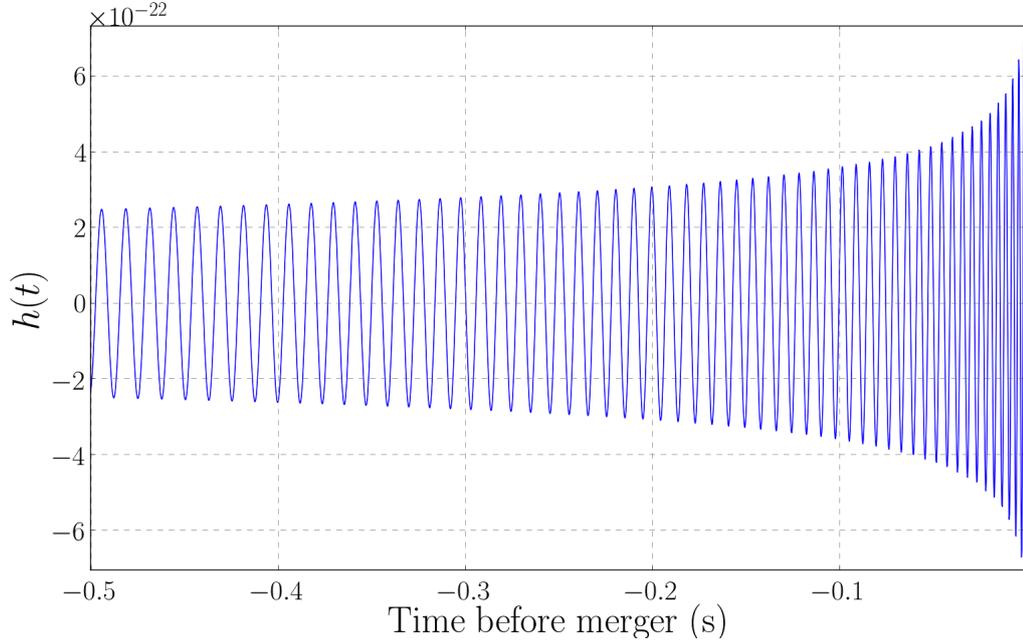


Figure 1.1: Inspiral phase of a gravitational-wave signal emitted from the coalescence of a face-on, overhead binary black hole system. The component masses of the binary are $m_1 = m_2 = 5 M_\odot$. The waveform’s amplitude and frequency increase up to the merger. This waveform was generated using the LIGO Scientific Collaboration Algorithm Library’s (LAL) implementation of the TaylorT4 waveform approximant [26].

(to leading order) scales like [29]

$$\dot{f}_{\text{gw}} \propto \left(\frac{M_c}{M_\odot} \right)^{5/3} f_{\text{gw}}^{11/3}. \quad (1.5)$$

As the frequency of the gravitational-wave signal increases, both $h_0^{\text{inspiral}}(t)$ and \dot{f}_{gw} increase. This endows the inspiral phase of the coalescence signal with a characteristic “chirp” signature, i.e., a signal which increases in both amplitude and frequency. An example of the inspiral phase of a gravitational-wave signal from binary coalescence is shown in Fig. 1.1.

The merger and ringdown phases of coalescence

After the inspiral phase the components of the binary merge and undergo quasi-normal ringing, with the gravitational-wave strain amplitude decreasing exponentially over a short time scale. The characteristic time scale is given by [30]

$$\mathcal{T} = \frac{2}{\pi} \frac{1}{f_0}, \quad (1.6)$$

$$f_0 = 10 \text{ kHz} \frac{M_\odot}{M}. \quad (1.7)$$

To leading order the ringdown waveform (for a binary with non-spinning components) has the form [30]

$$h^{\text{ringdown}}(t) \approx 10^{-20} \frac{1 \text{ Mpc}}{r} \frac{M}{M_\odot} e^{-t/\mathcal{T}} \cos(2\pi f_0 t). \quad (1.8)$$

For a double neutron star binary with component masses $m_1 = m_2 = 1.4 M_\odot$ the amplitude decreases by an e-fold after time $t \sim 5 \times 10^{-5}$ s.

Detecting the strain signal from a coalescing binary

Detecting gravitational waves thus requires measuring a signal whose amplitude is physically very small. Noise sources in the detectors typically cause strains in the instrument at the frequency of the gravitational-wave strain signal, which makes the task of detecting the strain from a gravitational wave challenging. The dominant noise sources for aLIGO include thermal (or Brownian) noise, quantum noise from the uncertainty in the photon count in the lasers (also called shot noise) and seismic noise from the Earth's seismic activity [2]. The strain spectrum (strain/ $\sqrt{\text{Hz}}$) of the various noise sources of aLIGO is shown in Fig. 1.2. Each noise source typically acts over a certain frequency range which is illustrated in Fig. 1.2. For example, thermal noise in the suspensions and mirrors is present over a frequency range between $\mathcal{O}(10 \text{ Hz})$ and $\mathcal{O}(1 \text{ kHz})$. Conversely, seismic noise typically dominates the noise below frequencies of $\mathcal{O}(10 \text{ Hz})$ which effectively sets the low frequency cut off of the instruments. The combined noise (shown as the black curve) is known as the detector noise amplitude spectral density

1. INTRODUCTION

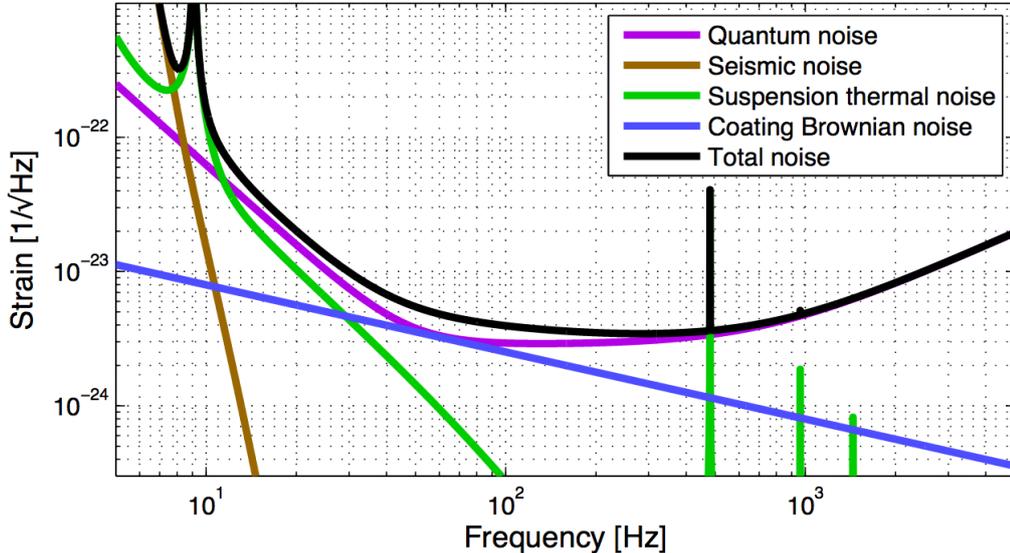


Figure 1.2: Anticipated Advanced LIGO noise sources generated using the online GWINC tool [31]. The combined noise (shown as the black curve) is known as the detector noise amplitude spectral density (ASD) and defines the overall sensitivity of the instrument. The noise ASD corresponds to the instrument in its high-powered zero-detuned configuration.

(ASD) and defines the overall sensitivity of the instrument. The peak strain-sensitivity of aLIGO can be read off from Fig. 1.2 and is around $10^{-23} \text{ Hz}^{-1/2}$ over a frequency range $10 \text{ Hz} \lesssim f \lesssim 10^3 \text{ Hz}$. In Fig. 1.3 we show the noise amplitude spectral densities for various Earth based gravitational-wave detectors. From Fig. 1.3 one can see the improvement in sensitivity between LIGO/Virgo, and their advanced configurations aLIGO/Virgo, which are around an order of magnitude more sensitive.

The problem of detecting signals in noise is a classic of signal analysis. The noise properties are well understood and over time scales around the typical length of a gravitational-wave signal it is effectively stationary [1; 29]. In addition, for many practical purposes it is often assumed that the noise will not be correlated with a gravitational-wave signal, for example, by considering it to be Gaussian. However, in practice real noise is not Gaussian and a large effort is spent on characterizing the non-Gaussian features of the detector noise (see e.g., [33; 34; 35]). Secondly, the functional form of gravitational-wave signals from

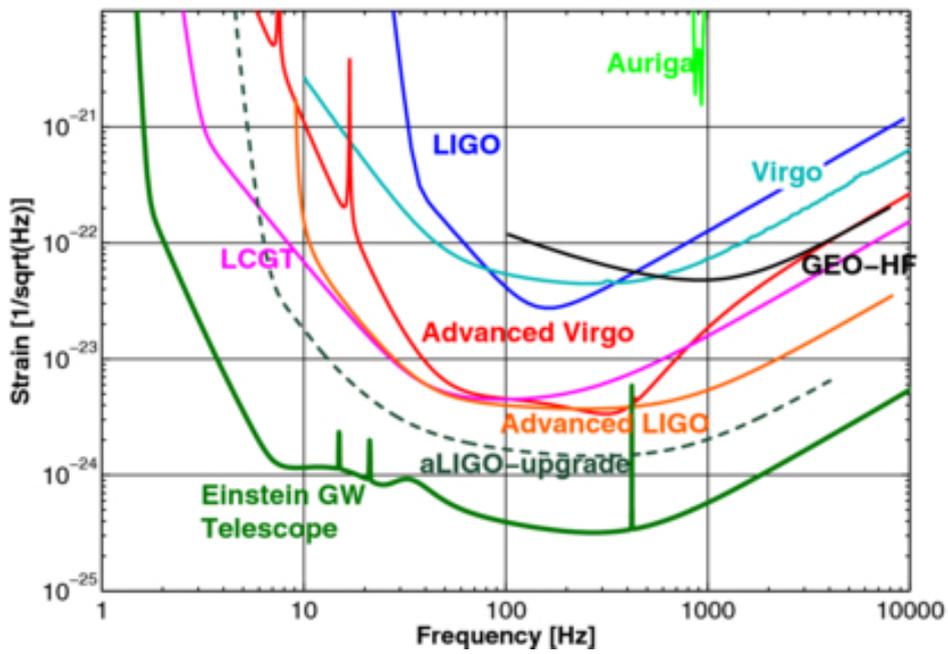


Figure 1.3: Noise amplitude spectral densities of various Earth based gravitational-wave detectors [32]. Note the increase in sensitivity by around an order of magnitude between LIGO and Virgo, and Advanced LIGO and Advanced Virgo.

1. INTRODUCTION

CBC sources can be modelled accurately [26], and one can filter the data for a signal of known functional form. This is known as “matched filtering” [5; 6; 7; 8] and is effective for detecting gravitational waves from CBC sources precisely because the gravitational-wave signal can be accurately modelled. In the presence of a gravitational-wave signal $h(t)$, the data $d(t)$ has the form

$$d(t) = n(t) + h(t), \quad (1.9)$$

where $n(t)$ is the noise. To detect gravitational waves using matched filtering one compares the data with a set of theoretical gravitational waveforms $\{T(t)\}$ (known as templates), normalized by the square of the detector’s ASD, which act as filters of the data. The filter which maximizes the detection-statistic is then chosen to be the candidate signal. A commonly used detection statistic is the signal-to-noise ratio (SNR) which is a ratio of the power of the gravitational-wave signal to that of the noise. The SNR of the data set d in Eq. (1.9) is [29]

$$\text{SNR} = 4\Re \frac{1}{\sigma} \int_0^\infty df \tilde{T}^*(f) \tilde{d}(f) / S_n(f), \quad (1.10)$$

where $\tilde{T}(f)$ is the Fourier-transform of the template $T(t)$ (and similarly for the data) and $S_n(f)$ is the detector power spectral density (PSD), which is the square of the ASD. The normalization σ is given by $\sigma = \left[\int_0^\infty df \tilde{T}^*(f) \tilde{T}(f) / S_n(f) \right]^{1/2}$. The optimal SNR is achieved when the template waveform is same as the signal in the data set and hence one requires the templates to be a faithful description of real gravitational waveforms. Using templates which are not accurate descriptions of gravitational waves contained in the data degrades the SNR, as the optimal SNR is recovered only when the template corresponds exactly to the signal. Using inaccurate templates also leads to losses in detection rates which can be understood as follows. The gravitational-wave signal from binary coalescence has a $1/r$ fall-off, c.f. Eq. (1.4), where r is the distance from the Earth to the source. Using a template which degrades the SNR by a fixed fractional amount ϵ can be compensated for by increasing the overall amplitude of the template. This is achieved by reducing the distance to the source by the same amount as the SNR loss. Reducing the distance to the source by a factor of ϵ

thus reduces the observable volume to a factor of ϵ^3 of its original size. Assuming an isotropic distribution of sources, the fractional loss in detection rates should thus scale like $(1 - \epsilon^3)$

One of the key considerations for a gravitational-wave search is the templates one uses to filter the data. For example, searches which target so called “low mass” systems with total masses less than $25 M_\odot$ [7] typically only require waveforms which describe the inspiral phase of coalescence. This can be understood by considering the frequency at which the binaries merge; a system with total mass of $5 M_\odot$ will merge at around 1 kHz, c.f. Eq. (2.1). For initial LIGO this inspiral phase sits in the frequency band over which the detector is most sensitive and hence the inspiral phase can dominate the SNR, c.f. Fig. 1.3. Conversely, heavier systems are liable to merge at frequencies at which the detector is most sensitive and so searches for these systems need to employ waveforms which also describe the merger and ringdown. The so-called “high mass” searches in initial LIGO [8] which aimed to detect compact binaries with masses up to $100 M_\odot$ employed such waveforms. Systems with even higher masses may have a very short detectable inspiral phase, or may have even merged at the point at which they are detectable. For such systems, the ringdown can be the dominant source of the SNR and hence one may require waveforms which describe solely the ringdown. The “ringdown search” during initial LIGO [30] employed such ringdown-only waveforms to search for binaries with total masses up to $390 M_\odot$.

Extracting the binary’s parameters

The gravitational waves encode the parameters of the source and for a binary system with non-spinning components there are nine such parameters: the two component masses, the distance to the Earth, the two angles which specify the sky position, the binary’s inclination, the “polarization phase”, the time of coalescence and the orbital phase at coalescence. Spinning binaries also have the two spin-vectors of each component which increases the number of parameters to fifteen. Because the filters are parameterized, the filter which maximizes the SNR can be used to estimate the parameters of the source. This is the maximum likelihood estimator of the source parameters [29] and in this fashion one can extract

1. INTRODUCTION

information about the source of a candidate gravitational-wave signal.

A limitation of these kinds of parameter estimates is that one only gets a “point estimate” of the parameters of the gravitational-wave source, the maximum likelihood estimator. A more informative approach would also provide a quantitative picture of the uncertainty of parameter estimates over the full domain of the parameter space. For this one subjects a stretch of data to a full Bayesian analysis and computes the posterior probability distributions of the source parameters of a candidate gravitational-wave signal, see e.g. [12; 36; 37]. This approach allows one to assign meaningful probability distributions to the measured properties of the candidate source. The Bayesian approach also allows for “model selection”, by comparing the null hypothesis that the data contains only noise, to the hypothesis that the data also contains a signal. This is the ratio of the probability that a data set is described by just noise, to the probability that the data contains a gravitational-wave signal. This is commonly known as the “Bayes factor” [12] and provides a useful way of quantifying the relative probability that the data contains a gravitational-wave signal.

Challenges of waveform modelling for detection and parameter estimation

There are a number of theoretical challenges associated with gravitational-wave detection and (Bayesian) parameter estimation. Firstly, in order to robustly detect and make parameter estimates of gravitational-wave sources, one requires templates which are accurate descriptions of real gravitational waves. While there are several schemes in which gravitational waveforms can be modelled, such as the post-Newtonian (PN) perturbative expansion and numerical solutions to the Einstein field equations (numerical relativity) [26], there is not yet a complete analytic description of a gravitational waveform for binary systems with arbitrary mass ratios and spins. For example, the PN expansion is valid when the binary system has components with comparable masses and the expansion is itself only formally valid up to a point before the binary’s components merge, when the expansion parameter, v/c becomes large. For orbital velocities $v/c = (M\pi f)^{1/3} \gtrsim 0.2$ the PN energy flux deviates significantly from more accurate numerical expressions,

see [e.g., 38; 39]. Hence waveforms computed using the PN expansion only describe the “inspiral” of the coalescence, ignoring the merger and the subsequent quasi-normal ringing of the merged objects. In contrast, numerical relativity is an approach which solves the Einstein field equations numerically and can describe the full coalescence. However, numerical relativity is extremely computationally expensive, with a single simulation taking between several hundred, to thousands of CPU-hours [40].

The limitations of both approaches owe to the difficulty of the two-body problem in general relativity. Because of this there has been a strong effort to produce waveforms which are a hybrid of analytic and numerical approaches, stitching together analytic fits and numerical expressions to produce an approximation scheme which can describe the full coalescence. An example of such a scheme which is commonly employed for gravitational-wave detection is the effective one-body numerical relativity (EOBNR) [26; 41] approach. Despite its success, it is only known to be accurate for binaries with mass ratios $q \equiv m_2/m_1 \geq 1/6$ though it is also formally valid in the test-particle limit $q \rightarrow 0$. For binaries with extreme mass-ratios ($q \ll 1$), consisting of a stellar-mass black hole and a super-massive black hole, several authors ([e.g., 42; 43]) have modelled the two body motion as that of a point particle in Kerr/Schwarzschild space-time, with corrections to account for the gravitational radiation-reaction from solutions to the Teukolsky equation. Like the PN approach, this also only describes the inspiral phase of the coalescence. More recently, the authors in [44] have added high-order-in-mass-ratio corrections to extreme mass-ratio inspiral waveforms so that they describe binaries with so called “intermediate mass-ratios”. For the work here we consider an intermediate mass-ratio q to be in the range $1/200 \leq q \leq 1/10$. However, to date only one numerical simulation of an intermediate mass-ratio binary exists [45] and analytic approximations to describe intermediate mass-ratio inspirals are at an early stage of development. This renders the problem of robustly detecting sources which consist of an intermediate-mass black hole and a small stellar mass companion difficult. For matched filtering to be successful, the template must be able to track the phasing of the gravitational wave in the data to within a wave-cycle, over many wave cycles. Errors in the phasing of the filters can render them ineffective and can thus lead to losses in detection rates.

1. INTRODUCTION

Template waveform generation can be a computational burden to performing matched-filtered searches or Bayesian parameter estimation. In order to perform a matched-filtered search for gravitational waves, or a full Bayesian parameter estimation, one requires a set of template waveforms which have to be computed numerically, known as a “template bank”. Templates generated using the PN or EOBNR approach are computed by first solving the underlying two-body equations of motion from which the gravitational waves are extracted. This is somewhat undesirable for gravitational-wave data analysis because of the sheer number of template waveforms required to construct a template bank for detection. In order to effectively search for gravitational waves, one typically constructs a dense template bank on the *intrinsic* parameter-space (the two masses and spins which control the phase evolution) to ensure that one does not lose a significant amount of SNR through under-dense gridding [46]. The typical number of template waveforms required for detection can be on the order of hundreds of thousands [47]. In addition, the accuracy requirements of Bayesian parameter estimation are somewhat stricter than for detection, and a standard analysis requires computing of millions of template waveforms [25]. Whereas in a gravitational-wave search one performs a maximization procedure to find the maximum-likelihood estimator, in Bayesian parameter estimation one is interested in computing the full posterior probability density function (PDF) of the source parameters. The PDF $p(\vec{\theta}|d, \mathcal{H})$ of a set of parameters $\vec{\theta}$ which parameterize a model, \mathcal{H} , assumed to describe a data set d , is related to the likelihood $p(d|\vec{\theta}, \mathcal{H})$, the *a priori* probability distribution of the source parameters $p(\vec{\theta}|\mathcal{H})$ and the “evidence” $p(d|\mathcal{H})$ via Bayes’ theorem [48]:

$$p(\vec{\theta}|d, \mathcal{H}) = \frac{p(\vec{\theta}|\mathcal{H}) p(d|\vec{\theta}, \mathcal{H})}{p(d|\mathcal{H})}. \quad (1.11)$$

As discussed earlier, the parameter vector $\vec{\theta}$ for coalescing binaries is nine dimensional for binaries with non-spinning components and fifteen dimensional for binaries with spinning components. It is immediately clear that a fine gridding in each dimension is not feasible and so one must make use of algorithms to preferentially sample regions of the parameter space where the bulk of the posterior probability is contained. There are currently several stochastic sampling algorithms which are dedicated to parameter estimation on gravitational waves

from CBC sources. These include Markov-chain Monte Carlo [13] and nested sampling [12] algorithms. A typical stochastic Bayesian analysis can take up to several weeks to run on a single stretch of data, around ten seconds in duration. The dominant cost of parameter estimation is waveform generation, and the cost of the analysis scales linearly with the cost of waveform generation.

Methods to improve the efficiency of detection and parameter estimation are therefore crucial for gravitational-wave astronomy. On the side of detection it has recently been shown that there is a large redundancy in the number of filters used for gravitational-wave searches [49] and the authors employed a singular value decomposition (SVD) of a template bank of waveforms used for a gravitational-wave search to show that the total number can be reduced by around an order of magnitude without incurring significant losses in SNR. In addition, the decomposition of the (discrete) template bank can be used to generate a continuous description of the bank (effectively an interpolation technique) [50]. The interpolation bypasses the need to perform ODE solving to generate template waveforms and could offer improvements to the efficiency of detection and parameter estimation when the computational cost is dominated by template waveform generation. It is all the more crucial to address the issue of the computational cost of template generation in the run-up to the advanced detector era as the in-band signal duration will be longer than in the initial configurations of the instruments. This will render computing templates more computationally expensive if not directly dealt with. For example, a binary neutron star system with total mass of $2.8 M_{\odot}$ at a gravitational-wave frequency of 40 Hz, roughly the low frequency cut off of initial LIGO, would have an in band signal duration of around 25 s (c.f. Eq. (1.2)). Conversely, the same system at a gravitational-wave frequency of 10 Hz, roughly the low frequency cut off of aLIGO, has an in band signal duration of around 20 min. Computing hundreds of thousands, or millions of these waveforms will render detection and parameter estimation an even more computationally expensive procedure.

1. INTRODUCTION

Overview of the thesis

Chapter II

Chapter II focuses on the detection of intermediate-mass black holes in aLIGO. As a detection channel we consider the coalescence of a stellar-mass black hole into an intermediate-mass black hole (IMBH) which we refer to as an intermediate mass-ratio coalescence (IMRAC). We consider binaries in which the IMBH has mass in the range $24 M_{\odot} - 200 M_{\odot}$ with a stellar-mass companion having masses in the range $1.4 M_{\odot} - 18.5 M_{\odot}$, and with mass ratios $1/140 \leq q \leq 1/10$. This range is well within the detectable mass-range of aLIGO. We investigate the relative contribution to the signal-to-noise ratio (SNR) of the three different phases of the coalescence - inspiral, merger and ringdown - of IMRACs as would be seen by aLIGO. We find that merger and ringdown play the dominant role over the bulk of the mass space though there is a small portion in which the inspiral portion of the coalescence signal dominates. We identify three regions in the mass-space in which (i) inspiral-only searches could be performed with losses in detection rates, L , in the range $10\% \lesssim L \lesssim 27\%$, (ii) searches based on inspiral-only templates lead to a loss in detection rates in the range $27\% \lesssim L \lesssim 50\%$, and (iii) templates that include merger and ringdown are essential to prevent losses in detection rates greater than 50%. In addition, we provide a comparison of existing inspiral-only waveforms over the mass range considered in the study. We consider waveforms based on the EOBNR and PN approximation schemes, and a waveform family developed specifically to describe intermediate mass-ratio inspirals [44] which we refer to as the “Huerta-Gair” waveform family after its authors. We measure the consistency of waveforms in different families by computing their normalized inner-products, maximized over the parameters which control the phase evolution of the waveforms. We find that the different families are $\geq 90\%$ consistent with each other which reinforces the importance of several numerical relativity simulations in this mass-range with which semi-analytical models of waveforms could be constructed in order to produce robust templates for detecting IMRACs. This study is restricted to binaries on (quasi-) circular orbits with non-spinning components.

The text of Chapter II is based on a paper by *Rory Smith*, Ilya Mandel and

Alberto Vecchio [19]. My contribution to this work was: (i) the implementation of the Huerta-Gair waveform family, (ii) implementation of all the calculations (iii) organized and led the project (iv) wrote the paper. The paper benefited from discussions with Eliu Huerta. This paper has arXiv number arXiv:1302.6049 [astro-ph.HE].

Chapter III

Chapter III is the first of three chapters on techniques to improve the efficiency of Bayesian parameter estimation. In Chapter III we provide a proof of principle of a technique which could be used to improve the efficiency of Bayesian parameter estimation on CBC sources when the computational cost of the analysis is dominated by template waveform generation. The work employs a singular value decomposition (SVD) of a template bank of waveforms which we use to interpolate templates for parameter estimation. The interpolated templates have already been employed in gravitational-wave searches [49; 50]. Interpolating template waveforms can be a less intensive computational procedure than generating templates via differential equation solving, for example. Because the accuracy requirements of gravitational-wave searches are typically less stringent than for parameter estimation we (a) study the bias in posterior parameter estimates incurred through using interpolated template waveforms, and (b) compute the relative computational cost of computing interpolated templates and non-interpolated templates. We find that the bias in mean posterior parameter estimates is negligible on the scale of statistical measurement uncertainty and that interpolating templates is around an order of magnitude faster than generating standard time-domain waveforms, e.g., TaylorT4 and EOBNR [26] waveforms. Our study uses a simple class of interpolated inspiral-only template waveforms and is restricted to a small patch of the parameter space. We consider future work required in order for interpolated template waveforms to be integrated within parameter estimation pipelines.

Chapter III grew out of a collaboration between *Rory Smith*, Kipp Cannon, Chad Hanna, Drew Keppel and Ilya Mandel while on a fellowship at the Perimeter Institute for Theoretical Physics. The text of Chapter III is based on a paper by

1. INTRODUCTION

Rory Smith, Kipp Cannon, Chad Hanna, Drew Keppel and Ilya Mandel [51]. My contribution to this work was: (i) writing the SVD-interpolation software, (ii) integrating the SVD-interpolation software into the LSC Algorithm Library (LAL), (iii) the comparison of PDFs generated using interpolated and non-interpolated waveforms in Sec. 3.5, (iv) the comparison of computational time of waveform families in Sec. 3.6, (v) organized and led the project, (vi) wrote the paper. This paper has arXiv number arXiv:1211.1254 [astro-ph.HE].

Chapter IV

Chapter IV builds on the previous chapter and represents continuing work on improving the efficiency of Bayesian parameter estimation. Here we extend the interpolation used in Chapter III to more general classes of template waveforms; inspiral-only time-domain waveforms, and time-domain waveforms which describe the full coalescence signal. We find that our implementation of the interpolation is successful for inspiral-only time-domain waveforms but fails for time-domain waveforms which describe the full coalescence.

All the work in Chapter IV is my own. It grew out of the work in Chapter III which was a collaboration with Kipp Cannon, Chad Hanna, Drew Keppel and Ilya Mandel, and benefited from discussions with the afore mentioned people.

Chapter V

Chapter V builds on the previous two chapters on improving the efficiency of Bayesian parameter estimation. While the previous two chapters focussed on interpolating template waveforms, here we show how to efficiently compute the continuous likelihood function on the three-dimensional subspace of parameters on which it has a non-trivial dependence – the chirp mass, symmetric mass ratio and coalescence time – via interpolation. We show that subsequently, sampling this interpolated likelihood function is a significantly cheaper computational process than directly evaluating the likelihood; we report improvements in computational time of between two and three orders of magnitude while keeping likelihoods accurate to $\lesssim 0.025\%$. Generating the interpolant of the likelihood function over a significant portion of the CBC mass space is computationally expensive but

highly parallelizable. This suggests that the overall wall time of generating the interpolant of the likelihood function can in principle be very small compared to the time of a full parameter estimation analysis.

The work in Chapter V is based on a paper by *Rory Smith*, Chad Hanna, Ilya Mandel and Alberto Vecchio. This work benefited from discussions with several people, in particular Kipp Cannon, Scott Field, Drew Keppel, Vivien Raymond, Manuel Tiglio and Alan Weinstein. This paper has arXiv number arXiv:1305.3798 [astro-ph.HE].

1. INTRODUCTION

Chapter 2

DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

2.1 Introduction

Observations of ultra-luminous X-ray sources and simulations of globular cluster dynamics suggest the existence of intermediate-mass black holes (IMBHs) [5; 20; 52; 53; 54; 55]. However, observational evidence for their existence is still under debate, see e.g. [56; 57]. Gravitational waves from binary coalescences involving IMBHs with masses $50 M_{\odot} \lesssim M \lesssim 500 M_{\odot}$ are potentially detectable by advanced detectors – including Advanced LIGO [58], Advanced Virgo [59], and KAGRA [60] – with a low frequency cutoff of around 10 Hz. If IMBHs do exist, one likely contribution to gravitational-wave detections is believed to be through the coalescence of a compact stellar-mass companion (black hole or neutron star) with an IMBH, at a possible rate of up to $\sim 10 \text{ yr}^{-1}$ [5; 14; 61], with masses of the IMBH in the range $50 M_{\odot} \lesssim M \lesssim 350 M_{\odot}$ [14]. We will denote these signals

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

as intermediate mass-ratio coalescences (IMRACs)¹.

Given that IMBHs in this mass range have proved extremely difficult to observe in the electromagnetic spectrum, gravitational-wave detections may provide the first unambiguous observations of such objects. Such observations would form an important channel for probing the dynamical history of globular clusters. Furthermore, Advanced LIGO/Virgo (aLIGO/Virgo) may be able to provide measurements of the quadrupole moment of a black hole [14; 15], which would allow a null-hypothesis test of the Kerr metric for IMBHs.

The gravitational waveform from the coalescence of two compact objects can be divided into three phases: a gradual inspiral, a rapid merger, and the quasi-normal ringdown of the resulting black hole. The relative contribution to the expected coalescence signal from inspiral, merger and ringdown is an important consideration for gravitational-wave searches. To leading Newtonian order the gravitational wave frequency at the inner-most stable circular orbit (ISCO) is

$$f_{\text{ISCO}} = 4.4 \text{ kHz} \left(\frac{M_{\odot}}{M} \right), \quad (2.1)$$

where M is the total mass of the binary. For advanced detectors with a low frequency cut-off of ~ 10 Hz, we may only have access to either the very late stages of the inspiral, or solely merger and ringdown for the heaviest IMRAC systems. While the power in the merger and ringdown is suppressed by a factor of the mass ratio relative to the power in the inspiral [62], the fact that IMRAC systems are liable to merge either in-band, or at the low frequency limit of the bandwidth, means that merger and ringdown may be significant over a large portion of the detectable mass-space. Additionally, for cases where IMRAC waveforms are inspiral-dominated, it is useful to know where inspiral-only searches could be targeted.

Detecting IMRACs through gravitational waves will require template gravitational-waveform families adapted to highly asymmetrical mass-ratio systems. However the development of numerical relativity simulations and perturbative techniques

¹ In the literature, the term frequently used for this class of objects is *intermediate mass-ratio inspirals* or IMRIs, see e.g. [14; 61]. However, in the context of ground-based observations, in particular with second-generation instruments, we will show that the full coalescence is important for these systems, and it therefore seems to be more appropriate to call them IMRAC.

in this regime is at an early stage which is potentially problematic. The issue of appropriate template waveform families is thus central to the detection of IMRACs through gravitational waves.

The effective-one-body approach, calibrated to numerical relativity, has led to template waveforms, known as EOBNR [41], that describe the full inspiral, merger and ringdown coalescence-signal for comparable mass-ratio binaries; EOBNR waveforms should also be accurate at extreme mass ratios. However, to date only one full numerical simulation exists for mass-ratio $q = 1/100$ binaries [45]. EOBNR waveforms have not yet been compared to NR simulations at such mass ratios, so their validity in the IMRAC regime remains to be demonstrated.

Meanwhile, in the context of extreme mass-ratio binaries, several authors have modelled the two-body motion by computing radiative and conservative self-force corrections to Kerr geodesic motion [42; 43]. Waveforms computed within this scheme are inspiral-only and are only developed to lowest order in the mass ratio. These waveforms have been adapted to describe intermediate mass-ratio inspirals by including higher-order-in-mass-ratio corrections in [63] and have been used to study the detection of intermediate mass-ratio inspirals in the Einstein Telescope [44]. We refer to these intermediate mass-ratio inspiral waveforms as the “Huerta Gair” (HG) waveform family after its authors. This waveform family should be physically well motivated to describe the inspiral of IMRACs.

Typically one does not have an exact representation of “true” gravitational-wave signals but requires templates which are sufficiently effective at filtering such signals. A common metric for quantifying how well approximate waveform families are at filtering gravitational-wave signals is known the “effectiveness”, or fitting factor [26]. This measures the fraction of the theoretical maximum signal-to-noise ratio (SNR) that could be recovered by using non-exact template waveforms.

The work in this Chapter focuses on the requirements of template waveforms for the detection of IMRACs in aLIGO. Firstly, by computing the effectiveness of inspiral-only template waveforms at filtering the full coalescence signal, we determine the relative importance of the inspiral and merger-ringdown phases. We identify three regions in the m_1 - m_2 plane in which: (a) inspiral-only searches are feasible with losses in detection rates L in the range $10\% \lesssim L \lesssim 27\%$, (b)

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

searches are limited by the lack of merger and ringdown in template waveforms and are liable to incur losses in detection rates in the range $27\% \lesssim L \lesssim 50\%$, and (c) where merger and ringdown are essential for searches in order to prevent losses in detection rates greater than 50%.

Secondly, to gain insight into the accuracy of the inspiral portion of IMRAC waveforms we compute the effectiveness of the inspiral-only portion of EOBNR waveforms at filtering gravitational-wave signals as described by the HG waveform family. We find that there is a non-negligible discrepancy between EOBNR and HG inspirals in the regime where inspiral-only searches could be considered sufficient. For reference we also compare EOBNR inspirals to a post-Newtonian (PN) [64; 65] waveform family known as TaylorT4 [26]. The PN expansion is liable to be a poor choice of approximant for IMRACs because of the large number of cycles spent at small radii. We find that EOBNR and HG are in better agreement with each other than to TaylorT4, as might be expected from the previous observation.

Our approach does not directly address the accuracy of template waveforms, because none of the waveforms considered have been matched to full numerical waveforms. However, assuming that the waveform families we consider “bracket” the correct gravitational waveforms in the IMRAC regime, this approach provides a useful heuristic for quantifying the effectiveness of existing gravitational waveforms for IMRAC searches. Further numerical relativity simulations will be important in the continuing development of accurate template waveforms for IMRACs.

Our analysis improves upon previous work to determine the detectability of IMRAC sources [66] which only considered the faithfulness of template waveforms, i.e., the effectiveness of template waveforms evaluated at the signal parameters. Additionally, that study only considered inspiral-only waveforms and focused on the LISA regime.

This chapter is organized as follows. In Sec. 2.2 we review the basics of matched filtering as applied to gravitational-wave detection. In Sec. 2.3 we describe the waveform families used in our study. In Sec. 2.4 we compute the contributions to signal-to-noise ratio (SNR) from the inspiral and merger and ringdown phases of EOBNR waveforms in the intermediate mass-ratio regime. In

Sec. 2.5 we study the effectiveness of inspiral-only waveforms to filter inspiral-merger-ringdown gravitational wave signals from IMRAC sources and identify the three regions in which different searches could be conducted. In Sec. 2.6 we compare the inspiral portion of EOBNR waveforms to HG and TaylorT4 waveforms. In Sec 2.7 we consider the implications of our results for future searches in advanced detectors.

2.2 Matched filtering of gravitational-wave signals

In order to quantify the requirements of gravitational waveforms for IMRAC searches, we summarize the key concepts of matched filtering which is the core principle behind gravitational-wave detection. The fundamental object is the optimal filter of the data which is used to define the maximum signal-to-noise ratio achievable from a stretch of data which contains a gravitational-wave signal.

The output of a gravitational-wave detector which contains a gravitational-wave signal will generally be of the form $s(t) = h(t) + n(t)$. Here $s(t)$ is a data-stream consisting of a gravitational-wave signal $h(t)$ and noise $n(t)$. We take the noise to be Gaussian and stationary for simplicity. We note that in general the noise will not be Gaussian and stationary in the detector network, though will not consider the effects of non-Gaussianity in the noise here.

One can filter the data for a signal $h(t)$ if we know the form of $h(t)$. Consider the following filtering procedure. We take a signal containing a rapidly oscillating function of time, $h(t)$, to be the signal buried in Gaussian and stationary noise. The convolution of the detector output, $s(t)$, with the signal, averaged over an observation time T , is

$$\frac{1}{T} \int_0^T dt s(t)h(t) = \frac{1}{T} \int_0^T dt h^2(t) + \frac{1}{T} \int_0^T dt n(t)h(t). \quad (2.2)$$

Both h and n are separately oscillating functions of time. The first term on the right hand side of Eq. (2.2) is positive definite and grows, for large values of T , as T . Averaged over the observation time T , its value will be the average amplitude

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

of the signal, $1/T \int_0^T dt h^2(t) \sim h_0^2$, where h_0 is the average amplitude of $h(t)$. Conversely, the noise $n(t)$ and the signal $h(t)$ are uncorrelated. The second term on the right hand side of Eq. (2.2) grows at most, for large T , as $T^{1/2}$ so that $1/T \int_0^T dt n(t)h(t) \sim (\tau/T)^{1/2} n_0 h_0$, where n_0 is the characteristic amplitude of the noise and τ is its characteristic period of the noise. Hence in the limit $T \rightarrow \infty$, the second term on the right hand side of Eq. (2.2) averages to zero and the signal is filtered out of the data-stream.

More formally, we work with the filtered detector output \hat{s} given by

$$\hat{s} = \int_{-\infty}^{\infty} dt s(t)K(t), \quad (2.3)$$

where $s(t)$ is a data stream and $K(t)$ is the filter function. A commonly used detection statistic is the ‘‘signal-to-noise’’ ratio (SNR), defined as the ratio of the expected value of \hat{s} when a signal is present, S , to the rms value of \hat{s} when the signal is absent, N . We thus seek the filter $K(t)$ which maximises the SNR. The expectation value of the filtered signal, S , is [29]

$$\begin{aligned} S &= \int dt \langle s(t) \rangle K(t), \\ &= \int dt h(t) K(t), \\ &= \int df \tilde{h}^*(f) \tilde{K}(f), \end{aligned} \quad (2.4)$$

where we have used the fact that $\langle n(t) \rangle = 0$ on the second line. The Fourier transform, $\tilde{h}(f)$ of $h(t)$, and its inverse are defined in Eq. (A.1). The rms value of \hat{s} when the signal is absent, N can be computed from [29]

$$\begin{aligned} N^2 &= \left[\langle s^2(t) \rangle - \langle s(t) \rangle^2 \right]_{h=0}, \\ &= \langle s^2(t) \rangle_{h=0}, \\ &= \int df \frac{1}{2} S_n(f) |\tilde{K}(f)|^2, \end{aligned} \quad (2.5)$$

where the quantity in the $S_n(f)$ is known as the detector ‘‘power spectral density’’

(PSD), and is defined by the ensemble average of the noise [67]:

$$\frac{1}{2}S_n(f) = \langle |\tilde{n}(f)|^2 \rangle \frac{1}{T}, \quad (2.6)$$

where T is the duration of an observation time over which the noise would typically be collected in an experiment. The SNR is thus given by [67]

$$\frac{S}{N} = \frac{\int_{-\infty}^{\infty} df \tilde{h}(f)\tilde{K}^*(f)}{\left[\int_{-\infty}^{\infty} df (1/2)S_n(f)|\tilde{K}(f)|^2 \right]^{1/2}}. \quad (2.7)$$

To find the filter K which maximizes Eq. (2.7) we define the “noise-weighted inner-product”:

$$(a|b) = 4 \Re \int_0^{\infty} df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}, \quad (2.8)$$

in terms of which the SNR, Eq. (2.7), is given by

$$\frac{S}{N} = \frac{(u|h)}{(u|u)^{1/2}}, \quad (2.9)$$

where $u(t)$ is a function with Fourier transform $\tilde{u}(f) = (1/2) S_n(f)\tilde{K}(f)$. Geometrically, we have a vector of unit norm $\hat{u} = u/(u|u)^{1/2}$ with which we are taking the inner-product with h . Thus, the SNR is maximized when \hat{u} is parallel to h . The optimal filter, K , which maximizes the SNR is therefore given by

$$\tilde{K}(f) = \text{const.} \frac{\tilde{h}(f)}{S_n(f)}. \quad (2.10)$$

The constant is irrelevant and the best filter is thus given by the signal itself. This is a crucial point as it states that we must know the form of the signal which we want to detect in order to obtain the optimal SNR. In terms of the noise-weighted inner-product the optimal SNR is expressed succinctly as

$$\left(\frac{S}{N}\right)_{\max} = (h|h)^{1/2}. \quad (2.11)$$

Central to the task of detecting gravitational waves from IMRACs is the choice

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

of model waveform h which we require to filter the data. In the section below we outline three model waveform families which we use in our study.

2.3 Waveforms

In this section we summarise the key concepts entering the construction of the waveforms used in this study. Throughout the chapter, for a binary system with individual component masses m_1 and m_2 (with $m_2 < m_1$) we define the total mass as $M \equiv m_1 + m_2$, and mass ratio and symmetric mass ratio as $q \equiv m_2/m_1$ and $\eta \equiv m_1 m_2 / (m_1 + m_2)^2$, respectively.

We first consider the family of waveforms constructed by calibrating the effective-one-body approach to numerical relativity (EOBNR) [41; 68; 69]. The EOBNR family describes the full inspiral-merger-ringdown signal; it is currently used in searches that reach the IMBH mass range, so far up to $100 M_\odot$ [8]. The free coefficients in the family have been fitted to comparable mass ratio numerical relativity simulations, and by construction this family is deemed to be faithful in the test particle limit. For this work, we use the implementation provided by the LIGO Scientific Collaboration Algorithm Library (LAL) that corresponds to the approximant EOBNRv2 [41].

We also consider a waveform family based on test particle motion in Kerr / Schwarzschild space-time with radiative and conservative self-force corrections which we refer to as the Huerta-Gair (HG) family [44]. The approximation scheme is constructed specifically to handle highly-asymmetrical mass-ratio binaries and is therefore a physically well motivated approximation scheme for intermediate mass-ratio inspirals. These waveforms have been compared against Teukolsky based waveforms and the match exceeds 95% over a large portion of the parameter space [63]. These waveforms have been used to study detection of intermediate mass-ratio inspirals in Einstein telescope in [44]. The Huerta-Gair waveforms describe only the inspiral portion of the coalescence signal. There is no corresponding LAL approximant for the HG waveform family.

Finally, as a reference we use the standard inspiral-only post-Newtonian approximation corresponding to the LAL approximant TaylorT4 at three-and-a-half post-Newtonian order in phase [26].

The TaylorT4 waveforms used here are computed in the so-called “restricted” amplitude approximation, which assumes the waveform amplitude to be zeroth post-Newtonian order and only includes the leading second harmonic of the orbital phase. We only include the leading second harmonic of the orbital phase in the EOBNR waveforms. We do not consider the effects of spin or eccentricity in any of the waveform families, as we restrict to circular orbits and non-spinning black holes.

2.3.1 EOBNR

EOBNR uses an effective one body (EOB) Hamiltonian to evolve the system’s trajectory through the inspiral and plunge phase, while numerical simulations and black hole perturbation theory are used to construct the merger and ringdown phases respectively. The PN Hamiltonian generates a waveform $h^{insp-plunge}(t)$ which is matched at some time, t_{match} , to a merger and ringdown waveform, $h^{merger-RD}$. The full EOBNR waveform is given by (see [26] and references therein),

$$\begin{aligned}
 h(t) &= h^{insp-plunge}(t)\Theta(t_{match} - t) \\
 &+ h^{merger-RD}(t)\Theta(t - t_{match}),
 \end{aligned}
 \tag{2.12}$$

where $\Theta(\cdot)$ is the Heaviside step function.

2.3.1.1 Inspiral-plunge waveform

To calculate the inspiral trajectory one solves Hamilton’s equations given the Hamiltonian of the EOB metric and the inspiral-plunge waveform at leading Newtonian-order is [26],

$$h(t)^{insp-plunge} = \mathcal{A} \hat{\omega}^{1/3} \sin [2\phi(t)],
 \tag{2.13}$$

where the “angular frequency” $\hat{\omega}$ is given by Eq. (3.23b) in [26]. The overall scaling \mathcal{A} depends on the source location with respect to the detector. In this

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

chapter we will use normalized template waveforms and so we leave this factor unspecified.

2.3.1.2 Merger-ringdown waveform

The merger-ringdown waveform is made up of a superposition of quasi-normal modes (QNMs) [41; 70; 71]. We only consider the dominant (2, 2) mode and the merger-ringdown waveform is given by

$$h^{\text{merger-RD}}(t) = \mathcal{A} \sum_{n=0}^{N-1} A_{22n} e^{-i\sigma_{22n}(t-t_{\text{match}}^{22})}, \quad (2.14)$$

where n is the overtone number of the Kerr QNM, N is the number of overtones, t_{match}^{22} is the time at which the (2, 2) mode is matched to the inspiral-plunge waveform and the coefficient A_{22n} is found by a matching procedure. We take $\sigma_{22n} \equiv \omega_n - i\alpha_n$ where $\omega_n > 0$ are the oscillation frequencies and $\alpha_n > 0$ are associated with the inverse of the decay times of the QNM. The full inspiral-merger-ringdown waveform is then given by matching the inspiral-plunge waveform, Eq. (2.13), to the merger-ringdown waveform, Eq. (2.14), waveform close to the EOB light-ring [72].

2.3.2 HG waveforms

HG waveforms aim to capture the main features of the binary inspiral when the mass-ratio of the system is small. Physically, this means we work in a test particle limit. For non-spinning bodies, this limit guarantees that the inspiralling particle instantaneously follows Schwarzschild geodesics while the flux of the orbital parameters is adiabatic. The orbital parameters (E , L_z) are evolved slowly using a prescription for the flux of energy, E , and angular momentum projected along the central black holes spin axis, L_z . Combining these fluxes with the geodesic equations for a test particle in Schwarzschild space-time yields the particles trajectory in Schwarzschild coordinates, (t, p, θ, ϕ) . Here we specialize to (quasi-)circular orbits around a non-spinning central black hole.

The angular momentum per unit mass of the test particle are

$$\frac{L_z}{m_2} = \left(\frac{p}{M}\right) \frac{M}{\sqrt{1 - (3M/p)}}, \quad (2.15)$$

and the inspiral trajectory is generated by the flux of the above quantity which is given by [63],

$$\begin{aligned} \dot{L}_z = & -\frac{32}{5} \frac{m_2^2}{M} \left(\frac{M}{r}\right)^{7/2} \left\{ 1 - \frac{1247}{336} \frac{M}{p} + 4\pi \left(\frac{M}{p}\right)^{3/2} \right. \\ & \left. - \frac{44711}{9072} \left(\frac{M}{p}\right)^2 + \left(c_1^a + c_1^b \left(\frac{M}{p}\right)^{1/2} + c_1^c \left(\frac{M}{p}\right) \right) \right\} \end{aligned} \quad (2.16)$$

The values of the coefficients are $c_1^a = -28.1517$, $c_1^b = 60.9607$, $c_1^c = 40.9998$.

Conservative self force corrections modify the orbital frequency and are given by, [63]

$$\begin{aligned} \frac{d\phi}{dt} := \Omega &= \frac{1}{M} \left(\frac{M}{p}\right)^{3/2} + \delta\Omega, \\ &= \frac{1}{M} \left(\frac{M}{p}\right)^{3/2} \left\{ 1 + \right. \\ &\quad \left. + m_2 \left(d_0 + d_1 \left(\frac{M}{p}\right) + d_{1.5} \left(\frac{M}{p}\right)^{3/2} + d_2 \left(\frac{M}{p}\right)^2 \right) \right\}. \end{aligned} \quad (2.17)$$

The values of the coefficients are $d_0 = 1/8$, $d_1 = 1975/896$, $d_{1.5} = \frac{-27}{10}\pi$, $d_2 = 1152343/451584$. To compute the inspiral trajectory at each instant in time we evolve L_z and ϕ according to Eq. (2.16) and Eq. (2.17) and compute p from L_z from Eq. (2.15). The gravitational-wave strain, $h(t)$, is given by,

$$h(t) = \mathcal{A}(\Omega(t) p(t))^2 \sin[2\phi(t)]. \quad (2.18)$$

The overall scaling \mathcal{A} depends on the source location with respect to the detector. In this chapter we will use normalized template waveforms and so we leave this factor unspecified.

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

2.3.3 TaylorT4

The TaylorT4 approximant is based on the post-Newtonian (PN) expansion (see section III B. of Ref. [26] and references therein). The orbital phase is given by solving the coupled differential equations of the orbital phase and velocity, $\phi(t)$ and $v(t)$, respectively,

$$\frac{d\phi}{dt} - \frac{v^3(t)}{M} = 0, \quad (2.19)$$

$$\frac{dv}{dt} = \mathcal{K}(t), \quad (2.20)$$

where the acceleration $\mathcal{K}(t)$ is given Eq. (A.14).

The gravitational wave strain $h(t)$ can then be expressed directly as

$$h(t) = \mathcal{A}v^2 \sin [2\phi(t)] . \quad (2.21)$$

The overall amplitude \mathcal{A} depends on the source location with respect to the detector. The amplitude is given in [73]. In this chapter we will use normalized template waveforms and so we leave this factor unspecified.

2.4 SNR from inspiral, merger and ringdown

In this section we consider the relative contributions of the different portions of the gravitational-wave coalescence signal to the SNR as a function of the IMRAC's mass. We take the instrument noise power spectral density (PSD) $S_n(f)$ to be the aLIGO high-power, zero-detuned noise PSD [74] which is shown in Fig. 2.1.

The expectation value of the optimal matched filtering SNR, in the case when the signal and template waveforms are identical, is given by [67]

$$\begin{aligned} \left(\frac{S}{N}\right)_{\max} &= (h|h)^{1/2}, \\ &= \left[4\Re \int_{f_{\min}}^{f_{\max}} \left(\frac{f|\tilde{h}(f)|}{\sqrt{fS_n(f)}} \right)^2 d\ln f \right]^{1/2}, \end{aligned} \quad (2.22)$$

where limits of integration correspond to the bandwidth of the detector. Writ-

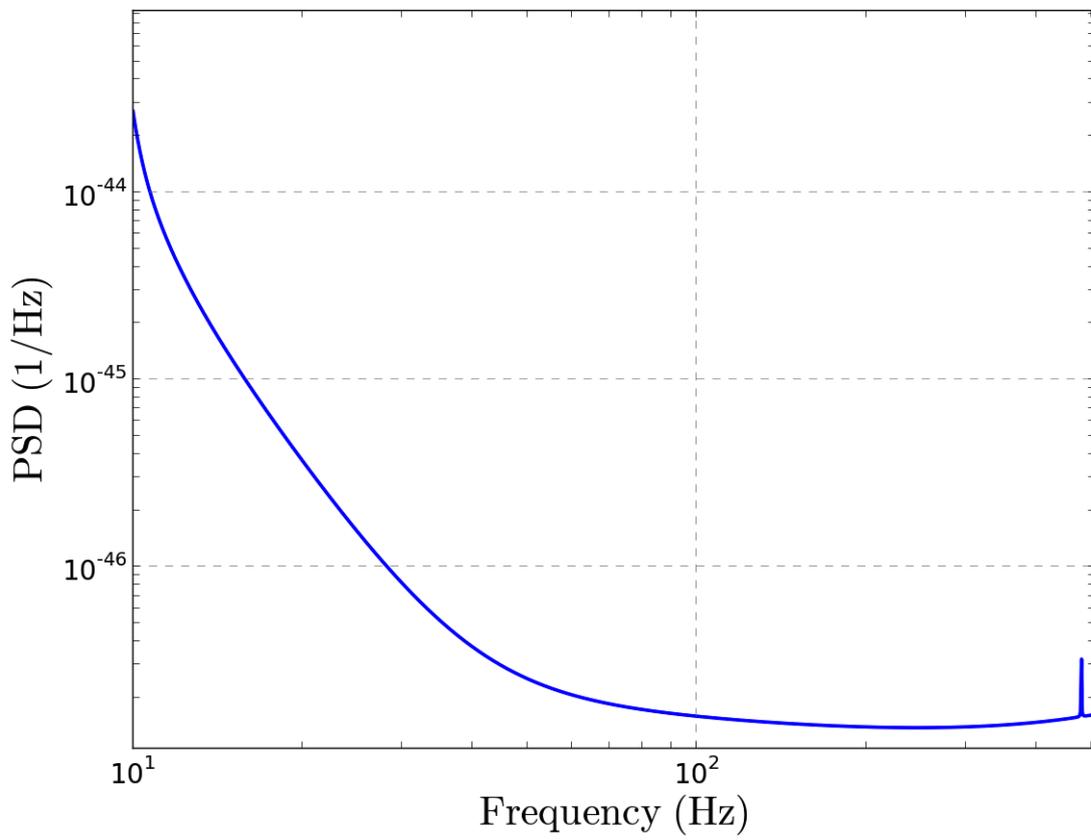


Figure 2.1: Advanced LIGO high-power, zero-detuned noise PSD [74]

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

ing the maximum SNR in the form above clearly separates it into contributions from the signal strain, $f|\tilde{h}(f)|$, and the root-mean-squared (rms) noise spectral amplitude, $\sqrt{fS_n(f)}$, which is the strain signal associated with the detector noise.

One can gain insight into the relative contributions to the SNR from inspiral, merger and ringdown by comparing the gravitational-wave strain to the noise rms value. In Fig. 2.2 we show the strain for selected overhead and face-on (i.e., optimally-located and orientated) IMRAC sources at a fiducial distance of 1Gpc as described by the EOBNR waveform family, and noise rms amplitude. The ISCO frequency of each signal is shown as a vertical line. The strain from merger and ringdown is thus the portion after the ISCO frequency. The contribution to the strain from merger and ringdown from binaries with component masses $(m_1, m_2) = [(200, 20) M_\odot, (200, 2) M_\odot]$, is greater than that of the noise rms amplitude (black curve with triangles). In general, systems with ISCO frequencies between 30 – 100 Hz, merge in the “bucket” of the noise curve, i.e., where the detector is most sensitive. For example, for the $(m_1, m_2) = (200, 20) M_\odot$ system (red dotted curve in Fig. 2.2), the merger and ringdown contribute the bulk of the SNR. Conversely, the inspiral contribution to the SNR is strongly suppressed for such massive systems.

In Fig. 2.3 we show the maximum SNR, Eq. (2.22), as a function of the binary’s total mass produced by inspiral-only and full EOBNR waveforms at four different mass-ratios in the range $1/200 \leq q \leq 1/10$. We construct inspiral-only EOBNR waveforms by Fourier transforming the full waveform into the frequency domain (see Eq. (A.1)) and truncating it at the ISCO frequency. We have considered the SNR for optimally-located and orientated sources at a fiducial distance of 1 Gpc. The lower bound mass of the smaller body is set to $m_2 = 1.4 M_\odot$ which is the canonical neutron-star mass.

We note that the IMRAC systems we are considering are liable to produce an appreciable SNR at distances greater than 1 Gpc. Such cosmological distances correspond to red shifts $z \geq 0.2$ [75]. Gravitational waves from binaries which propagate in an expanding FRW universe experience a Doppler shifting of their frequency such that the observed gravitational-wave frequency f_{obs} is related to the gravitational-wave frequency at the source, f_s , by $f_{\text{obs}} = (1/(1+z)) f_s$ [29]. Such a red shifted frequency corresponds to a binary with a chirp mass which

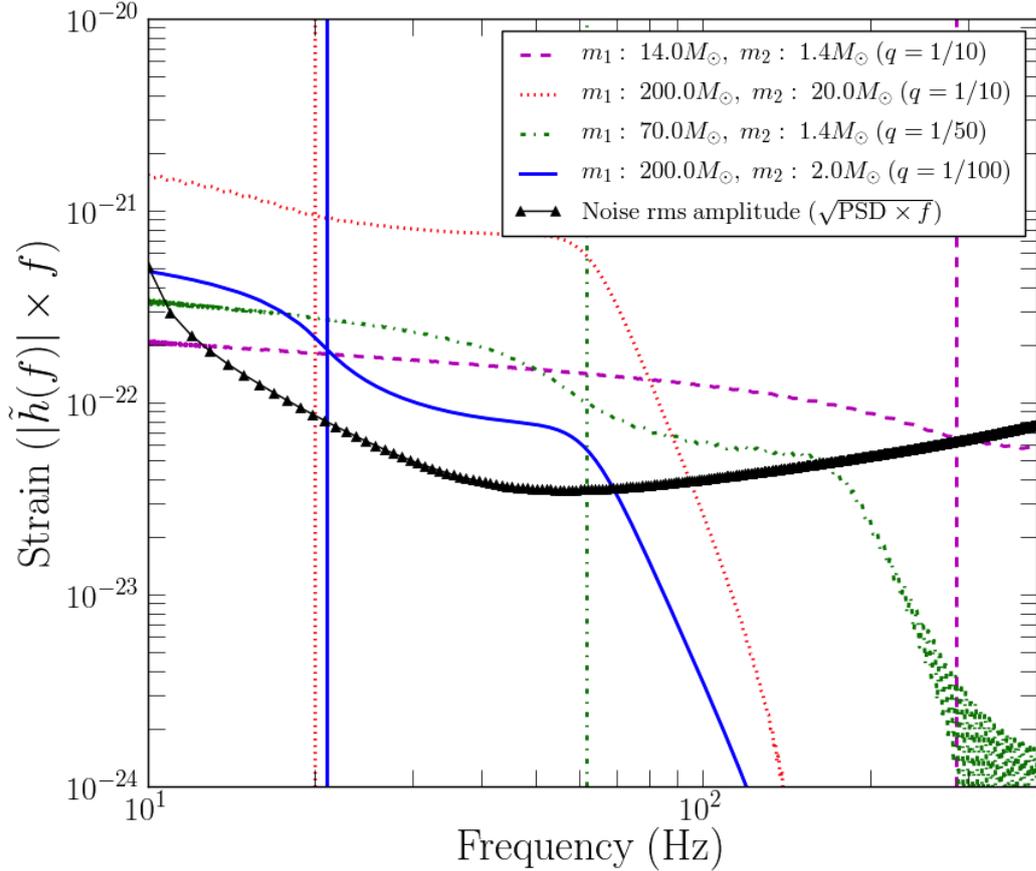


Figure 2.2: Strain of optimally-located and orientated IMRAC sources at a fiducial distance of 1 Gpc as described by EOBNR waveforms, and aLIGO noise. The corresponding ISCO frequency of each signal is shown as a vertical dashed line. The strain from the merger and ringdown from each source contributes after the ISCO frequency. For the sources with component masses $(m_1, m_2) = [(200, 20) M_\odot, (200, 2) M_\odot]$, the strain from merger and ringdown sits above the noise spectrum. The SNR from the full EOBNR waveform and from its inspiral-only portion are shown in Fig. 2.3.

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

is larger by a factor of $(1 + z)$ which is typically referred to as the “red shifted chirp mass” of the binary. Thus, systems at red shift $z = 0.2$ appear to be 1.2 times heavier. Hence, when we refer to binary masses in the following pages, we strictly mean the red shifted masses of the binary.

The lowest total mass for the $q = 1/50$, $1/100$ and $1/200$ subplots is set by fixing the mass of the smaller body to $m_2 = 1.4 M_\odot$. For the $q = 1/10$ subplot in Fig. 2.3 the smallest total mass is set to $M = 35 M_\odot$ as the inspiral phase accounts for $\sim 100\%$ of the SNR below this value. The lower limit of integration of Eq. (2.22) is 10 Hz and the upper limit is 2048 Hz, which is the Nyquist frequency of discretely sampled EOBNR waveforms generated at a sampling rate of $\Delta t = 4096$ s in the time-domain. We only consider systems with total masses such that the ISCO frequency is greater than 10 Hz (our low frequency cut-off). The highest total mass for each of the subplots in Fig 2.3 is set to $M = 300 M_\odot$ which ensures the ISCO frequency is greater than 10 Hz.

As anticipated from Fig. 2.2 there is a significant difference in the SNR between inspiral-only and full EOBNR waveforms which can be seen at all four mass-ratios. We also note that for systems with mass-ratios of $q = 1/10$ with total masses below around $M = 35 M_\odot$ the inspiral phase is the dominant source of SNR. If we consider 3% as a fiducial value of the difference between the full SNR and the one associated to the inspiral-only waveform – which leads to a loss in detection rates of 10% – this happens at $M \approx 35 M_\odot$ for $q = 1/10$. For binaries with $q = 1/50$, $1/100$ and $1/200$, the minimum difference in SNR between inspiral-only and full waveforms is $\approx 6\%$, 15% and 40% respectively for the mass ranges considered in Fig 2.3.

In summary, we have shown that inspiral-only templates will miss a significant portion of the total SNR of IMRAC signals over the bulk of the detectable mass-space. Future searches will therefore require templates that can match the full inspiral-merger-ringdown. However, there is a small region of the parameter space for which inspiral-only templates may suffice for searches, without inducing drastic losses in detection rates. In the following section we quantify the effectiveness of inspiral-only templates for searching for full coalescence signals in aLIGO.

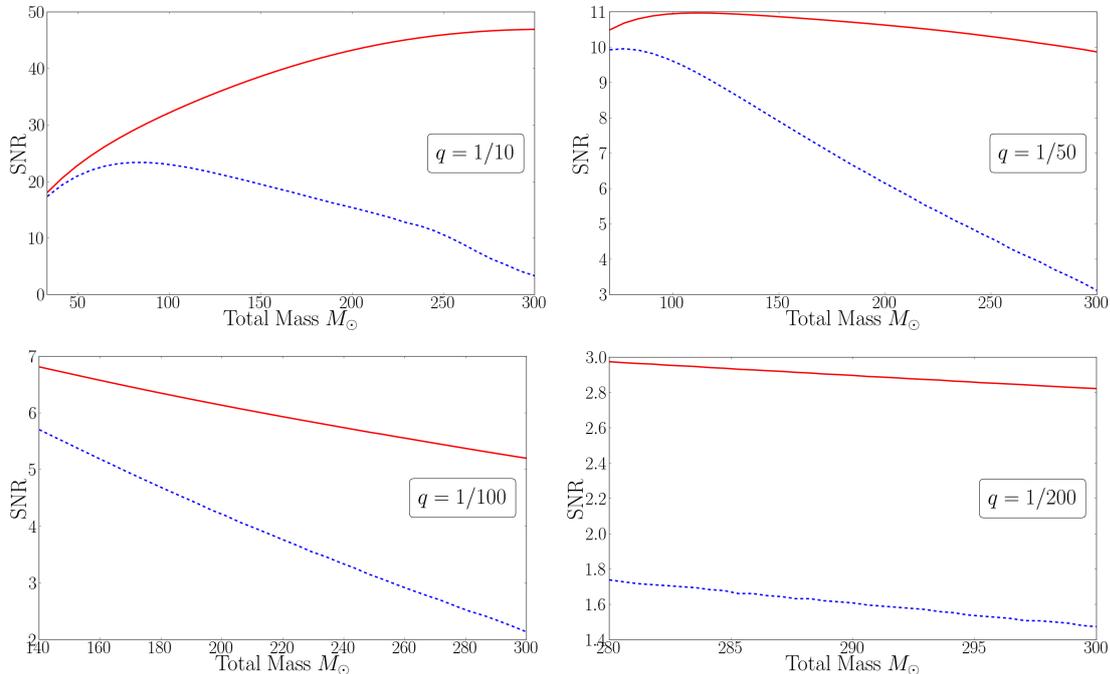


Figure 2.3: SNR of optimally-located and orientated IMRAC sources at a fiducial distance of 1 Gpc vs total mass for four different mass ratios; $q = (1/10, 1/50, 1/100, 1/200)$. The solid line is the SNR from full EOBNR waveforms and the dashed line from inspiral-only EOBNR waveforms truncated at the ISCO frequency in the frequency domain. The lower bound mass of the smaller body is set to $m_2 = 1.4 M_\odot$ which is the canonical neutron-star mass. The lowest total mass for the $q = 1/50, 1/100$ and $1/200$ subplots in Fig. 2.3 is set by fixing the mass of the smaller body to $m_2 = 1.4 M_\odot$. For the $q = 1/10$ subplot the smallest total mass is set to $M = 35 M_\odot$ as the inspiral accounts for $\sim 100\%$ of the SNR below this value. We only consider systems with total masses such that the ISCO frequency is greater than 10 Hz (our low frequency cut-off). The highest total mass for each of the subplots in Fig 2.3 is set to $M = 300 M_\odot$ which ensures the ISCO frequency is greater than 10 Hz. We find that there is a non-negligible contribution to the SNR from merger and ringdown in IMRAC signals above a total mass of around $M = 35 M_\odot$. The difference in SNR between inspiral-only and full waveforms is at the 3% level at around $M = 35 M_\odot$ at $q = 1/10$. For binaries with $q = 1/50, 1/100$ and $1/200$, the minimum difference in SNR between inspiral-only and full waveforms are at the 6%, 15% and 40% levels, respectively, over our mass range of interest. For IMRACs of astrophysical interest, more extreme mass ratios correspond to a greater total mass, which can place the merger and ringdown at a frequency where the detector has the greatest sensitivity.

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

2.5 Effectiveness of inspiral-only templates for IMRAC searches

We have shown in Sec. 2.4 that the SNR from merger and ringdown will provide a significant contribution to the total SNR over a broad portion of the IMRAC mass-space, c.f. Fig. 2.3. There is however a small portion of the parameter space where the SNR is dominated by the inspiral phase which can be seen, e.g. for $q = 1/10$ binaries in Fig. 2.3 with total masses $M \leq 35 M_{\odot}$. Thus it is important to quantify the effect of using inspiral-only templates to search for IMRAC signals which contain inspiral, merger and ringdown phases. The use of template waveforms that are not exact representations of the signals they filter degrades the SNR, as the optimal SNR can be recovered only when the template waveform corresponds exactly to the signal h , cf. Eq. (2.22). In practice however, we do not have access to an exact representation for h . Using a non-exact template waveform T to filter h caps the maximum recoverable SNR to [76]

$$\begin{aligned} \left(\frac{S}{N}\right) &= \max_{\vec{\theta}} \frac{(h(\vec{\lambda})|T(\vec{\theta}))}{(T(\vec{\theta})|T(\vec{\theta}))^{1/2}}, \\ &= \epsilon \left(\frac{S}{N}\right)_{max}, \end{aligned} \quad (2.23)$$

where $\vec{\lambda}$ and $\vec{\theta}$ represent the parameter vector of the signal and template, respectively. We define ϵ as the *effectiveness* of a template waveform family T at recovering the maximum SNR from a gravitational-wave signal h ; by definition $0 \leq \epsilon \leq 1$. This quantity is also referred to as the “fitting factor” in the literature [76]. It is convenient to define waveforms normalized to unit norm as $\hat{\mathbf{a}}(f) = \tilde{a}(f)/(a|a|)^{1/2}$ so that $(\hat{\mathbf{h}}|\hat{\mathbf{h}}) = (\hat{\mathbf{T}}|\hat{\mathbf{T}}) = 1$ and the effectiveness can be written succinctly as

$$\epsilon = \max_{\vec{\theta}} (\hat{\mathbf{h}}(\vec{\lambda})|\hat{\mathbf{T}}(\vec{\theta})). \quad (2.24)$$

Using normalized waveforms also has the advantage of eliminating the dependence of the waveforms on the source orientation and distance, which enter as an overall scaling.

Calculating the effectiveness, Eq. (2.24), requires maximizing over the component masses (m_1, m_2) and the time and phase at coalescence. We can efficiently maximize over the time and phase by Fourier transforming the integrand of the noise-weighted inner-product [77],

$$z(t_c) = 4 \int_{f_{min}}^{f_{max}} df \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} e^{2\pi i f t_c}, \quad (2.25)$$

which yields a complex time series whose elements correspond to the inner-product of a and b as one of the signals is time-shifted with respect to the other. We can efficiently find the time at coalescence, t_c , by finding the time at which the norm of this time series is a maximum. The phase at coalescence ϕ_c is then automatically given by finding the argument of the time-series at its peak amplitude. We thus modify the inner product $(a|b)$:

$$(a|b) \rightarrow (a|b)' = \max_{t_c} |z(t_c)|, \quad (2.26)$$

which we will adopt as the definition of the inner-product for the remainder of this chapter.

To compute the effectiveness of an inspiral-only IMRAC search we evaluate Eq. (2.24) for signals covering the IMRAC mass space. We take as our signal waveform, h , the full inspiral-merger-ringdown EOBNR waveform. We take the template, T , to be an *inspiral-only* EOBNR waveform, formed by truncating the full EOBNR waveform at the ISCO frequency in the frequency domain. With such signals and templates the effectiveness provides a measure of the maximum SNR which could be achieved through using an inspiral-only template to filter full coalescence-signals. To get a broad coverage of the IMRAC mass space we compute Eq. (2.24) for signals whose source masses cover the ranges $1.4 M_\odot \leq m_2 \leq 18.5 M_\odot$ and $24 M_\odot \leq m_1 \leq 200 M_\odot$, with mass ratios spanning the range $q := m_2/m_1 \in [1/140, 1/10]$. For each signal we evaluate the effectiveness, Eq. (2.24), where the template T describes the inspiral-only portion of an EOBNR waveform. We maximize over time and phase by maximizing the inner product of the signal with an inspiral-only EOBNR template, Eq. (2.26). The maximization over the masses is performed by finding the largest inner product between the

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

signal and a bank of template waveforms. The template bank is characterised by intrinsic parameters (\mathcal{M}, η) which are the chirp mass and symmetric mass-ratio of the system: $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, $\eta = (m_1 m_2) / (m_1 + m_2)^2$, where m_1 and m_2 are the component masses of the binary. The bank spans an extended mass range $3 M_\odot \leq \mathcal{M} \leq 30 M_\odot$ and $0.0065 \leq \eta \leq 0.082$. The results of the effectiveness of an inspiral-only IMRAC search are shown in Fig. 2.5.

Inspiral-only templates are $\sim 98\%$ effective at filtering full coalescence signals for total masses $M \lesssim 30 M_\odot$. Such systems have an ISCO frequency $150 \text{ Hz} \lesssim f_{\text{ISCO}}$ which is well within the peak sensitivity of the noise curve. However for the bulk of the mass space the effectiveness is below 75% . This is unsurprising given the SNR curves in Fig. 2.3 which clearly show the importance of the contribution of merger and ringdown to the SNR.

The loss in SNR incurred through using inspiral-only templates directly affects detection rates. Because the SNR scales inversely with the distance, the observable volume will scale with the cube of the effectiveness. Assuming that GW sources are isotropically distributed in the sky, the fractional loss in detection rates will be $1 - \epsilon^3$. The percentage loss in detection rates through using inspiral-only EOBNR templates to recover the full coalescence signal is also shown in Fig. 2.5. Over a broad portion of the mass-space inspiral-only templates incur losses in detection rates between $60 - 85\%$. As the total mass of the binary approaches $440 M_\odot$ the Schwarzschild ISCO frequency, Eq. (2.1), approaches 10 Hz which is near the low-frequency cut-off of the detectors. Hence the relative contribution of the inspiral phase to the coalescence signal of heavier systems diminishes until the only contribution is from the merger and ringdown. This is a striking indication of the need of merger and ringdown in IMRAC template waveforms. This suggests the importance of full numerical simulations in this regime in order to construct such a reliable waveform family including inspiral, merger, and ringdown phases.

We identify three regions in the m_1 - m_2 plane in which various searches could be constructed. The regions are defined by contours of constant effectiveness which are approximately given by $\mathcal{C} = (m_1/M_\odot)\sqrt{m_2/M_\odot}$, which are found purely empirically, with $1.4 M_\odot \leq m_2 \leq 18.5 M_\odot$ and mass-ratios $q \in [1/140, 1/10]$. The effectiveness is related to \mathcal{C} by $\epsilon \approx 1/100 \times (1.6 \mathcal{C} - 7.3 \times 10^{-3} \mathcal{C}^2)$.

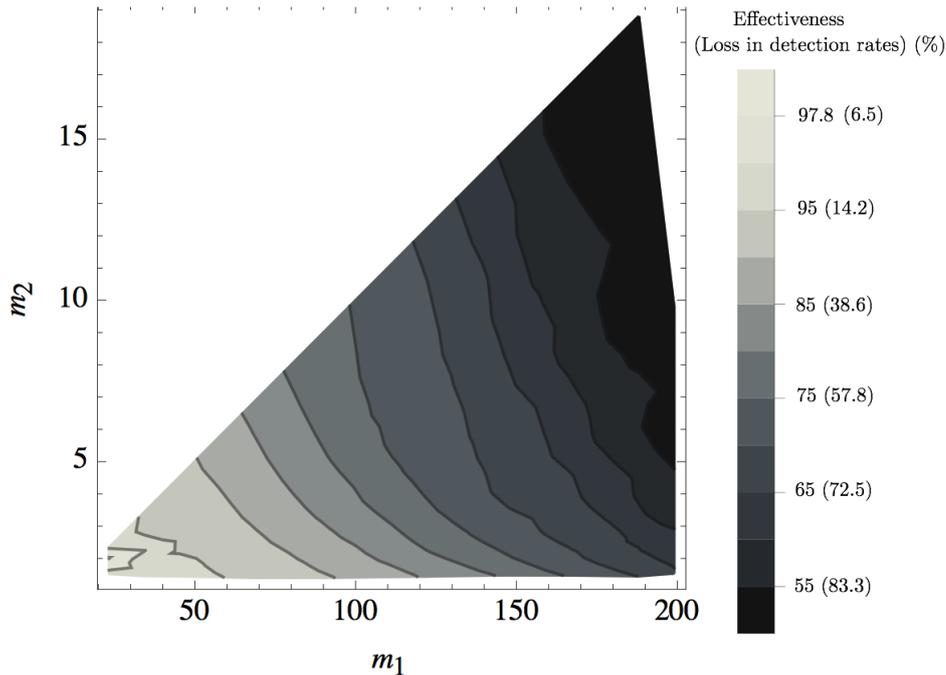


Figure 2.4: Effectiveness of inspiral-only EOBNR templates to filter full inspiral, merger and ringdown EOBNR signals as a function of the source component masses, and corresponding losses in detection rates. The diagonal corresponds to mass-ratio $q = 1/10$. Inspiral-only EOBNR templates are constructed by truncating the full waveform at the ISCO frequency in the frequency-domain. For the bulk of the parameter space inspiral-only templates are $\lesssim 75\%$ effective at filtering inspiral, merger and ringdown signals. Inspiral-only templates are $\sim 97 - 98\%$ effective for total masses $M_{\odot} \lesssim 30 M_{\odot}$. Inspiral-only templates within the 90%-effectiveness contour should be sufficient for IMRAC searches without incurring greater than 30% losses in detection rates.

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

Between the 97%- and 90%-effectiveness contours, the losses in detection rates are between $10\% \lesssim L \lesssim 27\%$ and so an inspiral-only search could be sufficient without incurring drastic losses in detections. The region bound from below in effectiveness by the 90%-effectiveness contour is defined by $\mathcal{C} \leq 100$, with the effectiveness increasing with decreasing \mathcal{C} . Between the 90%- and 80%- effectiveness contours the losses in detection rates are around $27\% \lesssim L \lesssim 50\%$. Thus, within this region searches will be limited by the lack of merger and ringdown in template waveforms, though an inspiral-only search would be feasible in principle. This contour is defined by $100 \lesssim \mathcal{C} \lesssim 150$. Below the 80%-effectiveness contour inspiral-only searches will incur losses in detection $50\% < L$ and so merger and ringdown will be crucial for searches. The region bound from above in effectiveness by the 80%-effectiveness contour which is defined by $150 \lesssim \mathcal{C}$, with effectiveness decreasing with increasing \mathcal{C} . The results are summarized in Table 2.1.

Effectiveness of inspiral-only search, ϵ (%)	Loss in detection rates, L (%)	Contours of constant effectiveness in m_1 - m_2 plane ($\mathcal{C} = (m_1/M_\odot)\sqrt{m_2/M_\odot}$) within mass range of interest	Implication for searches
$90\% \lesssim \epsilon \lesssim 97\%$	$10\% \lesssim L \lesssim 27\%$	$\mathcal{C} \lesssim 100$	Inspiral-only search sufficient but with non-negligible loss in detection rates.
$80\% \lesssim \epsilon \lesssim 90\%$	$27\% \lesssim L \lesssim 50\%$	$100 \lesssim \mathcal{C} \lesssim 150$	Inspiral-only search possible but limited by lack of merger and ringdown. Could potentially lose half of signals with inspiral-only templates.
$\epsilon \lesssim 80\%$	$50\% \lesssim L$	$150 \lesssim \mathcal{C}$	Merger and ringdown crucial for searches. Will miss over half of signals with inspiral-only templates.

Table 2.1: Effectiveness of inspiral-only searches, the corresponding loss in detection rates and the region in the m_1 - m_2 plane bounded by constant-effectiveness contours. For a given region in the m_1 - m_2 plane bounded by constant-effectiveness contours we summarize the implications for IMRAC searches.

2.6 Comparison of inspiral-only waveforms

We have shown that merger and ringdown are crucial for effective searches over a large portion of the IMRAC mass space, though there is a small region in which an inspiral-only search could be constructed without incurring losses in detection rates greater than around 27%. For this region, it is therefore important to study whether currently available waveforms are sufficiently accurate. The inspiral phase can be computed using perturbative expansions and thus it is interesting to quantify the consistency of different expansions.

To assess the effectiveness of the EOBNR inspiral, we employ a waveform family designed to approximate intermediate mass-ratio inspirals which we refer to as “Huerta-Gair” (HG) waveforms [44]. HG waveforms describe only the inspiral portion of the coalescence signal.

We repeat the study done in the previous Section using now the HG waveform family as the signal h and inspiral-only EOBNR as the template T . The results are reported over the whole parameter space in Fig. 2.5. For completeness, in Table 2.2 we also show the values of the effectiveness, Eq. 2.24, for selected mass combinations of EOBNR inspiral-only templates for filtering full EOBNR and HG signals respectively.

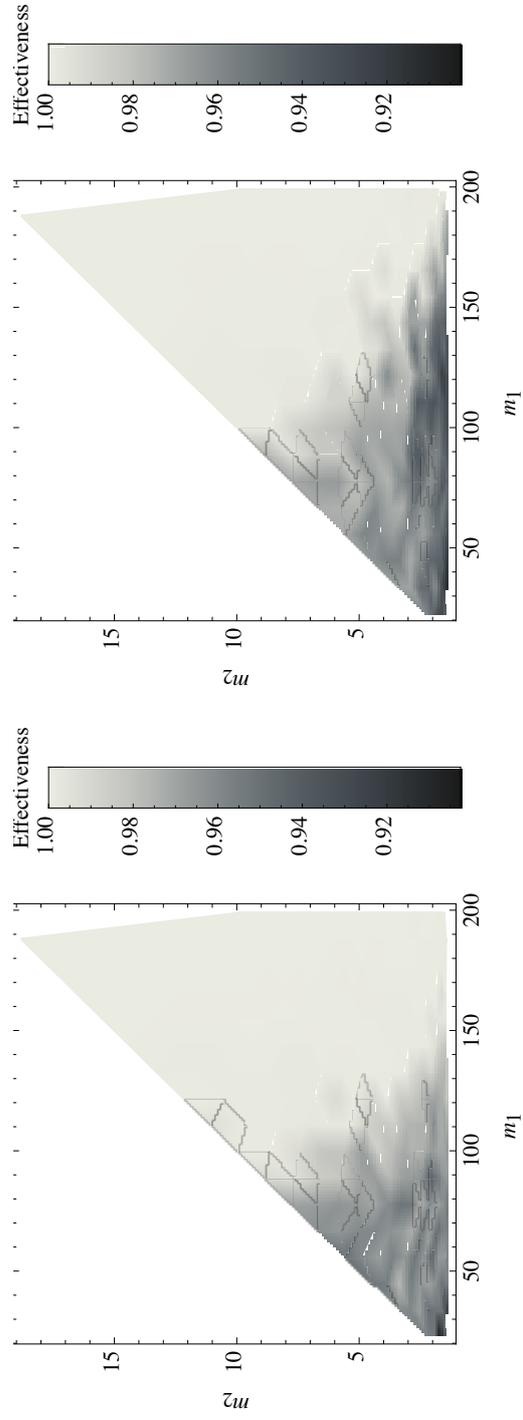


Figure 2.5: Effectiveness of inspiral-only EOBNR templates at filtering HG signals (left) and TaylorT4 signals (right) as a function of the source component masses encoded in the signal.

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

m_1 (M_\odot)	m_2 (M_\odot)	Signal waveforms		
		full EOBNR	Huerta-Gair	TaylorT4
50	5	0.90	0.95	0.96
100	5	0.76	0.97	0.97
200	5	0.53	0.99	0.99
50	1.4	0.96	0.96	0.90
100	1.4	0.86	0.98	0.94
200	1.4	0.67	0.98	0.98
5	5	0.99	0.89	0.99
20	20	0.99	0.92	0.99
100	100	0.52	0.51	0.50

Table 2.2: Summary of effectiveness of *inspiral-only* EOBNR template waveforms in recovering signals modelled using different waveform families – full EOBNR, Huerta-Gair and TaylorT4 – for selected component masses. Merger and ring-down become more prominent in the coalescence signal as the total-mass of the system is increased. The EOBNR inspiral is typically better at matching HG signals in the IMRAC regime than TaylorT4 signals. Results for equal mass-ratio systems are shown for reference below the horizontal line.

For the high-mass part of the mass-space the effectiveness of the EOBNR inspirals with respect to the HG waveforms is close to 100%. This is perhaps unsurprising because very high mass systems will have short inspirals and possible differences in the waveforms will not produce a significant degradation of SNR when matched over a small number of wave cycles. However, for lighter systems the effectiveness can be as low as 90%, which occurs in the region of mass space in which inspiral-only searches would be most feasible (see Table 2.1).

For reference we also compare inspiral-only EOBNR templates to TaylorT4 signal waveforms (which are inspiral-only). We construct signal waveforms on the same grid in $m_1 - m_2$ as for HG waveforms and use the same template bank of inspiral-only EOBNR waveforms. The results are summarized in the right panel of Fig. 2.5, and in Table 2.2 for selected masses. We find that the EOBNR inspiral has good filtering efficiency for the TaylorT4 waveform family. However, EOBNR is clearly a better match to the HG waveform family over a larger range of masses and mass ratios than to TaylorT4. This can be seen more clearly by comparing the subplots in Fig. 2.5. This is unsurprising given that the PN expansion is

unreliable at high velocities and highly asymmetrical mass-ratios. For orbital velocities $v/c = (M\pi f)^{1/3} \gtrsim 0.2$ the PN energy flux deviates significantly from numerical results, see [e.g., 38; 39]. A binary at its ISCO frequency has $v/c \sim 0.4$ which is well beyond the region of validity.

2.7 Discussion and conclusion

We have shown that over the bulk of the IMRAC mass space, merger and ring-down contribute significantly to the gravitational-wave coalescence signal. This happens despite the suppression of the power in the merger and ringdown in signals from binaries with very asymmetric mass ratios. The importance of merger and ringdown is due to the greater sensitivity to these waveform portions for high-mass signals, for which most of the inspiral may fall at frequencies below the detector’s sensitive band.

However there is a relatively large patch in mass space in which the inspiral-only waveforms are more than 90% effective. We identified three regions in which different searches could be considered appropriate based on thresholds of acceptable losses in detection rates. The mass space splits into a region in which inspiral-only searches could be feasible, incurring losses in detection rates of up to $\sim 27\%$; a region in which searches would be limited by lack of merger and ring-down in template waveforms, incurring losses in detection rates up to 50%; and a region in which merger and ringdown are critical to prevent losses in detection rates over 50%. The search regions are summarized in Table 2.1.

We have further shown that in the region of the IMRAC mass space in which inspiral-only searches are feasible, approximants adapted to asymmetric mass-ratio binaries are important, as here the binary is liable to have highly relativistic velocities $v \gtrsim 0.2$. We considered a waveform family designed to describe intermediate mass-ratio binaries which we referred to as the “Huerta-Gair” (HG) waveform family. By computing the effectiveness of inspiral-only EOBNR waveforms to filter signals described by the HG waveform family, we showed losses in recovered SNR could be as great as 10%. In Table 2.2 we summarize the effectiveness of the signal–template combinations used in the chapter.

We believe that template waveforms for IMRAC searches will benefit from

2. DETECTING INTERMEDIATE MASS-RATIO BINARIES WITH ADVANCED DETECTORS

calibration to several numerical simulations. We note that there already exists one very short numerical waveform of a $q = 1/100$ binary which we have not used in our study, and which EOBNR is not currently calibrated to [45].

Chapter 3

TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

3.1 Introduction

Astronomy and tests of fundamental physics with gravitational waves from compact binary coalescence (CBC) will ultimately be limited by our ability to estimate the binary's source parameters from the gravitational-wave signal [e.g., 78; 79; 80]. CBC sources with total masses in the range $2 M_{\odot} \lesssim M \lesssim 500 M_{\odot}$ will be amongst the prime sources for Advanced LIGO [58] and Advanced Virgo [59] when they begin operating around 2015 [5].

In a Bayesian treatment of parameter estimation, one is interested in the posterior probability distribution of the set of source parameters of the underlying model given an observed stretch of data. Waveform computation represents the majority of the computation cost in the Bayesian analysis of CBC sources, so the total computational cost scales roughly linearly with waveform generation. This

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

becomes burdensome when one needs to explore a large dimensional parameter space as the number of waveform computations is large, e.g., $\mathcal{O}(10^6)$ [25].

Recently, Cannon *et al.* [49] showed that a truncated singular value decomposition (SVD) can be applied to gravitational-wave template banks which span the two parameters describing the masses of the coalescing binary. The SVD decomposes the bank into a set of “basis templates” and projection coefficients. In general, the number of basis templates is much less than the total number of templates in the bank. Furthermore, the projection coefficients can be interpolated across the domain of the bank [50]. Template waveforms can thus be interpolated. Because the *intrinsic* parameter space of CBC sources with non-spinning components is two-dimensional (the two mass parameters), it is possible to set up the waveform computation for parameter estimation such that the waveform calculations are done by interpolation alone. However, the errors incurred from interpolation could, in principle, affect parameter-estimation accuracy.

In this chapter, we describe the application of SVD-interpolated waveforms to CBC parameter estimation. For a simulated data set containing a gravitational wave signal we provide a proof of principle that SVD-interpolated waveforms can be used for parameter estimation without significantly affecting the accuracy of the inferred probability distributions of the source parameters. We further show that the computational cost of using interpolated waveforms is around an order of magnitude less than that of commonly-used time-domain waveform families. This technique has the potential to increase the computational efficiency of CBC parameter estimation when the computational cost is dominated by waveform computation. Our application of the SVD is limited to a small patch of parameter space about the injected signal value.

This chapter is organized as follows. In Sections 3.2 and 3.4 we outline the principles of parameter estimation for CBCs and interpolating template waveforms based on the SVD, respectively. In Sec. 3.5 we describe the application of SVD-interpolated waveforms to parameter estimation. In Sec. 3.6 we compare the results of using interpolated and non-interpolated waveforms for parameter estimation and compare the computational cost of interpolation to using non-interpolated waveform families. In Sec. 3.7 we consider the future of using SVD-interpolated waveforms for parameter estimation and discuss the technical

requirements of implementing these waveforms in parameter-estimation pipelines.

3.2 CBC parameter estimation

The central quantity of interest in Bayesian parameter estimation is the posterior probability density function (PDF) of a set of parameters $\vec{\theta}$ which parameterize a model, \mathcal{H} , assumed to describe a data set d . The PDF is related to the likelihood function and prior probability via Bayes' theorem and is given by

$$p(\vec{\theta}|d, \mathcal{H}) = \frac{\mathcal{P}(\vec{\theta}|\mathcal{H}) \mathcal{L}(d|\vec{\theta}, \mathcal{H})}{p(d|\mathcal{H})}, \quad (3.1)$$

where $\mathcal{L}(d|\vec{\theta}, \mathcal{H})$ is the likelihood function and $\mathcal{P}(\vec{\theta}|\mathcal{H})$ is the prior probability which encodes our *a priori* belief in the distribution of $\vec{\theta}$. The quantity in the denominator, $p(d|\mathcal{H})$, is known as the evidence and is an overall normalization factor which we will not deal with here.

The CBC parameter vector $\vec{\theta}$ is high-dimensional. The phasing and amplitude of a waveform from a non-spinning coalescing compact binary source is controlled by two mass parameters, the chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ and symmetric mass ratio $\eta = (m_1 m_2) / (m_1 + m_2)^2$, where m_1 and m_2 are the component masses of the binary. In addition, a gravitational wave source with respect to the Earth is specified by location dependent parameters. These are the distance from the Earth D , inclination ι , right ascension α , declination δ , polarization phase φ and time and phase at coalescence, t_c and ϕ_c . In general, the CBC parameter vector $\vec{\theta}$ is nine-dimensional for circular binaries with non-spinning components. An illustration of a coalescing binary system, with the set of source parameters is shown in Fig. 3.1. One of the goals of gravitational-wave astronomy is to estimate the PDF of the parameters of a candidate gravitational wave source in order to assign a meaningful probability to our measurements of the source properties and demographics. To compute the right hand side of Eq. (5.1), we directly evaluate the likelihood function, $\mathcal{L}(d|\vec{\theta}, \mathcal{H})$. Under the hypothesis that the data, d , consists of Gaussian, stationary noise n and a gravitational-wave signal $h(\vec{\theta})$, the

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

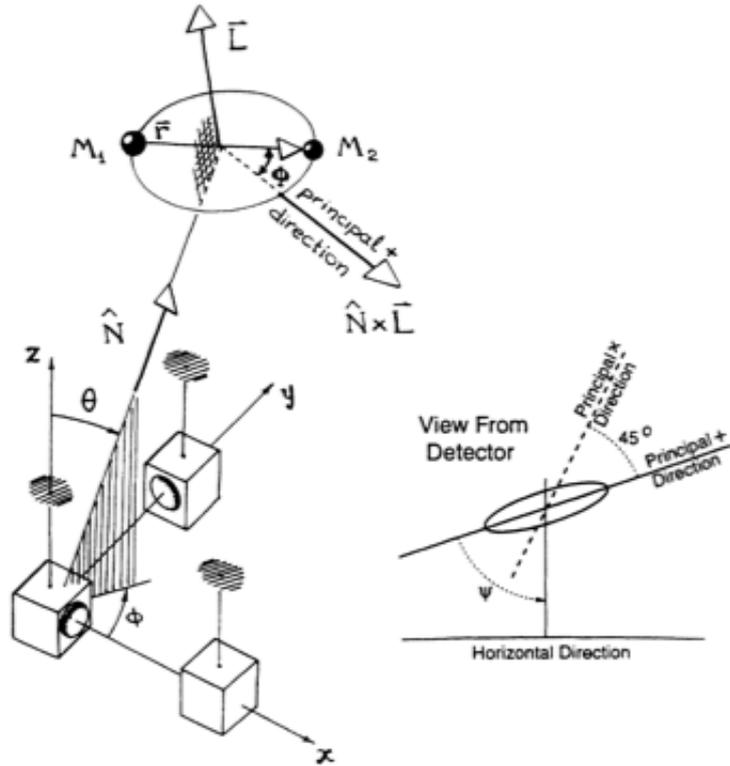


Figure 3.1: Cartesian coordinate system (x, y, z) attached to a ground based interferometer and a coalescing binary system in that frame [81]. The binary's components are labelled by two mass parameters m_1, m_2 . The sky position is determined by the two polar angles (θ, ϕ) which fixes the position vector \hat{N} . The inclination is given by $\cos i = \hat{N} \cdot \hat{L}$ where \hat{L} is the orbital angular momentum vector of the binary.

likelihood is a Gaussian [12]:

$$\mathcal{L}(d|\vec{\theta}, \mathcal{H}) = e^{-(d-h(\vec{\theta})|d-h(\vec{\theta}))^2/2}, \quad (3.2)$$

In the above, $(a|b)$ is the usual noise-weighted inner-product,

$$(a|b) = 4\Re \int_{f_{min}}^{f_{max}} df \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)}, \quad (3.3)$$

and $S_n(f)$ is the detector’s noise power spectral density (PSD). The limits of integration correspond to the bandwidth of the detector. A significant computational cost of evaluating the likelihood comes from computing the template waveform $h(\vec{\theta})$ at each point in the parameter space.

The full PDF is multi-dimensional and to get estimates on an individual parameter we work with the *marginalized* PDF of a single parameter $\theta_A \in \vec{\theta}$. Writing $\vec{\theta} = (\theta_A, \vec{\Theta})$, the marginalized one-dimensional PDF of θ_A is thus:

$$p(\theta_A|d, \mathcal{H}) = \int_{\vec{\Theta}} d\vec{\Theta} p(\vec{\theta}|d, \mathcal{H}). \quad (3.4)$$

To efficiently evaluate the likelihood function we typically use a stochastic sampling algorithm. Here we employ Markov-chain Monte Carlo (MCMC), whose application to gravitational-wave parameter estimation is described in detail in [13], and in Sec. 3.2.1.

We use the stationary phase approximation (SPA) inspiral waveforms for both our simulated signal and template model. This allows us to directly “inject” the signal waveform into simulated frequency-domain noise without performing an additional Fourier transformation, which could introduce spurious artifacts related to the abrupt in-band termination of a time-domain waveform. The post-Newtonian frequency-domain waveform for a face-on, overhead binary has the form

$$\tilde{h}(f) = \frac{\mathcal{A}(\mathcal{M}; f)}{D} e^{2\pi i f t_c - i\phi_c + i\Psi(\mathcal{M}, \eta; f)}, \quad (3.5)$$

where expressions for the phase $\Psi(\mathcal{M}, \eta; f)$ is given by Eq. (A.11) for TaylorF2 post-Newtonian waveforms at 3.5 post-Newtonian order and the amplitude is

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

given by [29]

$$\mathcal{A}(\mathcal{M}; f) = \frac{1}{\pi^{2/3}} \left(\frac{5}{24}\right)^{1/2} \mathcal{M}^{5/6} f^{-7/6}, \quad (3.6)$$

in geometric units, $G = c = 1$. In general, the template model will be determined by the needs of a particular analysis. Below, we describe the MCMC algorithm which we use to compute the posterior PDF, Eq. 5.1.

3.2.1 MCMC

Our goal is to compute the posterior PDF of the set of astrophysical parameters $\vec{\theta}$ which describe a compact binary, given a data set containing a gravitational-wave signal from such a source. The posterior PDF is given by Eq. (5.1). In principle, one could simply grid up the parameter space of interest, and compute the posterior PDF at each grid point. However, a number of problems become immediately apparent with this approach. Firstly, the compact binary parameter space is high dimensional (9 dimensional for binaries with non-spinning components, and 15 dimensional for binaries with spinning components). Assuming an equally spaced grid of N points in each dimension, one would require between $N^9 - N^{15}$ computations. Even a modest grid with $N \sim \mathcal{O}(100)$ points in each of the nine dimensions would require a minimum of $100^9 \sim 10^{18}$ computations. This coupled with the fact that evaluating the posterior PDF is an expensive computation renders this approach hopeless. Furthermore, the problems with the approach are exacerbated when the dimensionality is increased to 15, for obvious reasons. Hence, it is useful to consider algorithms which stochastically sample the parameter space preferentially, in regions of high posterior probability.

Ultimately we want the output of such an algorithm to be a set of sample points $\{\mathcal{P}\}$ which are drawn from the posterior PDF. In the context of MCMC, this set is known as a “Markov chain” [82], and the goal of the MCMC is to produce such a Markov chain via a stochastic process. The stochastic process we will consider is the classical Markovian random walk through the parameter space. By “Markovian” we mean that the random walk is such that given the i^{th} state in the Markov chain, the next state ($i + 1$) is determined solely by i , and no others. In the context of parameter estimation this means that the $(i + 1)^{th}$ sample of the posterior PDF in the Markov chain, $P(\vec{\theta}_{i+1})$, depends only on the

i^{th} sample, $P(\vec{\theta}_i)$. Thus, by Bayes' theorem we have

$$P(\vec{\theta}_{i+1}) = \frac{P(\vec{\theta}_i) P(\vec{\theta}_i|\vec{\theta}_{i+1})}{P(\vec{\theta}_{i+1}|\vec{\theta}_i)}. \quad (3.7)$$

The above equation provides a criterion known as ‘‘detailed balance’’ [82] which the Markovian random walk must obey:

$$P(\vec{\theta}_{i+1}) P(\vec{\theta}_{i+1}|\vec{\theta}_i) = P(\vec{\theta}_i) P(\vec{\theta}_i|\vec{\theta}_{i+1}). \quad (3.8)$$

In order to utilize the random walk, encoded in the detailed balance condition Eq. (3.8), to form the Markov chain of samples, we need to be able to specify a means of computing the conditional probability of state $(i + 1)$ given state i , $P(\vec{\theta}_{i+1}|\vec{\theta}_i)$, and vice versa. For practical purposes one ought to be able to propose a transition from state i to state $(i + 1)$ and accept the proposal provided that detailed balance is satisfied. The proposal should be arbitrary in the sense that it should reflect the requirements of a particular analysis. For example, when estimating the parameters of compact binaries we may want to explore known correlations or degeneracies between parameters, or pairs of parameters. We thus require that we have an additional distribution $g(\vec{\theta}_{i+1}|\vec{\theta}_i)$ which proposes which point in parameter space should be explored next. This is often referred to as a ‘‘jump proposal’’ in the literature [13; 25]. However, this alone is not sufficient to satisfy detailed balance, Eq. (3.8), because not all transitions will be allowed. Hence we can only accept transitions with a particular probability $A(\vec{\theta}_{i+1}|\vec{\theta}_i)$ which is known as the ‘‘acceptance probability’’. Thus we can write the conditional probability of state $(i + 1)$ given state i as

$$P(\vec{\theta}_{i+1}|\vec{\theta}_i) = g(\vec{\theta}_{i+1}|\vec{\theta}_i) A(\vec{\theta}_{i+1}|\vec{\theta}_i) \quad (3.9)$$

and the detailed balance condition, Eq. (3.8), is given by

$$P(\vec{\theta}_{i+1}) g(\vec{\theta}_{i+1}|\vec{\theta}_i) A(\vec{\theta}_{i+1}|\vec{\theta}_i) = P(\vec{\theta}_i) g(\vec{\theta}_i|\vec{\theta}_{i+1}) A(\vec{\theta}_i|\vec{\theta}_{i+1}) \quad (3.10)$$

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

Defining the ratio of acceptance probabilities,

$$r := \frac{A(\vec{\theta}_{i+1}|\vec{\theta}_i)}{A(\vec{\theta}_i|\vec{\theta}_{i+1})} = \frac{P(\vec{\theta}_{i+1}) g(\vec{\theta}_i|\vec{\theta}_{i+1})}{P(\vec{\theta}_i) g(\vec{\theta}_{i+1}|\vec{\theta}_i)}, \quad (3.11)$$

we accept a transition with probability

$$\text{acceptance probability} = \min(1, r). \quad (3.12)$$

We denote the ratio r as the ‘‘acceptance ratio’’. Together with a (user specified) jump proposal, the above condition provides us with all the information required to draw samples from the posterior distribution. We can compute the ratio $P(\vec{\theta}_{i+1})/P(\vec{\theta}_i)$ as this is just a ratio of likelihoods, which can be computed from Eq (3.2), and once we specify a jump proposal $g(\cdot)$, we can evaluate the acceptance ratio.

The method for drawing samples from the posterior PDF which we have explained above was first proposed by Metropolis and Hastings [83; 84] and hence is referred to as the Metropolis-Hastings algorithm for MCMC. In practice, specifying good jump proposals is extremely difficult and we do not discuss this topic here. For the purposes of using MCMC for compact binary parameter estimation, we use the implementation within the LAL library [85], known as lalinference.

3.3 Singular value decomposition of gravitational-wave template banks for CBC sources

The SVD which we will apply to gravitational waveforms for parameter estimation is adapted from a technique used for gravitational-wave detection [49; 50]. To detect gravitational-waves from binaries with non-spinning components, one filters the data with a large number of templates which span the space of the CBC chirp-mass and symmetric mass-ratio. The component masses of the gravitational-wave signal source are not known *a priori* and so one must scrutinize the continuous mass space of CBC signals to high fidelity. Templates in the space are distributed such that there is a fixed loss in SNR between nearest-neighbour templates which

is small. This is the so called “minimal match” between neighbouring templates, which is typically set to be greater than 97% [6; 7; 8; 49; 50] and corresponds to a loss in SNR of less than 3%. With such a template bank one might expect some correlation between the templates as the SNR between neighbouring templates differs by only 3% [49]. These correlations can be exploited by performing a SVD of a template bank which is used to generate an orthogonal set of basis templates and coefficients from which the original bank can be recovered. In particular, we will demonstrate that many of the basis vectors are redundant and hence the original bank can be recovered using a relatively small number of bases. Below we illustrate the SVD as applied to template waveforms. In Appendix B we provide supplementary material on the singular value decomposition, including an example of the SVD of a simple matrix.

The optimal filter for a matched filtered search is a gravitational-wave template, weighted by the inverse of the amplitude spectral density, a process known as “whitening”. We define the whitened frequency series as

$$\tilde{a}_w(f) = \frac{\tilde{a}(f)}{\sqrt{S_n(f)}}, \quad (3.13)$$

where $a(f)$ is the un-whitened frequency series. The study in this chapter deals exclusively with discretely sampled data and template waveforms. Because of this we find it useful to adopt the following notation for the data and waveforms.

We denote the discretely sampled whitened data-stream, and the α^{th} template waveform in a discrete template bank, as \vec{d} and \vec{h}_α respectively, where it is understood that these correspond to whitened versions of the raw, un-whitened data and template respectively. Both the data-stream and template are discretely sampled at a fixed sampling rate, Δt , in the time-domain. The elements of the data and waveforms at time $i \Delta t$ are d_i and $h_{\alpha i}$ respectively. We adopt the convention of [49] and define the SNR of the α^{th} template as the vector inner product

$$\rho_\alpha = \Re \vec{h}_\alpha \cdot \vec{d}, \quad (3.14)$$

where the definition of the vector inner product is given by Eq. (A.10). Note that this definition of the SNR implicitly assumes that the templates \vec{h}_α are normalized

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

to unit norm.

For non-spinning CBC sources the gravitational-wave phasing is controlled by the component masses of the binary. The template bank for a CBC search is generated by constructing a grid in the $\mathcal{M} - \eta$ plane such that for neighbouring templates indexed by α and α' we have $\vec{h}_\alpha \cdot \vec{h}_{\alpha'} > 0.95$. Instead of filtering the data using N real-valued filters we wish to remove the redundancy in the set of templates by reproducing ρ_α to the desired accuracy ρ'_α by linearly combining the output of a smaller set of real-valued filters, \vec{u}_μ . Hence, we want to truncate the number of unique filters such that [49]

$$\rho'_\alpha = \sum_{\mu=0}^{N'-1} A_{\alpha\mu} (\vec{u}_\mu \cdot \vec{d}) , \quad (3.15)$$

where \mathbf{A} is a real-valued “reconstruction matrix” and the total number of inner products required for the matched-filtered search is reduced from N to N' . To find the basis vectors \vec{u}_μ we compute the SVD of a real-valued $N \times L$ matrix \mathbf{H} , with $N \leq L$, whose columns are template waveforms:

$$\mathbf{H} = \{\vec{h}_0, \vec{h}_1, \dots, \vec{h}_{N-1}\} . \quad (3.16)$$

The SVD factors \mathbf{H} as [49]

$$H_{\nu j} = \sum_{\mu=0}^{N-1} v_{\nu\mu} \sigma_\mu u_{\mu j} , \quad (3.17)$$

where \mathbf{v} is a unitary matrix of reconstruction coefficients, $\vec{\sigma}$ is a vector of singular values and \mathbf{u} is a matrix of orthonormal basis vectors \vec{u}_μ . The matrix \mathbf{H} can be approximately reconstructed such that there are fewer basis vectors than templates in the original template bank. We define the approximate reconstruction of \mathbf{H} , \mathbf{H}' , as

$$H'_{\nu j} := \sum_{\mu=0}^{N'-1} v_{\nu\mu} \sigma_\mu u_{\mu j} , \quad (3.18)$$

where $N' < N$ which reduces the number of columns of \mathbf{u} . The reconstructed

SNR, Eq. (3.15), is then given by

$$\rho'_\alpha = \left(\vec{H}'_\alpha \right) \cdot \vec{d} \quad (3.19)$$

$$= \sum_{\mu=0}^{N'-1} (v_{\alpha\mu} \sigma_\mu) (\vec{u}_\mu \cdot \vec{d}) . \quad (3.20)$$

Comparing the above expression with Eq. (3.15), it follows that the SVD allows the α^{th} template waveform, \vec{h}_α , to be written as

$$\vec{h}_\alpha = \sum_{\mu=0}^{N'-1} M_{\mu\alpha} \vec{u}_\mu , \quad (3.21)$$

where we have defined the “projection coefficient” $M_{\nu\alpha}$ as

$$M_{\mu\alpha} := v_{\alpha\mu} \sigma_\mu . \quad (3.22)$$

The number of bases N' can be set by fixing the fractional loss in SNR which is incurred through truncating the sum in Eq. (3.18). The expected fractional loss in recovered SNR of *any* filter is given by [49]

$$\left\langle \frac{\delta\rho}{\rho} \right\rangle = \frac{1}{2N} \sum_{\mu=N'}^{N-1} \sigma_\mu^2 , \quad (3.23)$$

and the required number of filters N' is thus fixed by Eq. (3.23).

3.4 Interpolating template waveforms using the SVD

The accuracy requirements for gravitational-wave parameter estimation are typically higher than for searches. While a (dense) discrete template bank is sufficient for searches, parameter estimation requires that we are able to stochastically sample from a continuous parameter space in order to map the full likelihood-surface where the bulk of the posterior support is contained. Constructing a discrete tem-

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

plate bank is impractical for parameter estimation because one does not know *a priori* the shape of the surface. Conversely, for detection one is only interested in a single number, the maximum SNR, and the fractional loss in SNR of any stretch of data can be estimated from the error associated with the finite gridding of the bank.

We can construct a continuous representation of any template waveform in a given domain using the SVD of a template bank of gravitational waveforms. Recall that the SVD allows any template waveform to be written as a linear combination of basis waveforms and projection coefficients, c.f. Eq. (3.21). While this representation is discrete, the set of basis waveforms ought to be complete over the entire domain of the (discrete) template bank provided that neighbouring templates are sufficiently similar. It therefore follows that a template waveform $\vec{h}_{\alpha'}$, which is not in given template bank, should also be expressible as a sum of the same basis waveforms and its own projection coefficients. The projection coefficients can be found as follows. By the orthonormality of the bases, the inner product of $\vec{h}_{\alpha'}$ with a basis template \vec{u}_{μ} produces the μ^{th} projection coefficient associated with $\vec{h}_{\alpha'}$:

$$\vec{h}_{\alpha'} \cdot \vec{u}_{\mu} = M_{\mu\alpha'} . \quad (3.24)$$

This suggests that there is a continuous description of the projection coefficients across the domain of the template bank, and hence they ought to be amenable to interpolation. Because the template bank spans the $\mathcal{M} - \eta$ plane a single index α is somewhat cumbersome for enumerating the projection coefficients, which are function of two variables, \mathcal{M} and η . Without loss of generality we alter the notation of the projection coefficients and re-define the SVD of a template waveform as

$$\vec{h}(\mathcal{M}_k, \eta_l) = \sum_{\mu=0}^{N'-1} M_{\mu}(\mathcal{M}_k, \eta_l) \vec{u}_{\mu} , \quad (3.25)$$

where k and l enumerate points on a grid in the $\mathcal{M} - \eta$ plane. An “interpolated waveform”, \vec{h}' , can thus be constructed according to Eq. (3.25) from a linear combination of interpolated coefficients and basis waveforms:

$$\vec{h}'(\mathcal{M}, \eta) = \sum_{\mu} M'_{\mu}(\mathcal{M}, \eta) \vec{u}^{\mu} , \quad (3.26)$$

where $M'_\mu(\mathcal{M}, \eta)$ is an interpolated projection coefficient evaluated at an arbitrary point (\mathcal{M}, η) within the domain of the template bank. In order to interpolate the projection coefficients, we follow [50] and use Chebyshev polynomials of the first kind, although in principle other interpolation schemes could be used. The Chebyshev polynomials of the first kind are defined as [50]

$$T_J(x) = \frac{(x - \sqrt{x^2 - 1})^J + (x + \sqrt{x^2 - 1})^J}{2w}, \quad (3.27)$$

where $w = \sqrt{(1 + \delta_{J0})(J_{\max} + 1)}/2$ is a normalization constant and δ_{J0} is the Kroenecker delta. The polynomials $T_J(x)$ satisfy the discrete orthogonality condition

$$\sum_{j=0}^{J_{\max}} T_I(x_j) T_J(x_j) = \delta_{IJ}. \quad (3.28)$$

For interpolation, we use a net of points, scaled such that each dimension covers the interval $[-1, 1]$, located at the J_{\max}^{th} root of the Chebyshev polynomials, commonly referred to as interpolation nodes. This net of points is used in order to mitigate oscillatory behaviour at the boundary of the interpolation interval. For a single dimension, the nodes occur at the locations

$$x_j = \cos\left(\pi \frac{j + \frac{1}{2}}{J_{\max} + 1}\right). \quad (3.29)$$

In order to obtain the reconstruction coefficients at these locations, we project waveforms from these locations (scaled to the appropriate points in \mathcal{M} and η) onto the basis vectors. We generate waveforms at the (scaled) Chebyshev nodes using the following maps

$$\mathcal{M}[x_j] = \mathcal{M}_{\min} + \left(\frac{x_j + 1}{2}\right) \Delta\mathcal{M}, \quad (3.30)$$

$$\eta[y_l] = \eta_{\min} + \left(\frac{y_l + 1}{2}\right) \Delta\eta, \quad (3.31)$$

which project Chebyshev nodes on the interval $[-1, 1]$ onto the \mathcal{M} and η dimensions. Here \mathcal{M}_{\min} is the lower boundary in the \mathcal{M} -dimension and $\Delta\mathcal{M} =$

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

$\mathcal{M}_{\max} - \mathcal{M}_{\min}$ where \mathcal{M}_{\max} is the upper boundary of the \mathcal{M} -dimension, and vice versa for the η -dimension. From the values on the net, we interpolate to other positions in the $\mathcal{M} - \eta$ plane using 2D-Chebyshev interpolation for each set of reconstruction coefficients $M_\mu(\mathcal{M}, \eta)$. These reconstruction coefficients are projected onto the Chebyshev polynomials [50]

$$C_{KL\mu} = \sum_{k=0}^{K_{\max}} \sum_{l=0}^{L_{\max}} T_K(x_k) T_L(y_l) M_\mu(\mathcal{M}[x_k], \eta[y_l]) , \quad (3.32)$$

which provide weights for the 2D-Chebyshev interpolation. The values of K_{\max} and L_{\max} correspond to the number of templates in the \mathcal{M} and η directions respectively. The interpolated reconstruction coefficients are then given by a combination of the weights together with the Chebyshev polynomials [50]

$$M'_\mu(\mathcal{M}, \eta) = \sum_{K=0}^{K_{\max}} \sum_{L=0}^{L_{\max}} C_{KL\mu} T_K(x[\mathcal{M}]) T_L(y[\eta]) , \quad (3.33)$$

where the inverse maps $x[\mathcal{M}]$ and $y[\eta]$ project points on \mathcal{M} and η onto the interval $[-1, 1]$. The inverse maps are given by

$$x[\mathcal{M}] = \left(\frac{\mathcal{M} - \mathcal{M}_{\min}}{\Delta\mathcal{M}} \right) - 1 , \quad (3.34)$$

$$y[\eta] = \left(\frac{\eta - \eta_{\min}}{\Delta\eta} \right) - 1 . \quad (3.35)$$

Below we illustrate the application of the SVD to parameter estimation.

3.5 Parameter estimation using interpolated waveforms

We will compare the marginalized one-dimensional PDFs, Eq. (3.4), obtained by using an interpolated template waveform family to those generated using the standard, non-interpolated template waveform family for the same data set. For illustration we consider a toy example with five free parameters. In addition,

we consider data from a single detector, as opposed to a network, for simplicity. We generate a single-detector data set \vec{d} containing a signal waveform $h(\vec{\theta})$ and Gaussian stationary noise n with a power spectral density typical of initial LIGO [1]. By only having five free parameters we effectively set the prior on the other four to be delta functions centered on the signal values. We chose to fix the sky position (α, δ) and the inclination and polarization angle (ι, φ) of the template waveforms such that they are not searched over.

Because interpolation is carried out in the mass space only, we focus on exploring the effects of interpolation on mass parameters and parameters that are known to be very strongly correlated with masses (time and phase of coalescence and distance). If the accuracy of the recovery of these parameters is unaffected by interpolation, we do not expect the angular parameters to be affected, either. However, it is important to realize that the absolute accuracy with which some parameters, particularly distance, are estimated is improved by fixing sky location and orientation parameters and lifting corresponding degeneracies. Thus, the measurement uncertainties inferred below should not be considered typical. Since we demand that systematic biases from using interpolated templates are smaller than statistical measurement uncertainties, the improvement in the accuracy of distance measurement means that we are being conservative in evaluating the quality of SVD-interpolated parameter estimation.

The signal contained in the data set has source parameters $(\mathcal{M}, \eta, D, t_c, \phi_c) = (7.45 M_\odot, 0.247, 33 \text{ Mpc}, 0 \text{ s}, 2.16)$ and we have omitted the others which are not searched over. The signal has a signal-to-noise ratio $\text{SNR} = 14.8$. The frequency-domain data is sampled at $\Delta f = 1/32 \text{ Hz}$ with a maximum frequency of 512 Hz.

The choice of prior distribution is an important factor in Bayesian inference. The prior incorporates the ranges of the model parameters, and the *a priori* probability distribution on those parameters before the data are analysed. We use a uniform prior on $\log D$ and η with ranges $D \in [1 \text{ Mpc}, 100 \text{ Mpc}]$ as we have no *a priori* preferred distance to the source. We similarly use uniform priors on ϕ_c and t_c over the range $0 \leq \phi_0 \leq 2\pi$ and $-0.1 \text{ s} \leq t_c \leq 0.1 \text{ s}$, respectively.

We use a prior on \mathcal{M} of the form $\mathcal{P}(\mathcal{M}|\mathcal{H}) \propto \mathcal{M}^{-11/6}$ in the range $\mathcal{M} \in [7.20 M_\odot, 7.60 M_\odot]$ and a uniform prior on η in the range $\eta \in [0.175, 0.250]$. The

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

prior on \mathcal{M} corresponds to the Jeffreys prior for the waveform family described by Eq. (3.5) [12]. The Jeffreys prior is used to assign equal prior volume to points in the parameter space, which is necessary when dealing with a space which has non-trivial curvature.

For the mock data set we ran a MCMC comprising five parallel Markov chains in order to compute the PDF $p(\vec{\theta} = (\mathcal{M}, \eta, D, t_c, \phi_c)|d, \mathcal{H})$. The limits of integration of the likelihood function, Eq. (3.2), are fixed to $f_{\min} = 40$ Hz, $f_{\max} = 512$ Hz. To extract the posterior samples from the raw MCMC output we discard the first 10,000 samples as burn-in.

We measure the convergence of the parallel chains using the Gelman-Rubin R-statistic [86]. For well converged chains this should be close to $R = 1$ and we regard the MCMC to be complete once $R \leq 1.001$ for all parameters.

SVD Setup

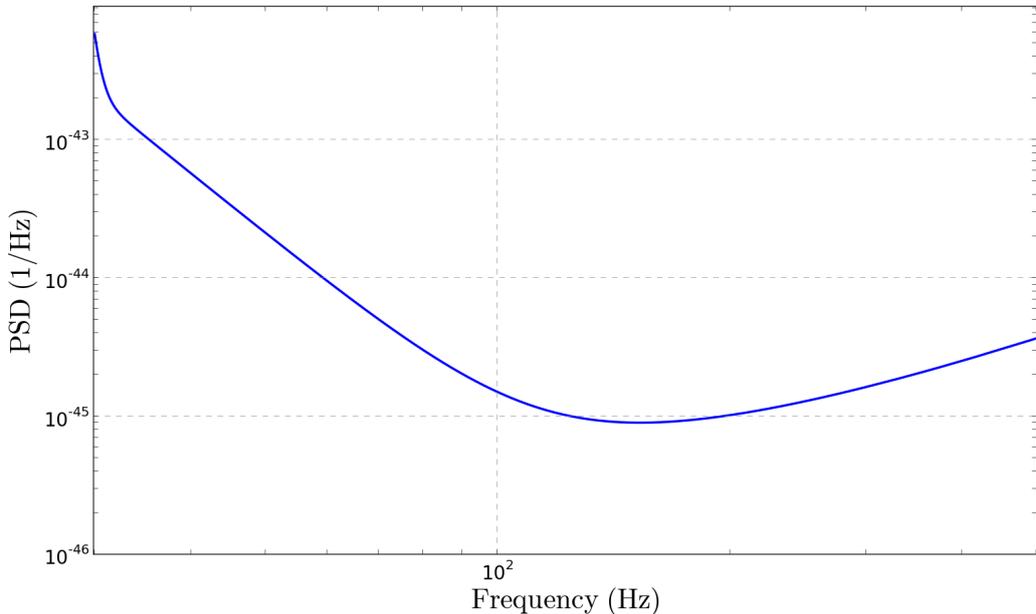


Figure 3.2: Noise PSD used to whiten template waveforms, Eq. (3.13), roughly matching that of initial LIGO [1].

The input to the SVD is a set of discretely sampled, whitened time-domain waveforms [50]. The frequency-domain SPA waveforms are first whitened in the frequency domain with the ASD and transformed into the time-domain (see Eq. (A.4) for our definition of the discrete Fourier transform and its inverse). We use a noise PSD roughly matching that of initial LIGO [1] shown in Fig. 3.2. To improve the interpolation we chose to align the waveforms such that they all go through frequency at which the noise PSD is a minimum at the same time. This is achieved by shifting the waveforms by a time t_{shift} in the frequency domain via the transformation

$$\tilde{h}(f) \Rightarrow \tilde{h}(f)e^{-2\pi if t_{\text{shift}}} . \quad (3.36)$$

The value of t_{shift} is computed from the time it takes for the binary to reach coalescence from the frequency at which the PSD is a minimum. This time is known as the “chirp time” and for the SPA waveform family it is given in [87]. The frequency at which the noise PSD is a minimum is $f_{\text{min}} = 152.84$ Hz.

By carrying out the interpolation in the time domain, we show that the technique can be applied to time-domain waveform families. Time-domain waveforms are typically computationally expensive for parameter estimation (see Sec. 3.6), so this approach allows us to assess the computational savings associated with interpolating them. It is also consistent with the work in [49; 50], where time-domain waveforms generated by Fourier transforming frequency-domain waveforms were interpolated.

We ensure that all templates are of the same length, equal to the next highest power of two of the longest time-domain waveform in the set, which in our case is 2 s. For the proof of principle we apply the SVD to a small patch in $\mathcal{M} - \eta$ space bounding the signal parameters. This region is set by the prior range on \mathcal{M} and η described above, chosen to be broad enough so that, for our data set, there is no likelihood support near the boundaries.

The number of computations for the SVD of a $N \times L$ matrix with $N \leq L$ scales like $\mathcal{O}(LN^2)$. For the purposes of constructing the SVD we have found it efficient to split the $\mathcal{M} - \eta$ space into four equally sized patches, with a separate SVD applied to each patch. The waveforms in each bank lie on an evenly spaced $(15 + 1) \times (15 + 1)$ grid. We add one to the grid in each dimension because the

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

interpolation can be unpredictable at the boundaries of the space. The grid is chosen by requiring that the normalized inner product between non-interpolated waveforms and interpolated waveforms generated within each patch is at least 99.9%. Such normalized mismatches of < 0.001 between interpolated and non-interpolated waveforms should ensure that parameter-estimation accuracy is not compromised as long as the signal-to-noise ratio does not exceed ~ 20 (so that twice the mismatch times the square of the SNR is less than unity [88; 89]), although parameter estimation could remain accurate at even higher SNRs. The patching is shown in Fig. 3.3. Each waveform in the template bank corresponds to a face-on, overhead binary generated at a fiducial distance of 1 Mpc. We choose to truncate the SVD so that the norm of any reconstructed template is conserved to a level of $\sim 10^{-5}$ (c.f. Eq. (3.23)). The truncated-SVD of the template bank in each patch uses 20 basis waveforms.

One subtlety of the implementation is that the interpolated waveforms are whitened time-domain filters. The likelihood function Eq. (3.2) is computed using un-whitened frequency series which in our example correspond to the Fourier transform of a time series. To recover the appropriate frequency series we first have to Fourier transform the interpolated waveform back into the frequency domain.

We further note that we do not use waveforms normalized to unit norm as input to the SVD, as was done in [49; 50]. Over large areas of the mass-space it is necessary to normalize the waveforms because there may be significant differences in the power of waveforms across the space and so the SVD can disproportionately weight waveforms with more power if they are not normalized to unit norm. We are not affected by this issue because all the waveforms in our example have roughly the same power due to the limited extent of the template bank in mass space. In general one would use waveforms normalized to unit norm for SVD over a larger parameter space; however, to recover the overall amplitude of the interpolated waveforms the normalization coefficients would themselves have to be interpolated.

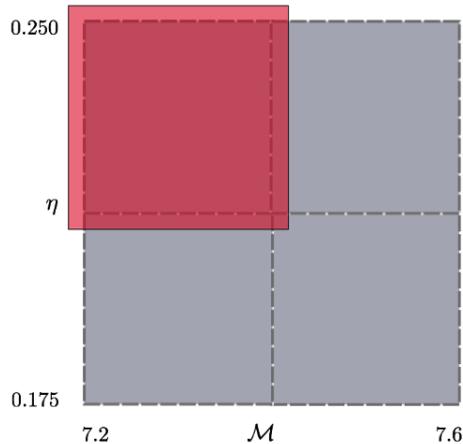


Figure 3.3: Patching of the parameter space. The space is patched into four regions. The SVD is applied to each patch (shown in red) which is padded at the boundary by one waveform and so contains a $(15+1)\times(15+1)$ grid of template waveforms. The SVD of the template banks consists of 20 basis templates in each patch.

3.6 Results: Comparison of parameter estimates using interpolated and non-interpolated waveforms

The marginalized PDFs for complete MCMC runs using non-interpolated and interpolated waveforms are shown in Fig. 3.4. We have omitted the marginalized one-dimensional PDF of the coalescence phase ϕ_c as it of little physical interest. Each run required around 1.5×10^6 waveform computations. The mean posterior values of the distributions along with the signal values are shown in Table 3.1.

The chirp-mass distribution computed using interpolated waveforms is clearly biased. This is corroborated by a two-sample KolmogorovSmirnov (K-S) test [90] which reveals that the two sets of samples are not consistent with arising from the same distribution with overwhelming odds. Nevertheless, the systematic bias in the mean posterior parameter estimate of chirp mass is a factor of four smaller than the statistical measurement uncertainty. Thus, we pass a commonly-used threshold for sufficient waveform-model accuracy [e.g., 89]. We note that the accuracy could be increased by, for example, using a higher density template bank

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

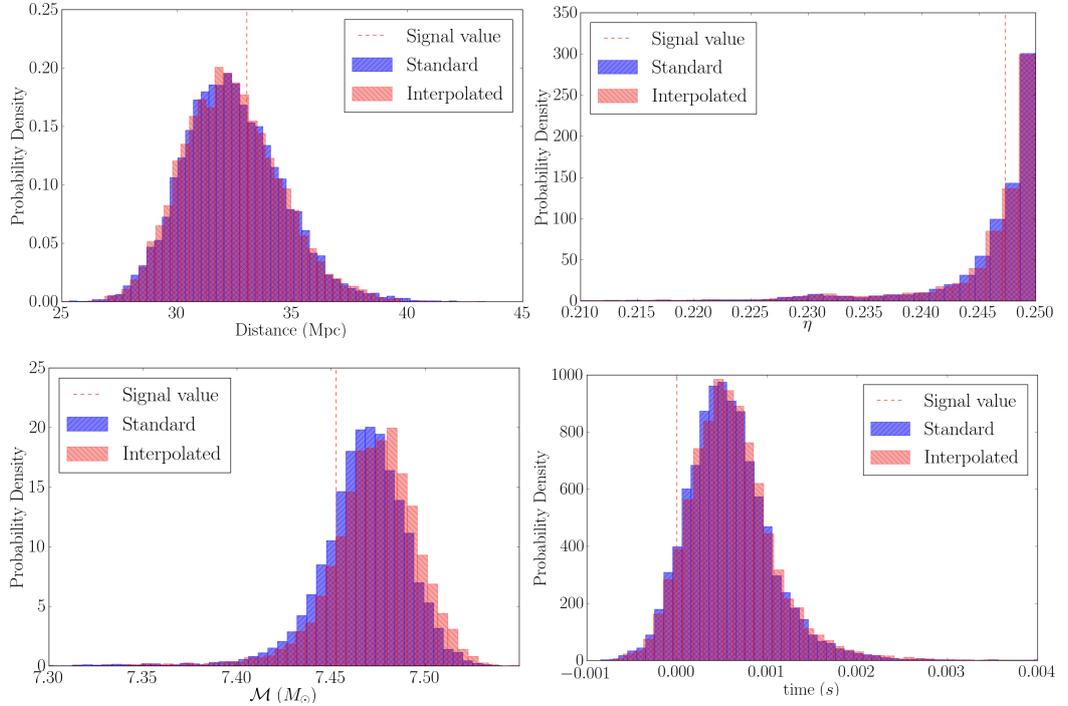


Figure 3.4: Marginalized PDFs, Eq. (3.4), produced using non-interpolated SPA waveforms (blue or dark grey) and interpolated SPA waveforms (red or light grey). The signal value is shown as the dashed red vertical line.

Param	Mean posterior value (interpolated SPA)	Mean posterior value (SPA)	Signal Value
$\mathcal{M} (M_{\odot})$	7.472 (2.5×10^{-2})	7.467 (2.5×10^{-2})	7.450
η	0.2457 (7.1×10^{-3})	0.2457 (7.2×10^{-3})	0.2473
D (Mpc)	32.39 (2.11)	32.40 (2.13)	33.00
t_c (s)	1.0×10^{-3} (4×10^{-4})	1.0×10^{-3} (4×10^{-4})	0

Table 3.1: Mean posterior parameter estimates (and standard deviations) of the marginalized PDFs using interpolated and non-interpolated SPA waveforms (Fig. 3.4).

or using waveforms normalized to unit norm as input to the SVD. In general the required accuracy can be estimated from the detection trigger SNR [89].

The two-sample K-S test marginally fails for the coalescence-time distribution, but there is no evidence of a systematic bias on the scale of statistical measurement errors. We find that the sets of posterior samples for the other two PDFs in Fig. 3.4, symmetric mass ratio and distance, are consistent with arising from same distribution as quantified by the K-S test.

Computational cost of template waveform generation

Two commonly-used time-domain waveform families which are relevant for parameter estimation are the inspiral-only post-Newtonian approximant TaylorT4 [26] and the effective-one-body family calibrated to numerical relativity (EOBNR, [26]) that includes inspiral, merger, and ringdown phases. The latter are typically more computationally intensive.

For technical reasons, our comparison uses interpolated waveforms which are based on the SPA approximation rather than time-domain waveform families. This is inconsequential because we Fourier transformed the SPA waveforms into the time-domain for the SVD procedure. We have observed that applying the truncated SVD to TaylorT4 time-domain templates in the mass space used to construct the SVD in Sec. 3.5 yields the same number of basis vectors as when applying it to inverse-FFT'd SPA templates. The computational cost of interpolation will therefore be identical.

Our measure to compare the computational costs of interpolated TaylorT4 and EOBNR waveforms is the time it takes to compute a single interpolated waveform. While this does not estimate the theoretical minimum number of FLOPs of the process, and is also hardware dependent, it does provide a useful heuristic for comparing the relative speed of each waveform family. Recall that the interpolated waveforms are a time-domain approximant and hence the comparison is to determine the computational savings for time-domain waveforms. We restrict our comparison to waveforms generated in the mass space in Fig. 3.3. The length of TaylorT4 and EOBNR waveforms will in general depend on the specific source masses. For a fair comparison we compare waveforms which have

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

approximately the same number of data points. Because EOBNR must be generated at a sampling rate of 4096 Hz, we ensure that the interpolated and TaylorT4 waveforms are sampled at this frequency. Each waveform is approximately 2 s in duration.

The results of the comparison are shown in Table 3.2. For reference we also show the computational time of standard SPA waveforms. We find that on average, the interpolated waveforms are ten times faster to generate than TaylorT4 and fifteen times faster than EOBNR, a significant increase in computational efficiency. However, for the waveform parameters considered here, inspiral-only waveforms could be generated at lower sampling rates than the 4096 Hz required for EOBNR waveforms; therefore, the cost of constructing interpolated or TaylorT4 waveforms can be around four times smaller relative to EOBNR than the values quoted in Table 3.2.

Waveform Family	Computational Time (T)
SPA	0.2
Interpolated	1
TaylorT4	10
EOBNR	15

Table 3.2: Computational time of template waveform generation in units of computational time of interpolated waveforms, T . EOBNR, TaylorT4 and interpolated waveform families are generated at a sampling rate of 4096 Hz and have a duration of 2 s. The interpolated waveforms consist of 20 pre-computed basis vectors. SPA waveforms are generated in the frequency domain; to ensure the SPA waveforms contain the same number of sample points they are generated at a sampling frequency $\Delta f = 1/2$ Hz and have a maximum frequency of 2048 Hz.

We also estimate the cost of pre-computing the SVD interpolation. We have previously noted that the computational cost of an SVD of an $N \times L$ matrix with ($N \leq L$) scales like $\mathcal{O}(N^2L)$. One also needs to compute the $N \times L$ matrix of template waveforms as input to the SVD. The cost of computing a waveform of length L is typically $\mathcal{O}(L)$, possibly with a very large pre-factor. Thus, the total cost of pre-computing interpolation coefficients will be less than $\mathcal{O}(N^2)$ times the cost of an individual waveform computation. For instance, in our example, $N = 16 \times 16 = 256$, so interpolation can reduce overall MCMC costs for any

time-domain waveform templates by an order of magnitude or more when the typical MCMC chain length of $\gtrsim 10^6$ samples is taken into account.

3.7 Discussion and conclusion

We have provided a proof of principle that interpolated waveforms can be used for parameter estimation without unacceptable loss of accuracy. Our example was restricted to a five-dimensional search over the source chirp mass \mathcal{M} and symmetric mass ratio η , the distance to the source D and the time and phase at coalescence t_c and ϕ_c . We further restricted the prior ranges on \mathcal{M} and η to $\mathcal{M} \in [7.20 M_\odot, 7.60 M_\odot]$ and $\eta \in [0.175, 0.250]$, respectively. The systematic biases observed when using interpolated waveforms are demonstrated to be smaller than statistical measurement uncertainties. Thus, SVD-interpolated waveforms satisfy the stringent waveform-model accuracy criteria imposed by parameter-estimation requirements.

The relative computational times of generating interpolated waveforms and time-domain TaylorT4 and EOBNR waveforms are shown in Table 3.2. Interpolated waveforms can be generated at around an order of magnitude more cheaply than TaylorT4 or EOBNR. This suggests that the computational cost of parameter estimation can be significantly reduced by employing SVD-interpolated waveforms for likelihood computation when the latter is dominated by the cost of waveform generation.

In order for interpolated templates to be viable for parameter estimation pipelines we need to apply the SVD-interpolation technique to a significantly larger region of the CBC mass space than in the example considered here. Searches of gravitational waves from low-mass systems look for binaries with a maximum total mass of $25 M_\odot$ and a minimum component mass of $1 M_\odot$ [7] and high mass searches target binaries with total masses between $25 M_\odot$ and $100 M_\odot$ [8]. To be able to apply our parameter estimation technique to triggers from such searches in a single step, without first determining the more limited mass region where there is significant likelihood support, we will need to efficiently patch the parameter space over a large mass range so that the computational cost of generating the SVD can be minimized. A necessary condition for setting up the SVD in all

3. TOWARDS RAPID COMPACT BINARY PARAMETER ESTIMATION USING INTERPOLATED WAVEFORMS

patches is that its computational cost, plus the cost of generating interpolated waveforms, must be less than the cost of performing the parameter estimation with non-interpolated waveforms. This will be the subject of future work.

For simplicity, we have considered analysing a single-detector data set. If one were to extend the use of interpolated waveforms to analysing data from a network of interferometers, one would require independent sets of interpolated waveforms for each data set of the network. This is because our technique relies on a SVD of template waveforms, whitened by the detector noise PSD. Because each detector in general has a different PSD, one would thus require separate sets of whitened filters, each whitened with the appropriate noise PSD. For a typical detector network of three interferometers, one would thus need to generate three template banks and perform three SVDs. While this increases the computational cost, we note that computing independent sets of template banks and SVDs is highly parallelizable, and hence the wall-time of computing three banks could easily be made the same as computing a single template bank.

Furthermore, we have to be able to extend the SVD to generic waveform families. Particularly interesting is the potential to extend the technique to EOBNR waveforms, which are currently expensive to generate, and waveform families which describe binaries with spinning components. The latter class of waveforms have an intrinsic parameter space with up to six more independent parameters (two spin vectors) and the current technique of interpolation within the intrinsic parameter space of waveforms may become costly in large-dimensional spaces. However, it is interesting to consider the potential to apply the technique to spin-aligned/anti-aligned waveforms [e.g., 87; 91] as this class of waveforms have only one extra parameter, the reduced spin of the binary. The analysis of data from advanced LIGO and Virgo detectors, which may have lower-frequency cutoffs close to 10 Hz [58], will require template waveforms that are several minutes in duration. This technique is likely to be highly relevant to parameter estimation in that context.

Chapter 4

PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

4.1 Introduction

Techniques to improve the efficiency of Bayesian parameter estimation are crucial in the run-up to the advanced-detector era. Interpolated gravitational waveforms based on the singular value decomposition have the potential to significantly increase the efficiency when parameter estimation is dominated by waveform computation.

For interpolated waveforms to be viable for parameter estimation on real data, a number of extensions of the technique demonstrated in Chapter III are required. While it has been demonstrated that interpolated waveforms can be used for parameter estimation without incurring significant bias in parameter estimates, the proof-of-principle study was limited to one waveform family and a small patch of parameter space. We begin this chapter by investigating the extension to generic waveform families which describe CBC sources with non-spinning components. In Chapter III we interpolated a waveform family constructed by Fourier transform-

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

ing frequency-domain SPA waveforms into the time-domain, which we used as templates for parameter estimation. We demonstrated that because the interpolation technique is the same regardless of the waveform family being interpolated one could expect to reduce the computational cost of generating time-domain waveforms, e.g. TaylorT4 and EOBNR [26], via interpolation.

It is occasionally useful for parameter estimation to utilize time-domain waveforms. For example, the NINJA project [92], and its sequel (NINJA2 [93]) currently underway, is an exercise in detecting and performing parameter estimation on (time-domain) waveforms generated within the framework of numerical relativity. Here EOBNR waveforms should be useful as a template waveform family because they have been matched to similar numerical waveforms. One of the reasons we did not initially apply the interpolation technique in Chapter III *directly* to time-domain waveforms was due to a number of potential issues which could complicate the interpolation. For example, because the waveforms are whitened in the frequency domain one first has to Fourier transform a time-domain waveform into the frequency domain to perform the whitening process. This will introduce artifacts into the Fourier series which can compromise the accuracy of the SVD, and hence interpolated waveforms. In addition, EOBNR waveforms may be difficult to interpolate because of the relative time-scales of the inspiral phase and the ringdown.

We first demonstrate that parameter estimation can be performed using interpolated TaylorT4 waveforms, which are inspiral-only time-domain waveforms. We find that while there are clear biases in the recovered posterior probability density functions, the bias of the mean posterior parameter estimates are much less than the scale of the statistical uncertainty and so we pass a common threshold for model waveform accuracy [89]. Furthermore, the improvement in computational time of generating interpolated TaylorT4 waveforms is around a factor of 10 with respect to non-interpolated TaylorT4 waveforms. In general, we find that applying the SVD to time-domain waveforms which describe the full inspiral, merger and ringdown phases of coalescence is challenging and the quality of interpolated EOBNR waveforms is too low for parameter estimation, but could suffice for gravitational-wave searches.

4.2 Interpolating inspiral-only time-domain waveforms

To demonstrate the subtleties of interpolating time-domain waveforms we first consider a simple inspiral-only time-domain waveform family known as TaylorT4. We refer the reader to Sec. 2.3.3 for details on the construction of the TaylorT4 waveform family. Spectral leakage is a general feature of Fourier transforming a time-series which has discontinuities from a hard beginning and ending. (Such a time-series can be thought of as the convolution of an infinite, smooth time-series with a box window). Thus, the Fourier transform of the convoluted time-series will contain additional high- and low-frequency components from the box window, which is a linear superposition of oscillating functions. This can be ameliorated by generating the time series at a lower sampling frequency which can move many of the artifacts to a region in which one is not interested.

There is also another problem associated with using time-domain waveforms which is symptomatic of our implementation. Inspiral-only time-domain waveforms are integrated from an initial frequency up to the ISCO frequency. Time-domain waveforms generated within the LIGO Scientific Collaboration Algorithm Library (LAL) are aligned such that the phase at last time-sample of different waveforms is the same. However, because the time series are discretely sampled, the frequency at the last time sample is in general not the same and is a reflection of inaccurate time-frequency evolution of the waveforms. This may be an issue for interpolating time-domain waveforms because the waveforms have to be aligned in frequency domain with the minimum of the noise PSD and so there may be inconsistencies in the alignment of the waveforms for the SVD. This can be ameliorated by noting that in the limit that the sampling frequency is infinitely fine, the frequency at the last sample will tend to the same value. Hence at higher sampling frequencies we may expect the time-frequency evolution of time-domain waveforms to be more accurate.

To study the effects of discrete sampling on interpolated waveforms, we compute the SVD of the waveforms over a range of sampling rates. As a measure of the quality of the interpolated waveforms we compute the inner products between interpolated and non-interpolated TaylorT4 waveforms on a uniform grid.

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

We chose to work with the “mismatch” between (normalized) interpolated and non-interpolated waveforms which is simply

$$\text{MM} = 1 - (h_{\text{interp}}|h_{\text{non-interp}}), \quad (4.1)$$

and is zero when $h_{\text{interp}} = h_{\text{non-interp}}$.

Recall that the input to the SVD is a set of whitened waveforms, in this case waveforms described by the TaylorT4 family. The waveforms are first whitened in the frequency domain with the square root of the noise PSD and transformed back into the time-domain for interpolation. We use the same noise PSD as in Chapter III which is typical of initial LIGO [1] and is shown in Fig. 3.2. The waveforms are generated from a starting frequency of 30 Hz in the time-domain which we use to mitigate the effects of spectral leakage in the Fourier-domain waveform when whitening, and are truncated at a low frequency cut-off of 40 Hz before back transforming into the time-domain for the SVD. To improve the interpolation we chose to align the waveforms such that they all go through frequency at which the noise PSD is a minimum at the same time. The frequency at which the noise PSD is a minimum is $f_{\text{min}} = 152.844$ Hz. This is achieved by time-shifting the waveforms in the frequency domain by an amount $-t_{\text{shift}}$ which can be calculated from the time it takes for the binary to reach coalescence from f_{min} . This time is calculated by computing Eq. (A.14) to find the frequency evolution of the binary as a function of time.

To compute the SVD, we work within the region of parameter space used in Chapter 3, which is defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}, 0.175 \leq \eta \leq 0.250$. We divide the parameter space into four equally sized patches as in Fig. 3.3. In each patch we construct a template bank of $(25 + 1) \times (25 + 1)$ whitened waveforms on a uniform grid. In Fig. 4.1 we histogram the set of mismatches between interpolated and non-interpolated waveforms, for waveforms generated at three different sampling rates: 512 Hz, 1024 Hz, 2048 Hz and 4096 Hz. The SVD in each case uses 20 basis templates per patch of parameter space.

To compute the mismatch between interpolated and non-interpolated TaylorT4 waveforms we generate the waveforms on a uniform 50×50 grid in the region of parameter space defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}, 0.175 \leq \eta \leq 0.250$.

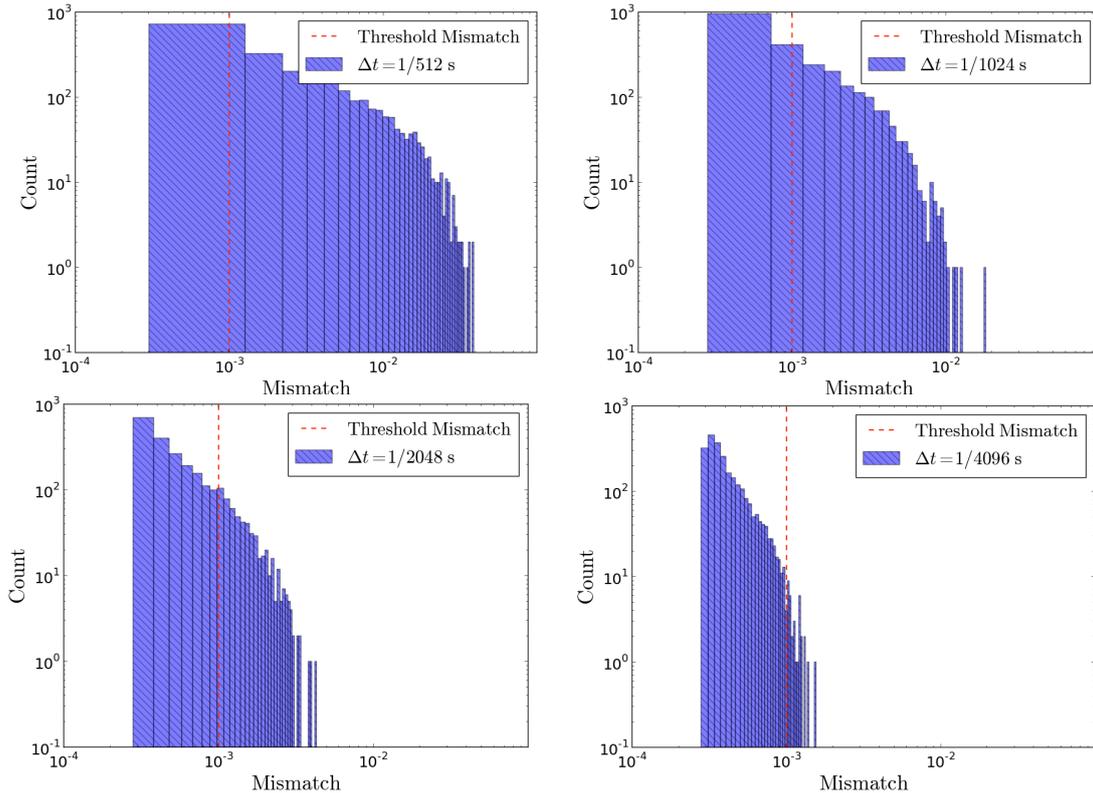


Figure 4.1: Histogram of mismatches between (normalized) interpolated and non-interpolated TaylorT4 waveforms generated on a uniform (50×50) grid in a patch of (\mathcal{M}, η) defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}$, $0.175 \leq \eta \leq 0.250$. Each panel corresponds to waveforms generated at different sampling rates: 512 Hz, 1024 Hz, 2048 Hz and 4096 Hz. We set a threshold mismatch of 10^{-3} between interpolated and non-interpolated waveforms, shown as a red vertical line, above which we deem interpolated waveforms to be insufficiently accurate for parameter estimation. The best case occurs for a sampling rate of 4096 Hz for which the majority of mismatches are less than the threshold.

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

We set a threshold mismatch of 10^{-3} above which the accuracy of interpolated waveforms may seriously compromise parameter estimation accuracy [89]. We find that the most accurate interpolated waveforms are those generated at a sampling frequency of 4096 Hz for which the vast majority of waveforms are more accurate than the threshold value. We note that there are around ~ 20 waveforms which are less accurate than the threshold, however this is unlikely to have a large impact over the bulk of the parameter space. We further note that increasing the number of templates in the template bank on which we perform the SVD has almost no effect in improving the quality of interpolated templates. Below we demonstrate the use of interpolated TaylorT4 waveforms as templates for parameter estimation.

4.2.1 Parameter estimation using interpolated TaylorT4 waveforms: 5-dimensional case

Our goal is to estimate the parameters of a gravitational-wave signal contained in Gaussian stationary noise using both the interpolated and non-interpolated TaylorT4 waveform family as templates. The gravitational-wave signal contained in the data is described by the non-interpolated TaylorT4 waveform family. We generate waveforms in the time-domain at a sampling frequency of $\Delta t = 1/4096$ s. The signal contained in the data set has source parameters $(\mathcal{M}, \eta, D, t_c, \phi_c) = (7.45 M_\odot, 0.247, 33 \text{ Mpc}, 0 \text{ s}, 2.16)$ and we have omitted the others which are not searched over. The signal has a signal-to-noise ratio $\text{SNR} = 14.8$. We use the same prior distributions as in Sec. 3.5. To recap, the prior distributions are set as follows.

4.2.1.1 Prior distributions

We use a uniform prior on $\log D$ and η with ranges $D \in [1 \text{ Mpc}, 100 \text{ Mpc}]$ as we have no *a priori* preferred distance to the source. We similarly use uniform priors on ϕ_c and t_c over the range $0 \leq \phi_c \leq 2\pi$ and $-0.1 \text{ s} \leq t_c \leq 0.1 \text{ s}$, respectively.

We use a prior on \mathcal{M} of the form $\mathcal{P}(\mathcal{M}|\mathcal{H}) \propto \mathcal{M}^{-11/6}$ in the range $\mathcal{M} \in [7.20 M_\odot, 7.60 M_\odot]$ and a uniform prior on η in the range $\eta \in [0.175, 0.250]$. The

prior on \mathcal{M} corresponds to the Jeffreys prior for the waveform family described by Eq. (3.5) [12].

For the mock data set we ran a MCMC comprising five parallel Markov chains in order to compute the PDF $p(\vec{\theta} = (\mathcal{M}, \eta, D, t_c, \phi_c)|d, \mathcal{H})$. The limits of integration of the likelihood function, Eq. (3.2), are fixed to $f_{\min} = 40$ Hz, $f_{\max} = 2048$ Hz. The upper limit of integration is the nyquist frequency of waveforms sampled in the time-domain at $\Delta t = 1/4096$ s. To extract the posterior samples from the raw MCMC output we discard the first 10,000 samples as burn-in.

We measure the convergence of the parallel chains using the Gelman-Rubin R-statistic [86]. For well converged chains this should be close to $R = 1$ and we regard the MCMC to be complete once $R \leq 1.001$ for all parameters.

4.2.1.2 SVD set up

To generate interpolated template waveforms we take the SVD of a template bank of whitened TaylorT4 waveforms over a region of parameter space defined by our prior range in \mathcal{M} and η : $\mathcal{M} \in [7.20 M_{\odot}, 7.60 M_{\odot}]$, $\eta \in [0.175, 0.250]$. We patch the parameter space into four equally sized regions as in Fig. 3.3 in which we generate a separate template bank. The waveforms in the template banks are generated at a sampling frequency of $\Delta t = 1/4096$ s. Each template bank uses $(25 + 1) \times (25 + 1)$ waveforms.

4.2.1.3 Results

The marginalized PDFs for complete MCMC runs using non-interpolated and interpolated TaylorT4 waveforms are shown in Fig. 4.2. We have again omitted the marginalized one-dimensional PDF of the coalescence phase ϕ_c as it of little physical interest. Each run required around 1×10^6 waveform computations to be fully converged. The mean posterior values of the distributions along with the signal values are shown in Table 4.1.

As with interpolated stationary phase approximation (SPA) waveforms, the chirp-mass distribution computed using interpolated TaylorT4 waveforms is clearly biased. However it is encouraging that the systematic bias in the mean posterior chirp mass estimate is around a factor of forty smaller than the statistical mea-

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

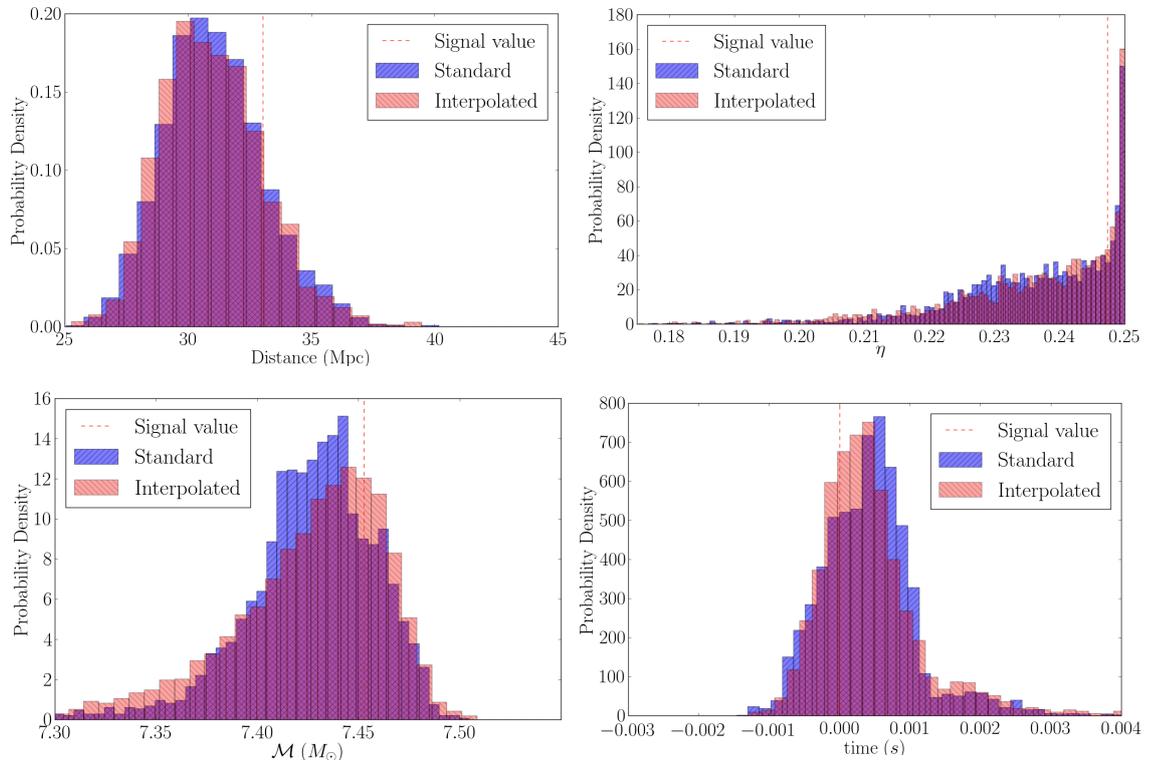


Figure 4.2: Marginalized PDFs, Eq. (3.4), produced using non-interpolated TaylorT4 waveforms (blue or dark grey) and interpolated TaylorT4 waveforms (red or light grey). The signal value is shown as the dashed red vertical line.

Param	Mean posterior value (interpolated TaylorT4)	Mean posterior value (TaylorT4)	Signal Value
$\mathcal{M} (M_{\odot})$	7.424 (3.7×10^{-2})	7.425 (3.3×10^{-2})	7.450
η	0.236 (1.3×10^{-2})	0.236 (1.2×10^{-2})	0.247
D (Mpc)	31.03 (2.4)	31.09 (2.6)	33.00
t_c (s)	1.0×10^{-3} (4×10^{-4})	1.0×10^{-3} (4×10^{-4})	0

Table 4.1: Mean posterior value parameter estimates (and standard deviations) of the marginalized PDFs using interpolated and non-interpolated TaylorT4 waveforms (Fig. 4.2).

surement uncertainty. Again we pass a commonly-used threshold for sufficient waveform-model accuracy [e.g., 89] and at a higher threshold than with interpolated SPA waveforms. In addition, the time distribution is also clearly biased, however as with the chirp mass, there is no bias of the mean posterior time parameter estimate on the scale of statistical measurements errors. This is also the case with the other two parameters, distance and symmetric mass ratio. We find that none of the sets of posterior samples for the PDFs in Fig. 3.4 are consistent with arising from same distribution as quantified by the K-S test.

We find that the reduction in computational time in using interpolated TaylorT4 waveforms is around a factor of 10, which is consistent with Table 3.2.

4.3 Interpolating inspiral-merger-ringdown waveforms

We now extend the interpolation technique to waveforms which describe the full inspiral, merger and ringdown phases of coalescence. We take as our model time-domain waveform EOBNR which was used in Sec. 2.3 and is described in detail in [26]. Unlike TaylorT4 the EOBNR family of waveforms must be generated

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

at a *minimum* sampling rate of 4096Hz in order to accommodate the ringdown phase. We use the implementation provided by the LIGO Scientific Collaboration Algorithm Library (LAL) that corresponds to the approximant EOBNRv2 [41].

An clear subtlety in applying the SVD to EOBNR is the alignment of the waveforms. Unlike SPA or TaylorT4 waveforms, there is no well defined end of EOBNR waveforms. Because of this we chose to align the peak amplitude of the waveforms with the minimum of the noise PSD in the frequency-domain. We use the same noise PSD as in Chapter III which is typical of initial LIGO [1] and is shown in Fig. 3.2. In order to perform the alignment we first have to compute the time-frequency evolution associated with EOBNR waveforms. This is computed within the waveform generation routine within LAL for the approximant EOBNRv2 [41]. We again take the reference frequency to be $f_{\min} = 152.84375$ Hz, which is the frequency at which the PSD is a minimum.

To determine the quality of interpolated EOBNR waveforms we compute the mismatch, Eq. (4.1), between interpolated and non-interpolated EOBNR waveforms. We apply the SVD to a set of whitened EOBNR waveforms generated in the same region of parameter space shown in Fig. 3.3 defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}, 0.175 \leq \eta \leq 0.250$. We divide the parameter space into four equally sized patches as in Fig. 3.3. As input to the SVD, we use $(35 + 1) \times (35 + 1)$ whitened waveforms in each patch's template bank. The waveforms are generated from a starting frequency of 30 Hz in the time-domain which we use to mitigate the effects of spectral leakage in the Fourier-domain waveform when whitening, and are truncated at a low frequency cut-off of 40Hz before back-transforming into the time-domain for the SVD. The SVD of each template bank uses 22 basis waveforms. To compute the mismatches we use a uniform grid of 50×50 waveforms over the parameter space defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}, 0.175 \leq \eta \leq 0.250$. We compute the mismatch for waveforms generated at the following sampling frequencies: 4096 Hz, 8192 Hz, 16384 Hz and 32768 Hz.

We observe that for the range of sampling frequencies the mismatches have a wide distribution about the threshold mismatch of 10^{-3} as is shown in Fig. 4.3. These mismatches suggest that the interpolated EOBNR waveforms are not suitable for parameter estimation as they are not sufficiently accurate. The waveforms we have considered have a relatively long (~ 2 s) inspiral followed by a

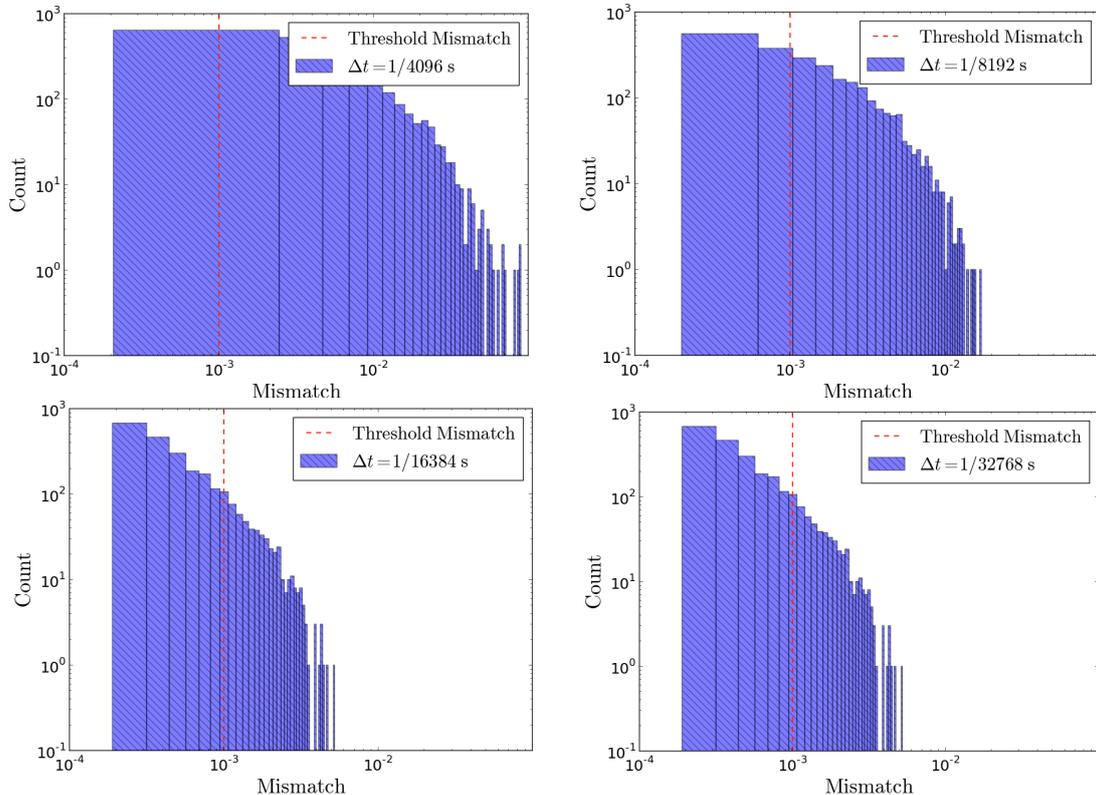


Figure 4.3: Histogram of mismatches between (normalized) interpolated and non-interpolated EOBNR waveforms generated on a uniform (50×50) grid in a patch of (\mathcal{M}, η) defined by $7.2 M_{\odot} \leq \mathcal{M} \leq 7.6 M_{\odot}$, $0.175 \leq \eta \leq 0.250$. Each panel corresponds to waveforms generated at different sampling frequencies: 4096 Hz, 8192 Hz, 16384 Hz and 32768 Hz. We set a threshold mismatch of 10^{-3} between interpolated and non-interpolated waveforms, shown as a red vertical line, above which we deem interpolated waveforms to be insufficiently accurate for parameter estimation. The best case occurs for sampling rate of 32768 Hz, though it is nearly identical to the 16384 Hz case. In all cases there is a wide distribution about the threshold mismatch of 10^{-3} which indicates that these interpolated EOBNR waveforms are not suitable for parameter estimation. We note that while these mismatches indicate that the interpolated EOBNR waveforms are not suitable for parameter estimation, they would be suitable to form a template bank for gravitational-wave searches.

4. PARAMETER ESTIMATION USING INTERPOLATED TIME-DOMAIN TEMPLATE WAVEFORMS

short ($\sim 10^{-2}$) ringdown. We note that increasing the number of basis vectors used for interpolation, *and* increasing the number of waveforms in the template bank on which we perform the SVD yields almost no improvement in the quality of interpolated waveforms.

While this is a set back for parameter estimation, there are alternative approaches to parameter estimation in which arbitrary waveforms could be used while still significantly decreasing its computational cost. We discuss these in the following chapter.

We note that while the interpolated EOBNR waveforms are insufficient for parameter estimation, they are sufficient to form a template bank for detection where the accuracy requirements are typically lower. For example, gravitational wave template banks are constructed such that neighbouring templates have a mismatch of around 3%.

4.4 Future work and Discussion

We have considered in detail the potential to use interpolated waveforms as one means to do this. For inspiral only waveforms we have found that the using interpolated waveforms leads to biases in the posterior probability distributions of the parameters. However, the bias in the mean posterior parameter estimates are negligible on the scale of statistical measurement uncertainties. In general, parameter estimation on real gravitational-wave signals from CBC sources requires that we are also able to use waveforms which describe the full inspiral merger and ringdown phases of coalescence. Our attempt to interpolate such waveforms has failed in this sense because the quality of the interpolated waveforms is insufficient for parameter estimation, though in principle could be used for searches. In the following chapter we consider how to circumvent the problems encountered in interpolating time-domain template waveforms by directly interpolating the likelihood function, without first computing an intermediate template waveform.

Chapter 5

RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

5.1 Introduction

In the previous two chapters we considered using interpolated template waveforms as a means to reduce the computational cost of parameter estimation. While the technique is promising in principle, we found that our implementation of interpolated waveforms is somewhat limited as we were unable to successfully interpolate time-domain waveforms which describe the full inspiral, merger and ringdown phases of compact binary coalescence. While there may be other, more robust, interpolation techniques which would allow generic template waveforms to be interpolated, we will not pursue them here. Instead, we will consider another means with which one could significantly reduce the computational cost of parameter estimation, again based on interpolation using the SVD.

We demonstrate that the SVD can be applied *directly* to the likelihood function itself which in turn allows the likelihood function to be directly interpolated. With this formulation of the likelihood function the only waveform computations

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

required for parameter estimation are those needed to construct the SVD of the likelihood function. We find that for a test-case, we are able to speed up the evaluation of the likelihood function by two orders of magnitude, whilst keeping likelihoods accurate to $\lesssim 0.025\%$. In addition, we will demonstrate that this improvement in computational time is conservative, and we discuss how additional computational savings would be achieved in practice. Crucially, the method we will present is totally transparent to the form of the template waveforms used for the likelihood computation.

Recall that the goal of Bayesian parameter estimation is to compute the posterior probability density function (PDF) of a set of parameters, $\vec{\theta}$, which underlie a model assumed to describe a data set d . The PDF is related to the likelihood function and prior probability via Bayes’ theorem:

$$p(\vec{\theta}|d) = \frac{\mathcal{P}(\vec{\theta}) \mathcal{L}(d|\vec{\theta})}{p(d)}, \quad (5.1)$$

where $\mathcal{L}(d|\vec{\theta})$ is the likelihood and $\mathcal{P}(\vec{\theta})$ is the prior probability which encodes our *a priori* belief in the distribution of $\vec{\theta}$. The quantity in the denominator, $p(d)$, is known as the “evidence”. Computing (5.1) requires evaluating the likelihood.

For binaries with non-spinning components $\vec{\theta}$ is nine dimensional. Exploring such a high dimensional space requires sophisticated stochastic Bayesian inference techniques [12; 13; 25] which preferentially sample the parameter space in regions of high posterior probability. The bulk of the computational cost of evaluating the likelihood function comes from computing template waveforms. Analyses on first-generation interferometer data require computing $\mathcal{O}(10^6)$ such waveforms [25; 51]. Sampling techniques such as Markov chain Monte Carlo (MCMC) [13; 25] and nested sampling [12; 94] evaluate likelihoods, and hence compute template waveforms, serially. Thus the total computational time to fully sample the parameter space scales linearly with the total time spent generating template waveforms. It can take hours to weeks to analyse a single stretch of data of a few seconds in duration, depending on the choice of the template waveform family. This problem will be exacerbated when analysing second-generation interferometer data as the waveforms will be forty times longer in duration if the starting frequency f_{\min} is changed from 40 Hz to 10 Hz.

For binaries with non-spinning components, the frequency-domain waveform $\tilde{h}(\vec{\theta}; f)$ has the schematic form

$$\tilde{h}(\vec{\theta}; f) = \sum_{\mu=+,\times} A_{\mu}(\vec{\theta}_L) \tilde{h}_0(\mathcal{M}, \eta; f) e^{2\pi i f t_c}, \quad (5.2)$$

where $A_{+,\times}$ denotes the (scalar) amplitudes of the “plus-” and “cross-” polarization states of the waveform. In general \tilde{h}_0 depends on the waveform family being used and can be computed by Fourier transforming the associated time-domain representation of the waveform family. The parameters which describe the binary are the chirp mass and symmetric mass-ratio, \mathcal{M} and η , the time at coalescence t_c and a set of parameters which describe the location and orientation of the binary $\vec{\theta}_L$.

Evaluating the likelihood function on the three-dimensional subspace of parameters (\mathcal{M}, η, t_c) represents the largest computational burden to parameter estimation on gravitational waves from CBC sources with non-spinning components because the likelihood function depends non-trivially on these parameters, and so requires a new waveform evaluation. In [51], we considered interpolation between waveforms over the mass parameter space as a way to reduce computational cost. Here, we demonstrate that the evaluation of an interpolated likelihood function over the (\mathcal{M}, η, t_c) subspace is a much faster computational procedure than the standard calculation of the likelihood (5.3) by using either full or interpolated waveforms. For the purposes of parameter estimation, one is not interested in template waveforms per se, but rather in the posterior probability distributions of the underlying parameters of the template waveforms that are assumed to describe the data. By directly using interpolated likelihood functions, one effectively bypasses template waveform generation during the sequential steps of an MCMC. This likelihood-interpolation technique is robust and could, in principle, be generalized to arbitrary template waveform families, in particular those that describe CBCs with spinning components.

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

5.1.1 Directly interpolating the likelihood function

We wish to generate a representation of the likelihood function over the continuous \mathcal{M}, η and t_c subspace. To achieve this we will interpolate the likelihood function over \mathcal{M}, η and t_c . The likelihood function that describes the probability of observing a data stream $d = h + n$ containing a given gravitational-wave signal $h(\vec{\theta}; t)$ and Gaussian and stationary noise $n(t)$ is [12]

$$\log \mathcal{L}(d|\vec{\theta}) = (d|h(\vec{\theta})) - \frac{1}{2} \left[(h(\vec{\theta})|h(\vec{\theta})) + (d|d) \right], \quad (5.3)$$

where $(a|b)$ is the usual noise-weighted inner product [46]. We define the complex-valued time-series corresponding to the inner product between two time series $a(t)$ and $b(t)$ as one is shifted by an amount t_c with respect to the other:

$$z[a, b](t_c) := 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} e^{-2\pi i f t_c}, \quad (5.4)$$

In the above, $\tilde{a}(f)$ is the Fourier transform of $a(t)$ and $S_n(f)$ is the detector noise power spectral density (PSD). The limits of integration are in general specified by the bandwidth over which an analysis is being conducted. In terms of $z(t_c)$ the inner products in (5.3) are succinctly expressed as

$$(h(\vec{\theta})|h(\vec{\theta})) = \Re \mathcal{A}(\vec{\theta}_L) z[h_0(\mathcal{M}, \eta), h_0(\mathcal{M}, \eta)](0), \quad (5.5)$$

$$(d|h(\vec{\theta})) = \Re \mathcal{B}(\vec{\theta}_L) z[d, h_0(\mathcal{M}, \eta)](t_c), \quad (5.6)$$

and $\mathcal{A}(\vec{\theta}_L)$ and $\mathcal{B}(\vec{\theta}_L)$ are known projections which contain the $\vec{\theta}_L$ dependence in the likelihood function.

We have previously interpolated template waveforms over the mass parameters [50; 51], and here we show that the same technique can be applied to interpolating the time series $z[d, h_0](t_c)$. The interpolation of $z[d, h_0](t_c)$ is based on the SVD of a set of (discretely sampled) time series distributed on a two-dimensional grid. In this case the two-dimensional grid spans \mathcal{M} and η and the time parameter is t_c . We use the notation $\vec{z}[d, h_0]$ to describe the discretely sampled $z[d, h_0](t_c)$. Recall that the SVD of the discretely sampled time series $\vec{z}[d, h_0]$ allows it to be written as a linear superposition of orthonormal basis vectors \vec{u}_μ and projection

coefficients M_μ [49]:

$$\vec{z}[d, h_0(\mathcal{M}, \eta)] = \sum_{\mu} M_\mu(\mathcal{M}, \eta) \vec{u}_\mu. \quad (5.7)$$

The coefficients M_μ can be interpolated over \mathcal{M} and η and we follow the method in [50] which uses Chebyshev polynomials of the first kind.

Interpolation of $z[h_0, h_0](0)$ over \mathcal{M} and η is simple as it is scalar valued and we again use Chebyshev polynomials of the first kind. Below we provide an example of the interpolation technique outlined here.

5.1.2 Likelihood interpolation: Examples

We compare interpolated likelihood functions to those generated by direct evaluation of waveforms and inner products. We consider two test cases: (i) the coalescence of binary black holes, and (ii) the coalescence of binary neutron stars.

We generate a discretely sampled, simulated data set \vec{d} for a single interferometer consisting of Gaussian and stationary noise \vec{n} and a gravitational-wave signal \vec{h} . The data set is 32 s in duration and has a sampling rate in the time domain of 4096 Hz. For binary black holes we model the gravitational-wave signal \vec{h} using the effective one-body approach calibrated to numerical relativity simulations (EOBNR) [41]. Such a gravitational-wave signal describes the full inspiral, merger and ringdown phases of coalescence. For binary neutron stars we model the gravitational waveform using a post-Newtonian (PN) model computed to 3.5 PN order in phase [26], which describes the inspiral phase of the coalescence only. We use an implementation of EOBNR and post-Newtonian waveforms from the LSC Algorithms Library (LAL) [85] corresponding to the approximants EOBNRv2 and TaylorT4 respectively.

Generating the interpolant of the likelihood function requires the following stages: (i) patch the mass space into smaller domains, (ii) generate a set of waveforms over a dense grid in each patch, (iii) filter the data with the template waveforms to compute the likelihoods, (iv) pack the likelihoods into a matrix and perform the SVD, (v) build the interpolant in each patch. Only after these stages have been completed can the interpolated likelihood function be sampled.

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

We first construct a discrete, uniform grid of template waveforms in $\mathcal{M} - \eta$ parameter space. We will use a small region around the parameters of the signal, as \mathcal{M} and η are typically constrained to $\lesssim 1\%$ and $\lesssim 10\%$, respectively, depending on the signal parameters and signal-to-noise ratio (SNR) [51; 95]. The region in $\mathcal{M} - \eta$ where the posterior has significant support can be found quickly during the burn-in phase of the MCMC, which requires a small fraction of the total number of samples necessary to evaluate the posterior probability distribution function.

We use the Chebyshev interpolation described in [50] to interpolate $z[h_0, h_0](0)$ for waveforms across the grid. To interpolate $\vec{z}[d, h_0]$ we first find the basis vectors \vec{u}_μ by constructing a matrix from the set of $\{\vec{z}[d, h_0]\}$, the columns of which correspond to a unique $\vec{z}[d, h_0]$ on the grid of waveforms, which we factor using the SVD. After performing the SVD, we truncate the number of basis vectors such that on average the norm of each \vec{z} is conserved to one part in 10^5 [50]. This can significantly reduce the number of basis vectors. We then apply the Chebyshev interpolation [50] to interpolate projection coefficients across the $\mathcal{M} - \eta$ grid.

5.1.2.1 Example 1: High-mass binary black holes

The signal is parameterized by $\vec{\theta}^s = (\mathcal{M} = 15.01 M_\odot, \eta = 0.205, D = 100 \text{ Mpc}, \iota = 0, \psi = 0, \alpha = 0, \delta = 0, t_c = 0.1 \text{ s}, \phi_c = 0)$. We use a noise PSD typical of initial LIGO [1]. The signal has an SNR of ≈ 15 . In order to interpolate the likelihood function across \mathcal{M}, η and t_c , we work within a small region of $\mathcal{M} - \eta$ space whose boundaries are given by $14.56 M_\odot \leq \mathcal{M} \leq 15.46 M_\odot$ and $0.143 \leq \eta \leq 0.25$. Assuming a statistical measurement uncertainty on \mathcal{M} and η of 1% and 10%, respectively, the parameter ranges correspond to a $\sim 3\sigma$ range about the signal value. Note that we cannot go about $\eta = 0.25$ in the η interval. We further restrict our range in t_c to be in a $\pm 0.2 \text{ s}$ window about the trigger time, which is a common time prior in Bayesian parameter estimation [12].

In Fig. (5.1) we compare a likelihood function generated via direct evaluation of inner products, to one which we have generated via SVD-interpolation. We find that we are able to reconstruct the log likelihood function by interpolation to within a fractional percentage error of at most 0.025%. While we have only plotted an interpolated likelihood function at the signal values of \mathcal{M} and η , the

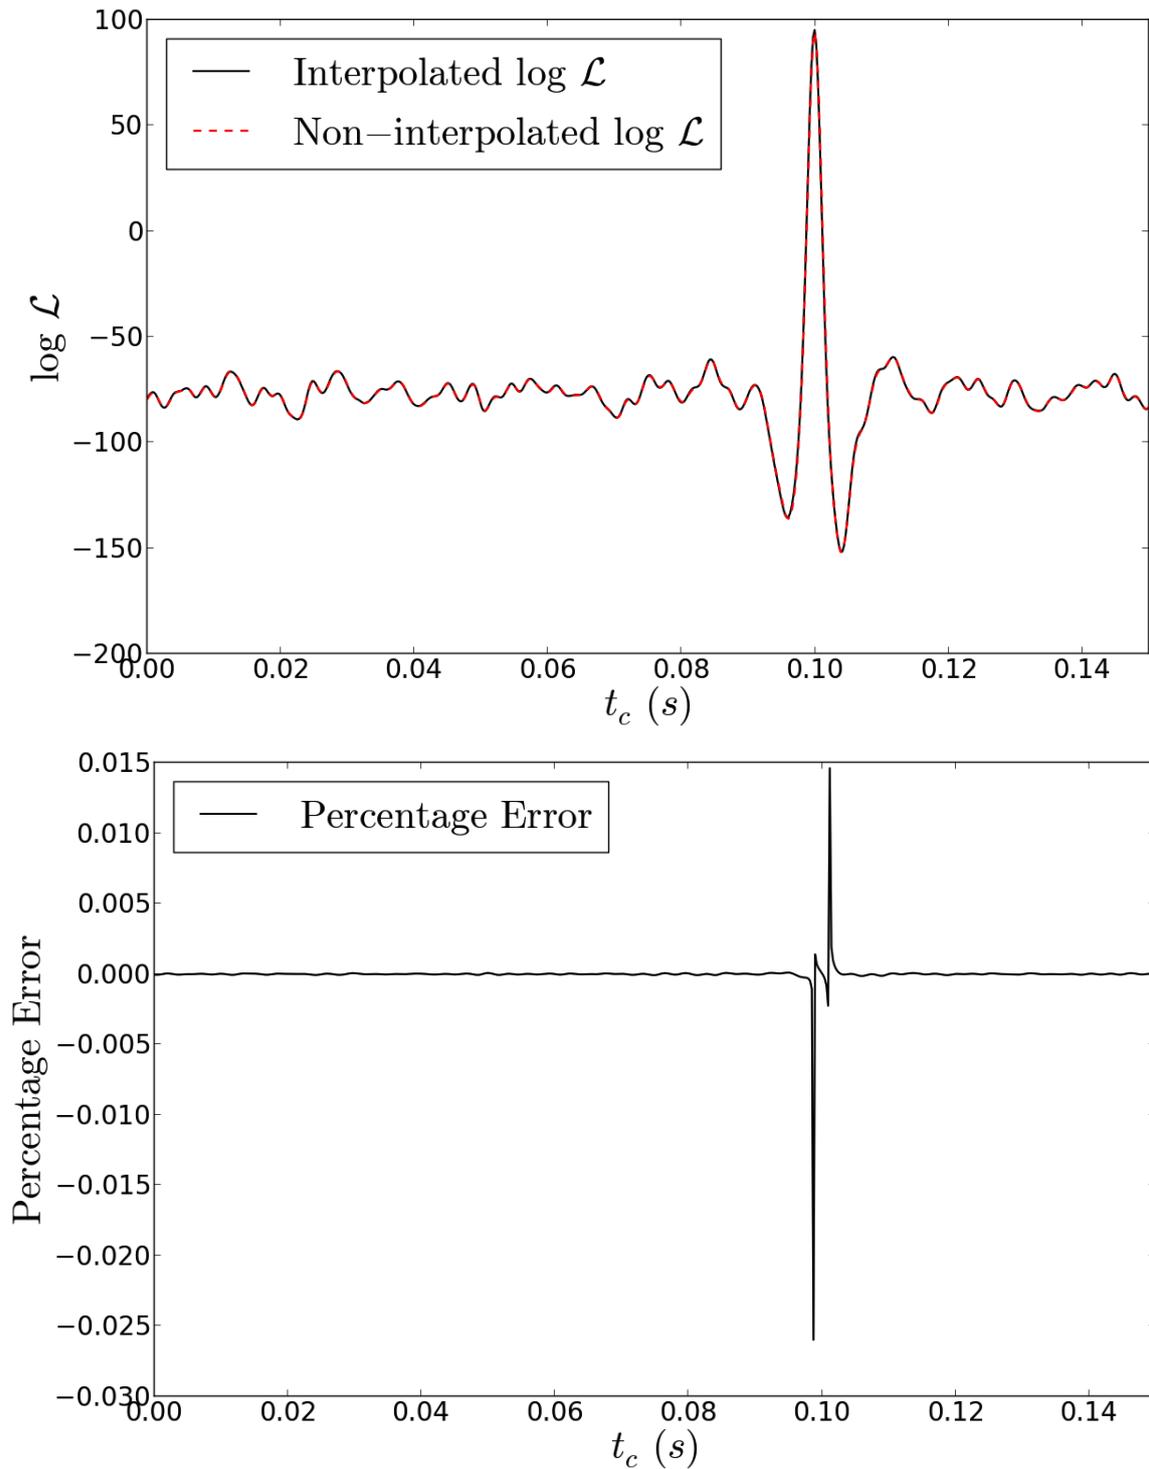


Figure 5.1: Interpolated and non-interpolated log likelihoods (top), and percentage error (bottom) for a data set containing a gravitational-wave signal from the coalescence of binary black holes.

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

errors quoted here are typical across the mass range we have considered. Meanwhile, for this waveform model and parameters, computing the likelihood via the interpolation procedure is around two orders of magnitude faster than generating a template waveform and directly evaluating the inner products in (5.3).

5.1.2.2 Example 2: Binary neutron stars

The signal is parameterized by $\vec{\theta}^s = (\mathcal{M} = 1.217 M_\odot, \eta = 0.2497, D = 20 \text{ Mpc}, \iota = 0, \psi = 0, \alpha = 0, \delta = 0, t_c = 0.1 \text{ s}, \phi_c = 0)$. We again use a noise PSD typical of initial LIGO [1], and the signal has $\text{SNR} \approx 15$. We interpolate the likelihood function over a small region of $\mathcal{M} - \eta$ space whose boundaries are given by $1.199 M_\odot \leq \mathcal{M} \leq 1.235 M_\odot$ and $0.212 \leq \eta \leq 0.25$. Assuming a statistical measurement uncertainty of 0.5% on \mathcal{M} and 5% on η , these parameter ranges correspond to a $\sim 3\sigma$ range about the signal value. Note that we cannot go about $\eta = 0.25$ in the η interval. We restrict our range in t_c to be in a $\pm 0.2 \text{ s}$ window about the trigger time.

Again, we find that we are able to reconstruct the log likelihood function to within a fractional percentage error of at most 0.025%. For the binary neutron star case, we find that computing the likelihood via interpolation is around three orders of magnitude faster than direct evaluation. This difference is larger than for the higher-mass binary black hole case because the waveform duration is significantly longer for binary neutron stars, whereas the cost of computing interpolated likelihoods remains fixed.

Below we discuss practical issues pertaining to incorporating interpolated likelihoods into real gravitational-wave parameter-estimation pipelines.

5.1.2.3 Practical considerations

For our interpolation technique to be viable for real data analyses, the total computational time of first constructing the interpolated likelihood function, and then sequentially sampling the interpolated likelihood function, must be less than the time for sequentially sampling the likelihood function directly.

Parameter estimation requires sampling the parameter space until the sampler has met its convergence criterion. The total number of likelihood evaluations

for convergence is typically $\sim \mathcal{O}(10^6)$ [25; 51]. When directly evaluating the likelihood function, the number of likelihood evaluations is a reasonable proxy for the number of waveform evaluations, which dominate the computational cost.

To sample the interpolated likelihood function there is an additional upfront cost of constructing the interpolant of the likelihood function. This cost will depend on the region of the parameter space over which the likelihood function needs to be interpolated and template waveforms must be computed. However, building the interpolant is highly parallelizable and computing it over an extended region of parameter space could be split into multiple independent subsets. This could greatly reduce the wall time of computing the interpolant. We have noted that one can restrict the range in parameter space over which the interpolant is built by using an MCMC to sparsely explore the parameter space in regions of high posterior probability. In practice, the number of samples for this “burn-in” is often $\sim \mathcal{O}(10^4)$ [51], and the likelihood has significant support in a relatively small patch in parameter space. The likelihoods computed during the burn-in evaluation could thus be stored for future interpolation.

One could also interpolate $\vec{z}(t_c)$ over many patches covering the parameter space in parallel. We have not investigated optimal patching, nor the required denseness of likelihood template calculation in order to generate a good basis for $\vec{z}(t_c)$; this will be the subject of future work.

Once the interpolant is constructed, the cost of sampling the parameter space will depend on that of computing the interpolated likelihood function. In our example we found that computing the interpolated likelihood function is between two and three orders of magnitude cheaper than directly evaluating the likelihood function, depending on the region in parameter space in which the likelihood function is being computed. The actual improvement will depend on the typical cost of waveform computation, which is a function of both the template waveform family used and the waveform parameters.

Here we had used the SVD to find a basis for the set of $\vec{z}(t_c)$. The SVD is not a unique technique for finding a basis set, and we note that the authors in [96] and [97] employ a greedy algorithm to efficiently generate a set of bases for gravitational waveforms which could in principle be applied to a set of $\vec{z}(t_c)$.

We have so far discussed interpolation in the mass parameters. It may also be

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

necessary to interpolate the quantities $z[d, h_0](t_c)$ in the t_c direction, because the coalescence time in a particular interferometer may lie in between discretely sampled time points. Second-order interpolation provides sufficient accuracy when the waveform is sampled at 4 kHz.

5.1.2.4 Discussion and conclusion

We have demonstrated a method to sample the CBC likelihood function via interpolation, with improvements of two to three orders of magnitude in efficiency. Our method utilizes a SVD of the likelihood function on a three-dimensional subspace of parameters: the chirp mass \mathcal{M} , symmetric mass-ratio η and time at coalescence t_c . The SVD factors the likelihood function into a set of scalar coefficients which describe a surface in \mathcal{M} and η , and a set of orthonormal basis vectors which describe how the surface is translated along t_c . The projection coefficients can be interpolated on the $\mathcal{M} - \eta$ plane and then trivially scaled by elements of the basis vectors to generate the likelihood at (\mathcal{M}, η, t_c) . This provides an efficient means to interpolate in three dimensions.

We note that while we have chosen an interpolation technique based on the SVD, it is by no means unique and other interpolation techniques have been applied to gravitational-wave data analysis [e.g., 97]. Notably, Mitra *et al.* [98] considered interpolating the matched-filtered output of gravitational-wave searches. They interpolated the signal-to-noise ratio, which is effectively a component of the likelihood function, and so their method could, in principle, be extended to interpolate likelihood functions. The key difference with our approach is that we use a decomposition of the likelihood as a function of time, while [98] treat it as a scalar quantity. This provides us with an efficient means of reducing the total data needed for interpolation, exploiting correlations along the t_c direction by ranking the basis vectors in order of importance in reconstructing the likelihood function. Hence, we can effectively exorcise redundant information based on our accuracy requirements. The number of bases needed to approximately reconstruct the likelihood to high accuracy using the SVD is generally small compared to the number of raw likelihood vectors which we decompose.

Likelihood interpolation appears to be more robust than waveform interpola-

tion [51], and so utilizing interpolated likelihood functions may also be a stepping stone to tackling the more difficult issue of rapidly estimating the parameters of binaries with spinning components.

5. RAPIDLY EVALUATING THE COMPACT BINARY LIKELIHOOD FUNCTION VIA INTERPOLATION

Chapter 6

SUMMARY

We have considered two aspects of gravitational-wave astronomy; the ability to detect intermediate-mass ratio coalescences (IMRACs) with Advanced LIGO (aLIGO), and methods to improve the efficiency of parameter estimation on gravitational waves from coalescing compact binaries using the singular value decomposition (SVD).

In Chapter II we studied the requirements of model waveforms for detecting IMRACs using aLIGO. We considered binaries in which the intermediate-mass black hole has mass in the range $24 M_{\odot} - 200 M_{\odot}$ with a stellar-mass companion having masses in the range $1.4 M_{\odot} - 18.5 M_{\odot}$ and with mass ratios $1/140 \leq q \leq 1/10$. In this mass space we studied the contributions to the signal-to-noise ratio (SNR) from the three stages of binary coalescence; inspiral, merger and ringdown. We have shown that over the bulk of the mass space, merger and ringdown contribute significantly to the coalescence signal, despite the suppression of the power in the merger and ringdown in the coalescence signal. This is because the systems we have considered can merge in the frequency range in which aLIGO is most sensitive. For the mass space we considered, the inspiral portion of the coalescence signal may fall at frequencies below the detectors sensitive band and hence the bulk contribution to SNR comes from the merger and ringdown. However, we have also shown that despite the importance of merger and ringdown, there is a small portion of the mass space in which the inspiral portion dominates the coalescence signal. We identified three regions in the mass space in which different searches could be considered appropriate based on thresh-

6. SUMMARY

olds of acceptable losses in detection rates. The mass space splits into a region in which inspiral-only searches could be feasible, incurring losses in detection rates of up to $\sim 27\%$; a region in which searches would be limited by lack of merger and ringdown in template waveforms, incurring losses in detection rates up to 50% ; and a region in which merger and ringdown are critical to prevent losses in detection rates over 50% . The search regions are summarized in Table 2.1.

We also considered the accuracy of model waveforms which are currently used for gravitational-wave searches for high-mass binaries in the IMRAC regime. Model waveform families based on the effective one-body approach, matched to numerical relativity - known as EOBNR [26] - are only known to be accurate at describing binaries with mass ratios $q > 1/6$ and so to detect IMRACs one needs waveforms which are able to describe binaries which have smaller mass ratios. We thus compared the EOBNR waveforms to a waveform family designed for such small mass-ratio binaries which we refer to as the “Huerta-Gair” (HG) [44; 63] waveform family. Because HG waveforms only describe the inspiral portion of the coalescence signal, we only compared the EOBNR inspiral to HG waveforms. By computing the fitting factor between these waveform families, we found that EOBNR-inspirals are $\geq 90\%$ effective at describing the HG inspiral. This implies that EOBNR waveforms may be suitable for conducting a search for gravitational waves from IMRACs. However, this result reinforces the need to develop approximation schemes which faithfully describe IMRACs. In particular, due to the importance of merger and ringdown in the coalescence signal, numerical relativity simulations in this regime will be important in order to produce template waveforms which describe the full coalescence signal.

In Chapters III, IV and V we considered methods to improve the efficiency of parameter estimation of gravitational waves from coalescing binaries using the singular value decomposition (SVD). We applied the SVD in two ways: (i) to interpolate template waveforms and (ii) to directly interpolate the likelihood function. The SVD has already been used to interpolate template waveforms for gravitational-wave searches [49] which reduces the computational cost of waveform generation in searches. Because the waveform accuracy requirements of parameter estimation are typically higher than for gravitational-wave searches, in Chapter III we investigated the potential to use interpolated waveforms in pa-

parameter estimation as a means to reduce the computational cost associated with waveform generation. We provided a proof-of-principle that interpolated template waveforms can be used for parameter estimation. While using interpolated waveforms leads to biases in the posterior probability density functions (PDFs), the bias in the mean posterior parameter estimates are smaller than the scale of statistical measurement uncertainty and thus we pass a common accuracy threshold for model waveform accuracy [89]. Furthermore, we provided a comparison of the computational times of various waveform families. We found that in general, interpolating template waveforms is at least an order of magnitude faster than using non-interpolated waveforms and so using interpolated template waveforms for parameter estimation is a promising technique to improve its efficiency.

Despite the potential of using interpolated waveforms, the study in Chapter III was limited to one waveform family. We used the SVD to interpolate a class of waveforms known as TaylorF2 [26] which are already efficient to generate. This was because of the simplicity of implementing TaylorF2 waveforms. In Chapter IV we considered interpolating generic classes of template waveform. This extended the proof-of-principle in Chapter III to include inspiral-only time-domain waveforms and time-domain waveforms which describe the full coalescence signal (inspiral, merger and ringdown). We found that we were able to use an interpolated inspiral-only time-domain waveform family known as TaylorT4 [26] as templates for parameter estimation. We reported negligible biases in mean parameter estimates on the scale of statistical measurement uncertainty. We found that interpolated TaylorT4 waveforms are around an order of magnitude faster to generate than non-interpolated TaylorT4 waveforms, which is consistent with our results in Chapter III. We also were able to interpolate time-domain waveforms which describe the full coalescence signal. We used a waveform family which uses the “effective one-body” approach, matched to numerical simulations of the merger and ringdown, known as EOBNR [26]. However, we found that we were not able to generate the interpolated templates at sufficient accuracy for parameter estimation, though they could be considered accurate enough for gravitational-wave searches. There were a number of reasons for the loss in accuracy. Firstly, the relative time scales of the inspiral and ringdown are significantly different. For the mass range we considered the inspiral phase is around 2 s in

6. SUMMARY

duration while the ringdown is $\mathcal{O}(10^{-2})$ s in duration. In addition, we encountered a difficulty in our implementation of interpolated waveforms which makes it difficult to extend to time-domain waveforms. When performing the SVD of a set of template waveforms we found it useful to first align the waveforms in time at the frequency at which the noise PSD is a minimum. The time-frequency evolution of time-domain waveforms can be inaccurate unless one uses a very high sample rate, and hence the alignment of the waveforms we used can be off by a small amount. This finding was reinforced by the fact that at high sampling rates, the accuracy of interpolated time-domain waveforms becomes higher.

In order to address the issue of efficient parameter estimation with generic waveforms, in Chapter V we showed that the SVD enables one to directly interpolate the likelihood function. We analytically formulated the interpolated likelihood function and showed that parameter estimation based on an interpolated likelihood function solves numerous computational issues. For example, waveform generation is only required in order to generate the SVD of the likelihood function and evaluating the interpolated likelihood function requires no new waveform computations. In addition likelihood interpolation should be totally transparent to the template waveform family being used. We reported (conservative) improvements in the efficiency of evaluating the interpolated likelihood function between two and three orders of magnitude, whilst keeping likelihoods accurate to $\lesssim 0.025\%$. We believe that this technique is highly relevant to parameter estimation in aLIGO because the in-band signals are liable to be around two orders of magnitude longer than initial LIGO which will exacerbate the cost of template waveform generation, and hence parameter estimation. Further work on directly interpolating the likelihood function is well merited and is a promising approach to rapidly estimating the parameters of compact binaries.

Appendix A

FREQUENTLY USED QUANTITIES

A.1 Definitions of discrete representations of continuous functions

Here we define the conventions that we adopt for discrete representations of continuous functions, which are frequently used.

For a real-valued function $a(t)$, the convention for the Fourier Transform is

$$\tilde{a}(f) = \int_{-\infty}^{+\infty} dt a(t) e^{-2\pi i f t}, \quad (\text{A.1})$$

with the inverse transform

$$a(t) = \int_{-\infty}^{+\infty} df \tilde{a}(f) e^{+2\pi i f t}. \quad (\text{A.2})$$

For the case of discrete time/frequency series, consider a (discrete) real-valued time series \vec{a} sampled at discrete time intervals $\Delta t = 1/(2f_{\text{Nyq}})$, where f_{Nyq} is the *Nyquist frequency*. The discrete time series has a duration T in the time-domain. The total number of data samples of the discrete, real-valued time-series is therefore $N = T/\Delta t = 2Tf_{\text{Nyq}}$. The data points at discrete time t_j and

A. FREQUENTLY USED QUANTITIES

frequency f_k of \vec{a} , and its (discrete) Fourier transform \tilde{a} , are therefore:

$$a(t_j) = a(j\Delta t) = a_j \quad (\text{A.3})$$

$$\tilde{a}(f_k) = \tilde{a}(k\Delta f) = \Delta t \times \tilde{a}_k, \quad (\text{A.4})$$

where we have defined the Fourier series as:

$$\tilde{a}_k = \sum_j a_j e^{-2\pi i j k / N}, \quad (\text{A.5})$$

$$a_j = \frac{1}{N} \sum_k \tilde{a}_k e^{+2\pi i j k / N}. \quad (\text{A.6})$$

We approximate the inner product $(a|b)$ between two (real) functions a and b in the discrete case as:

$$\begin{aligned} (a|b) &= 4\Re \int_0^\infty \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} df, \quad (\text{A.7}) \\ &\approx 4\Re \Delta f \sum_{k=-}^{N/2} \frac{\tilde{a}^*(f_k)\tilde{b}(f_k)}{S_n(f_k)}, \quad (\text{A.8}) \end{aligned}$$

where S_n is the detector noise power spectral density (PSD). The inner product can further simplified by using “whitened” frequency series. We define a whitened frequency series as the raw frequency series weighted by the square root of the noise PSD:

$$\tilde{a}_w(f_k) = \frac{\tilde{a}(f_k)}{\sqrt{S_n(f_k)}}, \quad (\text{A.9})$$

Using whitened waveforms, we define the vector inner product between two discretely sampled, whitened time-series \vec{a}_w and \vec{b}_w , as

$$\vec{a}_w \cdot \vec{b}_w := 4 \Delta f \sum_{k=0}^{N/2} \tilde{a}_w^*(f_k)\tilde{b}_w(f_k), \quad (\text{A.10})$$

which in general is complex. The standard definition of the inner product, Eq. A.8, is just the real part of the vector inner product.

A.2 Post-Newtonian phasing formulae

We provide useful formulae for post-Newtonian gravitational waveforms. We use v as our expansion parameter which is given by $v = (\pi M f)^{1/3}$ in geometric units, $G = c = 1$, where M is the total mass of the binary.

A.2.1 TaylorF2

For the frequency-domain representation of the gravitational-wave phasing, we use the formula corresponding to the approximant TaylorF2. The phasing up to 3.5PN order is given by [26]

$$\begin{aligned}
\psi_{3.5}^{(\text{F2})}(f) = & 2\pi f t_c - \phi_c - \frac{\pi}{4} + \frac{3}{128 \eta v^5} \left[1 + \frac{20}{9} \left(\frac{743}{336} + \frac{11}{4} \eta \right) v^2 - 16\pi v^3 \right. \\
& + 10 \left(\frac{3058673}{1016064} + \frac{5429}{1008} \eta + \frac{617}{144} \eta^2 \right) v^4 + \pi \left(\frac{38645}{756} - \frac{65}{9} \eta \right) \\
& \left. \left\{ 1 + 3 \log \left(\frac{v}{v_{\text{iso}}} \right) \right\} v^5 + \left\{ \frac{11583231236531}{4694215680} - \frac{640}{3} \pi^2 - \frac{6848 \gamma}{21} - \frac{6848}{21} \right. \right. \\
& \left. \left. \log(4v) \left(-\frac{15737765635}{3048192} + \frac{2255 \pi^2}{12} \right) \eta + \frac{76055}{1728} \eta^2 - \frac{127825}{1296} \eta^3 \right\} v^6 \right. \\
& \left. + \pi \left(\frac{77096675}{254016} + \frac{378515}{1512} \eta - \frac{74045}{756} \eta^2 \right) v^7 \right],
\end{aligned} \tag{A.11}$$

where $\eta = m_1 m_2 / (m_1 + m_2)^2$ is the symmetric mass ratio, m_1 and m_2 are the component masses, and $\gamma = 0.577216\dots$ is the Euler constant.

A.2.2 TaylorT4

The TaylorT4 orbital phase is given by solving the coupled differential equations of the orbital phase and velocity, $\phi(t)$ and $v(t)$, respectively,

$$\frac{d\phi}{dt} - \frac{v^3}{M} = 0, \tag{A.12}$$

$$\frac{dv}{dt} = \mathcal{K}(t), \tag{A.13}$$

A. FREQUENTLY USED QUANTITIES

where the acceleration $\mathcal{K}(t)$ is given by [26]

$$\begin{aligned}
\mathcal{K}(t) = & \frac{32}{5} \frac{\eta}{M} v^9 \left[1 - \left(\frac{743}{336} + \frac{11}{4} \eta \right) v^2 + 4\pi v^3 + \left(\frac{34103}{18144} + \frac{13661}{2016} \eta + \frac{59}{18} \eta^2 \right) v^4 \right. \\
& - \left(\frac{4159}{672} + \frac{189}{8} \eta \right) \pi v^5 + \left(\frac{16447322263}{139708800} + \frac{16}{3} \pi^2 - \frac{1712}{105} \gamma \right. \\
& + \left. \left(\frac{451}{48} \pi^2 - \frac{56198689}{217728} \right) \eta \frac{541}{896} \eta^2 - \frac{5605}{2592} \eta^3 - \frac{856}{105} \log(16v^2) \right) v^6 \\
& \left. - \left(\frac{4415}{4032} - \frac{358675}{6048} \eta - \frac{91495}{1512} \eta^2 \right) \pi v^7 \right]. \tag{A.14}
\end{aligned}$$

Appendix B

SINGULAR VALUE DECOMPOSITION OF $M \times N$ MATRICES

This appendix provides supplementary material to Sec. 3.3. We give examples of how the singular value decomposition (SVD) can be performed in practice for simple $M \times N$ matrices. In the examples, we will provide a geometric interpretation of the singular value decomposition.

B.1 Definitions of the singular value decomposition of $M \times N$ matrices

Consider a $M \times N$ matrix \mathbf{A} , where M is the number of rows of \mathbf{A} and N is the number of columns. For simplicity we assume $M \geq N$. The singular value decomposition of \mathbf{A} is a factorization of the form [99]

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \tag{B.1}$$

where \mathbf{U} is an orthonormal $M \times N$ matrix, \mathbf{V} is an orthonormal square $N \times N$ matrix and $\mathbf{\Sigma}$ is a diagonal square $N \times N$ matrix. The matrices \mathbf{U} , \mathbf{V} and $\mathbf{\Sigma}$ are defined as follows [100]:

B. SINGULAR VALUE DECOMPOSITION OF $M \times N$ MATRICES

- The diagonal elements of $\mathbf{\Sigma}$ are known as the “singular values” of \mathbf{A} and are denoted by σ_i . The singular values correspond to the positive square roots of the eigenvalues of the associated matrix $\mathbf{K} = \mathbf{A}^T \mathbf{A}$. For convenience, the singular values are ordered such that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N$.
- The N columns of \mathbf{V} and \mathbf{U} are known as the right- and left-singular vectors of \mathbf{A} respectively. The right-singular vectors, \vec{v}_i are the normalized eigenvectors of $\mathbf{A}^T \mathbf{A}$ and the left-singular vectors, \vec{u}_i , are the normalized eigenvectors of $\mathbf{A} \mathbf{A}^T$ and satisfy

$$\mathbf{A} \vec{v}_i = \sigma_i \vec{u}_i, \quad (\text{B.2})$$

$$\mathbf{A}^T \vec{u}_i = \sigma_i \vec{v}_i. \quad (\text{B.3})$$

Below we provide a step-by-step example of computing the SVD of a simple matrix \mathbf{A} .

B.1.1 Step-by-step example

Compute the SVD of the matrix \mathbf{A} given by

$$\mathbf{A} = \begin{bmatrix} -2 & 2 \\ 1 & 2 \end{bmatrix}. \quad (\text{B.4})$$

Step 1: Compute the singular values of \mathbf{A}

To compute the singular values we first need to compute $\mathbf{A}^T \mathbf{A}$:

$$\mathbf{K} = \mathbf{A}^T \mathbf{A} = \begin{bmatrix} 5 & -2 \\ -2 & 8 \end{bmatrix}. \quad (\text{B.5})$$

The eigenvalues, λ_i , are given by solving for the characteristic polynomial,

$$\det(\mathbf{K} - \lambda \mathbf{I}_2) = 0, \quad (\text{B.6})$$

where \mathbf{I}_2 is the 2×2 identity matrix. It is easy to show that the eigenvalues are

$$\lambda_1 = 9, \quad \lambda_2 = 4, \quad (\text{B.7})$$

and hence the singular values σ_i are

$$\sigma_1 = 3, \quad \sigma_2 = 2 \quad (\text{B.8})$$

Step 2: Compute the right-singular vectors of \mathbf{A}

The normalized eigenvectors, \vec{v}_i , of \mathbf{K} are the right-singular vectors of \mathbf{A} . It is simple to show that these are given by

$$\vec{v}_1 = \frac{1}{\sqrt{5}} \begin{bmatrix} -1 \\ 2 \end{bmatrix}, \quad \vec{v}_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} -2 \\ -1 \end{bmatrix}. \quad (\text{B.9})$$

Step 3: Compute the left-singular vectors of \mathbf{A}

Having computed the singular values and right-singular vectors of \mathbf{A} , we can easily compute its left-singular vectors using Eq. (B.2). Since $\sigma_1, \sigma_2 \neq 0$ we can immediately calculate \vec{u}_1 and \vec{u}_2 from Eq. (B.2):

$$\vec{u}_1 = \frac{1}{\sigma_1} \mathbf{A} \vec{v}_1 = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 \\ 1 \end{bmatrix}, \quad \vec{u}_2 = \frac{1}{\sigma_2} \mathbf{A} \vec{v}_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 \\ -2 \end{bmatrix}. \quad (\text{B.10})$$

Step 4: Writing the singular value decomposition of \mathbf{A}

We now have all the information to write the singular value decomposition of \mathbf{A} . The three matrices \mathbf{U} , \mathbf{V} and $\mathbf{\Sigma}$ are formed thus: The i^{th} column of \mathbf{U} and \mathbf{V} are given by the u_i and v_i respectively, and the diagonal elements of $\mathbf{\Sigma}$ are ordered such that $\Sigma_{11} = \sigma_1$, $\Sigma_{22} = \sigma_2$ etc... Hence,

$$\mathbf{U} = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} -\frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \end{bmatrix}, \quad \mathbf{\Sigma} = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix}. \quad (\text{B.11})$$

The singular value decomposition of \mathbf{A} is thus

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \end{bmatrix}, \quad (\text{B.12})$$

and it can be easily verified that this representation is equivalent to Eq. (B.4).

B. SINGULAR VALUE DECOMPOSITION OF $M \times N$ MATRICES

B.1.2 Geometric picture of the SVD

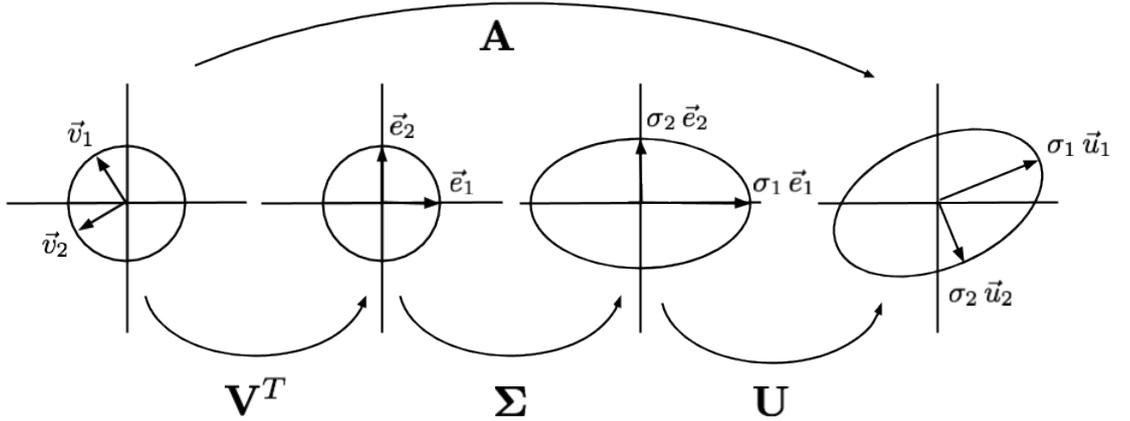


Figure B.1: Action of the matrix \mathbf{A} , Eq. (B.4), on the unit circle. The transformation of the vectors \vec{v}_1 and \vec{v}_2 , Eq. (B.9), is shown explicitly.

The SVD has a deep connection with geometry (see e.g. [100; 101]) which we will touch upon here. The canonical example of how the SVD relates to geometry is to consider the effect of a matrix \mathbf{A} on the unit circle, by looking at the contributions of each matrix of the SVD of \mathbf{A} for \mathbf{A} given by Eq. (B.4). The idea is to think of \mathbf{V} and \mathbf{U} as rotations and reflections, and $\mathbf{\Sigma}$ as a stretching. Fig. B.1 illustrates the sequence of transformations induced by acting the matrix \mathbf{A} on the unit vectors \vec{v}_1 and \vec{v}_2 , and all other vectors on the unit circle when \mathbf{A} . We can understand how \vec{v}_1 and \vec{v}_2 are transformed by solving the following system of linear equations

$$\mathbf{A} \vec{v}_i = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \vec{v}_i = \vec{x}_i, \quad (\text{B.13})$$

for $i = (1, 2)$. We already know the solution to this equation, which is given in Eq. (B.2), however it is useful to provide a geometric interpretation of the solution. When \vec{v}_1 and \vec{v}_2 are multiplied by \mathbf{V}^T , the two vectors undergo a rotation and

become unit vectors along the x- and y-axes respectively:

$$\mathbf{V}^T \vec{v}_1 = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} -\frac{1}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} \end{bmatrix} = \vec{e}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad (\text{B.14})$$

$$\mathbf{V}^T \vec{v}_2 = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} -\frac{2}{\sqrt{5}} \\ -\frac{1}{\sqrt{5}} \end{bmatrix} = \vec{e}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (\text{B.15})$$

The matrix Σ then stretches \vec{e}_1 and \vec{e}_2 :

$$\Sigma \vec{e}_1 = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \vec{e}_1 = 3 \vec{e}_1 = \sigma_1 \vec{e}_1, \quad (\text{B.16})$$

$$\Sigma \vec{e}_2 = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \vec{e}_2 = 2 \vec{e}_2 = \sigma_2 \vec{e}_2, \quad (\text{B.17})$$

Finally, $\sigma_1 \vec{e}_1$ and $\sigma_2 \vec{e}_2$ are projected onto $\sigma_1 \vec{u}_1$ and $\sigma_2 \vec{u}_2$:

$$\mathbf{U} \sigma_1 \vec{e}_1 = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix} \sigma_1 \vec{e}_1 = \sigma_1 \begin{bmatrix} \frac{2}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} \end{bmatrix} = \sigma_1 \vec{u}_1, \quad (\text{B.18})$$

$$\mathbf{U} \sigma_2 \vec{e}_2 = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix} \sigma_2 \vec{e}_2 = \sigma_2 \begin{bmatrix} \frac{1}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} \end{bmatrix} = \sigma_2 \vec{u}_2. \quad (\text{B.19})$$

Hence, \mathbf{A} projects the vectors \vec{v}_1 and \vec{v}_2 onto $\sigma_1 \vec{u}_1$ and $\sigma_2 \vec{u}_2$, respectively. It is easily verified that $\sigma_1 \vec{u}_1$ and $\sigma_2 \vec{u}_2$ define the major and minor axes of an ellipse. One can see how the matrix \mathbf{A} transforms all the other vectors which describe points on the unit circle, by realising that any other vector on the unit circle will simply correspond to a rotation of \vec{v}_1 or \vec{v}_2 by a fixed angle. Hence any point on the circle is projected onto a point on the ellipse with major and minor axes defined by $\sigma_1 \vec{u}_1$ and $\sigma_2 \vec{u}_2$.

B. SINGULAR VALUE DECOMPOSITION OF $M \times N$ MATRICES

REFERENCES

- [1] D. Sigg and the LIGO Scientific Collaboration, “Status of the LIGO detectors,” *Class. Quant. Grav.*, vol. 25, p. 114041, 2008. [1](#), [8](#), [63](#), [64](#), [65](#), [76](#), [82](#), [90](#), [92](#)
- [2] G. M. Harry and the LIGO Scientific Collaboration, “Advanced LIGO: the next generation of gravitational wave detectors,” *Classical and Quantum Gravity*, vol. 27, p. 084006, 2010. [1](#), [7](#)
- [3] K. Somiya, “Detector configuration of KAGRA the Japanese cryogenic gravitational-wave detector,” *Classical and Quantum Gravity*, vol. 29, p. 124007, 2012. [1](#)
- [4] B. Iyer, T. Souradeep, C. Unnikrishnan, S. Dhurandhar, S. Raja, A. Kumar, and A. S. Sengupta, “LIGO-India Tech. rep.,” 2011. [1](#)
- [5] J. A. et. al The LIGO Scientific Collaboration, “Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors,” *Class. Quant. Grav.*, vol. 27, p. 173001, 2010. [1](#), [2](#), [10](#), [21](#), [49](#)
- [6] J. A. et al. The LIGO Scientific Collaboration and The Virgo Collaboration, “Search for gravitational waves from compact binary coalescence in LIGO and Virgo data from S5 and VSR1,” *Phys. Rev. D*, vol. 82, p. 102001, 2010. [1](#), [10](#), [57](#)
- [7] J. A. et al. LIGO Scientific Collaboration, Virgo Collaboration, “Search for gravitational waves from low mass compact binary coalescence in LIGOs

REFERENCES

- sixth science run and Virgos science runs 2 and 3,” *Phys. Rev. D*, vol. 85, p. 082002, 2012. [1](#), [10](#), [11](#), [57](#), [71](#)
- [8] J. A. et al. The LIGO Scientific Collaboration, “Search for Gravitational Waves from Binary Black Hole Inspiral, Merger and Ringdown in LIGO-Virgo Data from 2009-2010,” *arXiv:1209.6533*, 2012. [1](#), [10](#), [11](#), [28](#), [57](#), [71](#)
- [9] J. A. et al. The LIGO Scientific Collaboration and The Virgo Collaboration, “An upper limit on the stochastic gravitational-wave background of cosmological origin,” *Nature*, vol. 460, pp. 990–994, 2009. [1](#)
- [10] J. A. et al. The LIGO Scientific Collaboration and The Virgo Collaboration, “Searching for a stochastic background of gravitational waves with the laser interferometer gravitational-wave observatory,” *The Astrophysical Journal*, vol. 659, p. 918, 2007. [1](#)
- [11] J. Taylor and J. Weisberg, “A new test of general relativity - gravitational radiation and the binary pulsar psr 1913+16,” *Ap. J.*, vol. 253, pp. 908–920, Feb. 1982. [1](#)
- [12] J. Veitch and A. Vecchio, “Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network,” *Phys Rev. D*, vol. 81, p. 062003, 2010. [2](#), [12](#), [15](#), [53](#), [64](#), [79](#), [86](#), [88](#), [90](#)
- [13] M. van der Sluys, I. Mandel, V. Raymond, V. Kalogera, C. Roever, and N. Christensen, “Parameter estimation for signals from compact binary inspirals injected into LIGO data,” *Class. Quant. Grav.*, vol. 26, p. 204010, 2009. [2](#), [15](#), [53](#), [55](#), [86](#)
- [14] D. A. Brown, J. Brink, H. Fang, J. R. Gair, C. Li, G. Lovelace, I. Mandel, and K. S. Thorne, “Prospects for detection of gravitational waves from intermediate-mass-ratio inspirals,” *Phys. Rev. Lett.*, vol. 99, p. 201102, 2007. [2](#), [21](#), [22](#)
- [15] C. L. Rodriguez, I. Mandel, and J. R. Gair, “Verifying the no-hair property of massive compact objects with intermediate-mass-ratio inspirals in

REFERENCES

- advanced gravitational-wave detectors,” *Phys. Rev. D*, vol. 85, p. 062002, 2012. [2](#), [22](#)
- [16] R. OShaughnessy, C. Kim, V. Kalogera, and K. Belczynski, “Constraining population synthesis models via empirical binary compact object merger and supernova rates,” *The Astrophysical Journal*, vol. 672, p. 479, 2008. [2](#)
- [17] B. Kiziltan, A. Kottas, and S. Thorsett, “The neutron star mass distribution,” *arXiv:1011.4291*, 2010. [2](#)
- [18] J. A. et al. The LIGO Scientific Collaboration and The Virgo Collaboration, “Search for gravitational waves from binary black hole inspiral, merger, and ringdown,” *Phys. Rev. D*, vol. 83, p. 122005, Jun 2011. [2](#)
- [19] R. Smith, I. Mandel, and A. Vecchio, “Studies of waveform requirements for intermediate mass-ratio coalescence searches with advanced detectors,” *arXiv:1302.6049*, 2012. [2](#), [17](#)
- [20] S. Farrell, N. Webb, D. Barret, O. Godet, and J. Rodrigues, “An intermediate-mass black hole of over 500 solar masses in the galaxy ESO 243-49,” *Nature*, vol. 460, p. 7251, 2009. [2](#), [21](#)
- [21] T. J. Maccarone, A. Kundu, S. E. Zepf, and K. L. Rhode, “A black hole in a globular cluster,” *Nature*, vol. 445, p. 183185, 2007. [2](#)
- [22] A. Patruno, S. Portegies Zwart, J. Dewi, and C. Hopman, “The ultraluminous X-ray source in M82: an intermediate-mass black hole with a giant companion,” *MNRAS Lett.*, vol. 370, pp. L6–L9, 2006. [2](#)
- [23] C. D. Bailyn, R. K. Jain, P. Coppi, and J. A. Orosz, “The mass distribution of stellar black holes,” *The Astrophysical Journal*, vol. 499, no. 1, p. 367, 1998. [3](#)
- [24] W. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, “The mass distribution of stellar-mass black holes,” *The Astrophysical Journal*, vol. 741, p. 103, 2010. [3](#)

REFERENCES

- [25] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Rover, and N. Christensen, “Degeneracies in Sky Localisation Determination from a Spinning Coalescing Binary through Gravitational Wave Observations: a Markov-Chain Monte-Carlo Analysis for two Detectors,” *Class. Quant. Grav.*, vol. 26, p. 114007, 2009. [3](#), [14](#), [50](#), [55](#), [86](#), [93](#)
- [26] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, “Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors,” *Phys. Rev. D*, vol. 80, p. 084043, 2009. [4](#), [6](#), [10](#), [12](#), [13](#), [17](#), [23](#), [24](#), [28](#), [29](#), [32](#), [69](#), [74](#), [81](#), [89](#), [98](#), [99](#), [103](#), [104](#)
- [27] P. C. Peters and J. Mathews, “Gravitational Radiation from Point Masses in a Keplerian Orbit,” *Phys. Rev.*, vol. 131, pp. 435–440, 1963. [4](#)
- [28] B. Schutz, *A First Course in General Relativity*. Cambridge University Press, 2009. [4](#)
- [29] M. Maggiore, *Gravitational Waves Volume 1: Theory and Experiments*. Oxford University Press, 2008. [4](#), [5](#), [6](#), [8](#), [10](#), [11](#), [26](#), [34](#), [54](#)
- [30] J. A. et al. The LIGO Scientific Collaboration and The Virgo Collaboration, “Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data,” *Phys. Rev. D*, vol. 80, p. 062001, 2009. [7](#), [11](#)
- [31] L. S. Finn, “Gravitational wave interferometer noise calculator,” 2013. [8](#)
- [32] S. Hild, “Beyond the second generation of laser-interferometric gravitational wave observatories,” *Classical and Quantum Gravity*, vol. 29, p. 124006, 2012. [9](#)
- [33] B. Allen, J. D. E. Creighton, E. E. Flanagan, and J. D. Romano, “Robust statistics for deterministic and stochastic gravitational waves in non-Gaussian noise: Frequentist analyses,” *Phys. Rev. D*, vol. 65, p. 122002, Jun 2002. [8](#)
- [34] R. Biswas, P. Brady, J. Burguet-Castell, K. Cannon, J. Clayton, A. Dietz, N. Fotopoulos, L. Goggin, D. Keppel, C. Pankow, L. Price, and R. Vaulin,

- “Detecting transient gravitational waves in non-Gaussian noise with partially redundant analysis methods,” *arXiv:1201.2964*, 2012. [8](#)
- [35] R. Biswas, P. Brady, J. Burguet-Castell, K. Cannon, J. Clayton, A. Dietz, N. Fotopoulos, L. Goggin, D. Keppel, C. Pankow, L. Price, and R. Vaulin, “Likelihood-ratio ranking of gravitational-wave candidates in a non-Gaussian background,” *arXiv:1201.2959*, 2012. [8](#)
- [36] P. Jaranowski and A. Krolak, “Gravitational-Wave Data Analysis. Formalism and Sample Applications: The Gaussian Case,” 2012. [12](#)
- [37] I. Mandel, “Parameter estimation on gravitational waves from multiple coalescing binaries,” *Phys. Rev. D*, vol. 81, p. 084029, 2010. [12](#)
- [38] E. Poisson, “Gravitational radiation from a particle in circular orbit around a black hole. vi. accuracy of the post-newtonian expansion,” *Phys. Rev. D*, vol. 52, pp. 5719–5723, Nov 1995. [13](#), [47](#)
- [39] N. Yunes and E. Berti, “Accuracy of the post-newtonian approximation: Optimal asymptotic expansion for quasicircular, extreme-mass ratio inspirals,” *Phys. Rev. D*, vol. 77, p. 124006, Jun 2008. [13](#), [47](#)
- [40] T. W. Baumgarte and S. L. Shapiro, *Numerical Relativity: Solving Einstein’s Equations on the Computer*. Cambridge University Press, 2010. [13](#)
- [41] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, “Inspiral-merger-ringdown multipolar waveforms of nonspinning black-hole binaries using the effective-one-body formalism,” *Phys. Rev. D*, vol. 84, p. 124052, Dec 2011. [13](#), [23](#), [28](#), [30](#), [82](#), [89](#)
- [42] J. R. Gair and K. Glampedakis, “Improved approximate inspirals of test-bodies into kerr black holes,” *Phys. Rev. D*, vol. 73, p. 064037, 2006. [13](#), [23](#)
- [43] S. Babak, H. Fang, J. R. Gair, K. Glampedakis, and S. A. Hughes, ““kludge” gravitational waveforms for a test-body orbiting a kerr black hole,” *Phys. Rev. D*, vol. 75, p. 024005, 2007. [13](#), [23](#)

REFERENCES

- [44] E. Huerta and J. R. Gair, “Intermediate-mass-ratio-inspirals in the einstein telescope: I. signal-to-noise ratio calculations,” *Phys Rev. D*, vol. 83, p. 044020, 2011. [13](#), [16](#), [23](#), [28](#), [44](#), [98](#)
- [45] H. Nakano, Y. Zlochower, C. O. Lousto, and M. Campanelli, “Intermediate-mass-ratio black-hole binaries: numerical relativity meets perturbation theory,” *Phys. Rev. D*, vol. 84, p. 124006, 2011. [13](#), [23](#), [48](#)
- [46] B. J. Owen, “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing,” *Phys. Rev. D*, vol. 53, pp. 6749–6761, 1996. [14](#), [88](#)
- [47] I. W. Harry, B. Allen, and B. S. Sathyaprakash, “Stochastic template placement algorithm for gravitational wave data analysis,” *Phys. Rev. D*, vol. 80, p. 104014, Nov 2009. [14](#)
- [48] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. [14](#)
- [49] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, “Singular value decomposition applied to compact binary coalescence gravitational-wave signals,” *Phys. Rev. D*, vol. 82, p. 044025, 2010. [15](#), [17](#), [50](#), [56](#), [57](#), [58](#), [59](#), [65](#), [66](#), [89](#), [98](#)
- [50] K. Cannon, C. Hanna, and D. Keppel, “Interpolating compact binary waveforms using the singular value decomposition,” *Phys. Rev. D*, vol. 85, p. 081504, 2012. [15](#), [17](#), [50](#), [56](#), [57](#), [61](#), [62](#), [65](#), [66](#), [88](#), [89](#), [90](#)
- [51] R. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel, “Towards rapid parameter estimation on gravitational waves from compact binaries using interpolated waveforms,” *arXiv:1211.1254*, 2012. [18](#), [86](#), [87](#), [88](#), [90](#), [93](#), [95](#)
- [52] S. W. Davis, R. Narayan, Y. Zhu, D. Barret, S. A. Farrell, O. Godet, M. Servillat, and N. A. Webb, “The Cool Accretion Disk in ESO 243-49 HLX-1: Further Evidence of an Intermediate-mass Black Hole,” *ApJ*, vol. 734, p. 111, June 2011. [21](#)

-
- [53] A. K. H. Kong, R. Di Stefano, and F. Yuan, “Evidence of an Intermediate-Mass Black Hole: Chandra and XMM-Newton Observations of the Ultraluminous Supersoft X-Ray Source in M101 during Its 2004 Outburst,” *ApJ. Lett.*, vol. 617, pp. L49–L52, Dec. 2004. [21](#)
- [54] H. Feng and R. Soria, “Ultraluminous X-ray sources in the Chandra and XMM-Newton era,” *New Astronomy Reviews*, vol. 55, pp. 166–183, Nov. 2011. [21](#)
- [55] G. Fabbiano, “Populations of X-Ray Sources in Galaxies,” *Annual Review of Astronomy & Astrophysics*, vol. 44, pp. 323–366, Sept. 2006. [21](#)
- [56] M. C. Miller, “Intermediate-mass black holes as LISA sources,” *Class. Quant. Grav.*, vol. 26, p. 094031, 2009. [21](#)
- [57] M. C. Miller and E. J. M. Colbert, “Intermediate-mass black holes,” *Int. J. Mod. Phys. D*, vol. 13, pp. 1–64, 2004. [21](#)
- [58] G. Harry and The LIGO Scientific Collaboration, “Advanced LIGO: the next generation of gravitational wave detectors,” *Class. Quant. Grav.*, vol. 27, p. 084006, 2010. [21](#), [49](#), [72](#)
- [59] The Virgo Collaboration, “Advanced Virgo Baseline Design note VIR027A0,” tech. rep., 2009. [21](#), [49](#)
- [60] K. Somiya, “Detector configuration of KAGRA the Japanese cryogenic gravitational-wave detector,” *Classical and Quantum Gravity*, vol. 29, no. 12, p. 124007, 2012. [21](#)
- [61] I. Mandel, D. A. Brown, J. R. Gair, and M. C. Miller, “Rates and characteristics of intermediate-mass-ratio inspirals detectable by advanced ligo,” *Astrophys. J.*, vol. 681, p. 1431, 2008. [21](#), [22](#)
- [62] E. E. Flanagan and S. A. Hughes, “Measuring gravitational waves from binary black hole coalescences. i. signal to noise for inspiral, merger, and ringdown,” *Phys. Rev. D*, vol. 57, pp. 4535–4565, 1998. [22](#)

REFERENCES

- [63] E. A. Huerta and J. R. Gair, “Influence of conservative corrections on parameter estimation for extreme-mass-ratio inspirals,” *Phys. Rev. D*, vol. 79, p. 084021, 2009. [23](#), [28](#), [31](#), [98](#)
- [64] T. Futamase and Y. Itoh, “The Post-Newtonian Approximation for Relativistic Compact Binaries,” *Living Rev. Relativity*, vol. 10, 2007. [24](#)
- [65] L. Blanchet, “Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries,” *Living Rev. Relativity*, vol. 9, 2006. [24](#)
- [66] I. Mandel and J. R. Gair, “Can we detect intermediate mass ratio inspirals?,” *Class. Quant. Grav.*, vol. 26, p. 094036, 2009. [24](#)
- [67] C. Cutler and E. Flanagan, “Gravitational waves from merging compact binaries: How accurately can one extract the binaries parameters from the inspiral waveform?,” *Phys. Rev. D*, vol. 49, p. 2658, 1994. [27](#), [32](#)
- [68] T. Damour, “Coalescence of two spinning black holes: An effective one-body approach,” *Phys. Rev. D*, vol. 64, p. 124013, Nov 2001. [28](#)
- [69] T. Damour, P. Jaranowski, and G. Schäfer, “Effective one body approach to the dynamics of two spinning black holes with next-to-leading order spin-orbit coupling,” *Phys. Rev. D*, vol. 78, p. 024009, Jul 2008. [28](#)
- [70] T. Damour and A. Gopakumar, “Gravitational recoil during binary black hole coalescence using the effective one body approach,” *Phys. Rev. D*, vol. 73, p. 124006, 2006. [30](#)
- [71] A. Buonanno and T. Damour, “Transition from inspiral to plunge in binary black hole coalescences,” *Phys. Rev. D*, vol. 62, p. 064015, 2000. [30](#)
- [72] A. Buonanno, Y. Pan, J. G. Baker, J. Centrella, B. J. Kelly, S. T. McWilliams, and J. R. van Meter, “Toward faithful templates for non-spinning binary black holes using the effective-one-body approach,” *Phys. Rev. D*, vol. 76, p. 104049, 2007. [30](#)
- [73] L. Blanchet, G. Faye, B. R. Iyer, and S. Sinha, “The third post-newtonian gravitational wave polarizations and associated spherical harmonic modes

- for inspiralling compact binaries in quasi-circular orbits,” *Classical and Quantum Gravity*, vol. 25, p. 165003, 2008. [32](#)
- [74] The LIGO Scientific Collaboration, “Tech. Rep. LIGO-T0900288-v3 LIGO Project,” 2009. [32](#), [33](#)
- [75] B. S. Sathyaprakash, B. F. Schutz, and C. V. D. Broeck, “Cosmography with the einstein telescope,” *Classical and Quantum Gravity*, vol. 27, no. 21, p. 215006, 2010. [34](#)
- [76] T. A. Apostolatos, “Search templates for gravitational waves from precessing, inspiraling binaries,” *Phys. Rev. D*, vol. 52, pp. 605–620, Jul 1995. [38](#)
- [77] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, “Findchirp: An algorithm for detection of gravitational waves from inspiraling compact binaries,” *Phys. Rev. D*, vol. 85, p. 122006, 2012. [39](#)
- [78] T. G. F. Li, W. D. Pozzo, S. Vitale, C. V. D. Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio, “Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence,” *Phys. Rev. D*, vol. 85, p. 082003, 2012. [49](#)
- [79] I. Mandel and R. O’Shaughnessy, “Compact binary coalescences in the band of ground-based gravitational-wave detectors,” *Class. Quant. Grav.*, vol. 27, p. 114007, 2010. [49](#)
- [80] T. Bulik and K. Belczynski, “Constraints on the binary evolution from chirp mass measurements,” *ApJ*, vol. 589, p. L37, 2003. [49](#)
- [81] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, “Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries,” *Phys. Rev. D*, vol. 49, pp. 6274–6297, Jun 1994. [52](#)
- [82] B. A. Berg, *Markov Chain Monte Carlo Simulations and Their Statistical Analysis*. Singapore, World Scientific, 2004. [54](#), [55](#)

REFERENCES

- [83] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller, “Equation of state calculations by fast computing machines,” *Journal of Chemical Physics*, vol. 21, no. 6, p. 10871092, 1953. [56](#)
- [84] W. Hastings, “Monte carlo sampling methods using markov chains and their applications,” *Biometrika*, vol. 51, no. 1, p. 97109, 1970. [56](#)
- [85] The LIGO Scientific Collaboration, “LSC Algorithm Library (LAL),” 2013. [56](#), [89](#)
- [86] A. Gelman and D. B. Rubin, “Inference from iterative simulation using multiple sequences,” *Statistical Science* **7**, 4, 457-472 (1992), vol. 7, pp. 457–472, 1992. [64](#), [79](#)
- [87] P. Ajith, “Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins,” *Phys. Rev. D*, vol. 84, p. 084037, Oct 2011. [65](#), [72](#)
- [88] L. Lindblom, B. J. Owen, and D. A. Brown, “Model waveform accuracy standards for gravitational wave data analysis,” *Phys. Rev. D*, vol. 78, p. 124020, 2008. [66](#)
- [89] F. Ohme, “Analytical meets numerical relativity - status of complete gravitational waveform models for binary black holes,” *Class. Quant. Grav.*, vol. 29, p. 124002, 2012. [66](#), [67](#), [69](#), [74](#), [78](#), [81](#), [99](#)
- [90] W. J. Conover, *Practical Nonparametric Statistics*. John Wiley and Sons, 1971. [67](#)
- [91] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, “Prototype effective-one-body model for nonprecessing spinning inspiral-merger-ringdown waveforms,” *Phys. Rev. D*, vol. 86, p. 024011, Jul 2012. [72](#)
- [92] B. A. et al., “Testing gravitational-wave searches with numerical relativity waveforms: results from the first Numerical INjection Analysis (NINJA) project,” *Classical and Quantum Gravity*, vol. 26, p. 165008, 2009. [74](#)

REFERENCES

- [93] “NINJA BBH Project 2 Home Page,” 2013. [74](#)
- [94] J. Skilling, “Nested sampling,” *AIP Conf. Ser. Vol. 735*, p. 395, 2004. [86](#)
- [95] The LIGO Scientific Collaboration and The Virgo Collaboration, “Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network,” *arXiv:1304.1775*, 2013. [90](#)
- [96] S. Field, C. Galley, F. Herrmann, J. Hesthaven, E. Ochsner, and M. Tiglio, “Reduced basis catalogs for gravitational wave templates,” *Physical Review Letters*, vol. 106, no. 22, p. 221102, 2011. [93](#)
- [97] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, “Gravitational wave parameter estimation with compressed likelihood evaluations,” *arXiv:1304.0462*, 2013. [93](#), [94](#)
- [98] S. Mitra, S. V. Dhurandhar, and L. S. Finn, “Improving the efficiency of the detection of gravitational wave signals from inspiraling compact binaries: Chebyshev interpolation,” *Phys. Rev. D*, vol. 72, p. 102001, Nov 2005. [94](#)
- [99] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jung-man, P. Alken, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual, 3rd ed., for version 1.12*. Network Theory Ltd, United Kingdom, 2009. [105](#)
- [100] G. Golub and C. Van Loan, *Matrix Computations*. The Johns Hopkins University Press (3rd edition), 1996. [105](#), [108](#)
- [101] J. Demmel, *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997. [108](#)