Davies, Gareth Stephen (2015) *Bayesian searches for continuous gravitational waves in the frequency domain.* PhD thesis.

Bayesian searches for continuous gravitational waves in the frequency domain

Gareth Stephen Davies

MPhys

Submitted in fulfilment of the requirements

for the Degree of Doctor of Philosophy

School of Physics & Astronomy

College of Science & Engineering

University of Glasgow

August 2015

Abstract

This thesis concerns the analysis of continuous gravitational waves from neutron stars with non-axisymmetric rotational motion using data from ground-based interferometric gravitational wave detectors, the development of a computationally efficient algorithm for analysis of this data and the use of this algorithm in followup searches, which were previously too computationally expensive to consider.

Presented in this thesis is a new, computationally efficient method for down sampling data from gravitational-wave detectors for use in targeted searches for continuous waves from rapidly rotating neutron stars, which we call Spectral Interpolation, or *SplInter*. This method is tested and compared to the existing analysis in various situations, including an end-to-end pipeline comparison utilising hardware injections. The limits of the application of the SplInter algorithm are explored, including a study into its use in the analysis of continuous waves from neutron stars in binary systems.

Next, a search is presented for signals from known pulsars in LIGO science run 6 data, with a comparison to a similar, previously performed analysis. This search produces upper limits on the amplitude of gravitational waves from the sources which are compared to the previous analysis and observational limits.

The first Bayesian follow up search for candidate continuous gravitational-wave signals from all-sky and directed searches is presented. These are also the first searches using the Bayesian pipeline to search in more than one frequency parameter, and over sky position. We search for a prospective signal from a directed search for a possible neutron star, and for possible signals from a deep all-sky search pointed along the Orion spur.

The results of this work and the future of these algorithms and searches are discussed, with emphasis on enhancing the algorithm so that it is applicable to all sources.

Contents

List of Tables

List of Figures

Acknowledgements

> *I don't know half of you half as well as I should like; and I like less than half of you half as well as you deserve.*
>
> Bilbo Baggins, Lord of the Rings, J.R.R Tolkien

There are many people who deserve a lot of credit and thanks for their help during my time at the University of Glasgow, both directly in my work and indirectly for encouragement and support. I shall try to thank those who deserve it without going over the top – after all it was still me who did the work.

First and foremost I thank my supervisor, Graham Woan, whose guidance, encouragement, support and understanding have been vital for my work and allowed me to explore projects without getting too hung up on details.

Next, I'd like to thank Siong Heng, my second supervisor, for providing general advice and helping me to socialise at conferences. I wish that I had been able to realistically and sensibly use Siong's suggestion to name the Spectral Interpolation algorithm 'SpIDERMAN' (SPectral Interpolation for Downsampling, Efficent Recovery in the Manipulation of Analysis of Neutron stars). I'd also like to thank my masters project supervisor, David Burton at Lancaster University, for introducing me to the world of gravitational waves, and helping me to begin to pretend to understand general relativity.

There aren't any sufficient words to convey the debt of gratitude I owe to Matt Pitkin for his help during my PhD. Barely a day went by without me asking him for advice on programming, statistics, physics, writing, presentations or simple life skills.

He has been a mentor, office mate and a font of wisdom and knowledge, but above all he has been a friend. I know I'm not the first PhD student at Glasgow that has been helped massively by Matt and I certainly won't be the last.

I'd like to thank my fellow students in the Institute for Gravitational Research at Glasgow, many of whom have been inspirations in terms of the hard work and dedication they have put into their work, and many of whom have provided me with fantastic memories of my time as a PhD student. Remembering the look on Ignacio (Nacho) Santiago Prieto's face after he passed his viva has kept me going when I've wondered if the effort was worth it. Josh Logue has been wonderful in welcoming me to Glasgow, even if I had to explain to more than a few people at conferences over the years that we were not the same person. Craig Lawrie, Colin Gill, Mervyn Chan and Brynley Pearlstone have always been great to talk to. A special thanks also goes to Yiming Hu, one of the happiest, hardest working and most intelligent people I have ever met.

In particular I'd like to thank Erin MacDonald and Jade Powell for being excellent office mates, Jade is one of the hardest workers I have ever seen, and puts me to shame, and Erin's consistent optimism and passion for science was infectious.

Chris Messenger has been a great source of knowledge, without whom I would never have been able to begin to think I could understand the $\mathcal{F}$ statistic, and I have had quite a few interesting conversations with Xilong Fan over the years.

Outside of the data analysts, I'd like to thank Neil Gordon, Jan Hennig and Sean Leavey for a wonderful post-conference road trip in California. Thanks also to Morag Casey, for giving me the chance to tutor undergraduate students and to all of my fellow tutors.

To all of the people above, I say this: it has been wonderful getting to know you over the years, thank you for your help and encouragement and I wish you all the best for your future.

As an addition in my corrections, I would like to thank Martin Hendry and Patrick Sutton for examining my thesis, and Ken Strain for convening my viva. Thank you very much for making my viva an oddly enjoyable experience.

Outside of Glasgow, I want to thank the continuous-waves group in the LIGO and

Virgo Scientific collaborations, particularly Ra Inta and Vladimir Dergachev for your help in providing the followup candidates, and Reinhard Prix and Damir Buskulic for your suggestions in the review of the Spectral Interpolation algorithm.

My parents, Eric and Hilary Davies have been a constant support throughout my life and have always helped but not pushed. They first piqued my interest in science, and have always encouraged my curiosity for how the world works.

Thanks to the rest of my family, both biological and otherwise, for their support throughout my PhD. My sisters, Megan and Deborah, were clever enough and talented enough to allow me to go about my business without any weight of expectation. Sarah's parents, brothers and grandparents have all been there to help and support us throughout our time here, for which I will always be grateful.

Without my amazing partner Sarah Jane Cabourn, I would not be here, in fact I would barely be anywhere. Sarah is my partner in the truest sense of the word, my best friend and my better half. I could not list all that Sarah has done for me without at least another 195 pages. Sarah, thank you so much for everything.

In memory of my grandparents, Edna, Ethel, Fred and Roy.
I wish you could have been here to see this.
I hope that you would have been proud of me.

Also in memory of Lawrence Harold Cabourn,
who has been like a grandfather to me
21st February 1929 – 22nd October 2015.

Author's Declaration

It doesn't matter if it's been said, it's never been said by me.

Josh Homme, Queens of the Stone Age, 2009

I, Gareth Stephen Davies, confirm that the work presented in this thesis is my own. Where information or figures have been derived from other sources, appropriate references have been given.

Chapter 1 gives an introduction to the concept of gravitational waves, a brief derivation of the gravitational-wave equation from general relativity, gravitational-wave sources and detectors. All figures are either my own or appropriately referenced. Information comes from referenced sources.

Chapter 2 discusses and introduces Bayesian formalism and parameter estimation, as well as more specific discussions on continuous gravitational waves and their analysis. The only figure is my own, containing data from the sources specified. Information in the text is cited as appropriate.

Some of the work in chapters 3 and 4 is being prepared for publication in [34]. Figures 3.3.1 to 3.3.3, 3.4.3 and 4.1.17, were created by myself for use in both this thesis and the planned publication, and are used with permission of the other authors. Figures 4.1.6, 4.1.9, 4.1.22, 4.1.23 and 4.2.1 are also created by myself and parts of which or slightly different versions of which are to be utilised in the mentioned paper, reproduced with permission of the other authors.

Chapter 3 introduces continuous waves in the frequency domain, the Spectral Interpolation method and algorithm, and explains some of the details with regard to the implementation of the algorithm. Information is cited as appropriate, and figures are all my own.

Chapter 4 presents tests performed to compare the Spectral Interpolation algorithm to the heterodyne routine. Citations are made as appropriate, and all figures are my own.

Chapter 5 contains results of a search for known pulsars in LIGO data, and includes tables (5.2 and 5.3) indicating $h_0^{95\%}$ upper limits from the heterodyne and MCMC search (as in [2]) and from Spectral Interpolation and nested sampling, which are my own work. The heterodyne and MCMC results have been reproduced with permission of the main author. All other results and discussion are my own.

Chapter 6 describes a follow up analysis of candidate continuous-wave signals from all-sky and directed searches. These candidates come from [6, 52], and are used with permission of the authors. All figures are my own and information has been referenced as appropriate.

The conclusions given in chapter 7 are my own analysis of the work and results in this thesis. The future work stated for the Spectral Interpolation algorithm is an interpretation of discussions between myself, Graham Woan, Matt Pitkin and Reinhard Prix. The rest of the discussion of the future prospects of the work is my own.

The appendix contains discussions and figures which are relevant to the discussion in the main text, but have been cut out for brevity or aesthetics. All figures are my own and citations are given where relevant.

I would like to make a quick note with regard to the writing style in this thesis. Throughout this work, the plural pronoun ‘we’ is utilised, this is in order to indicate a collaborative effort between the author and the reader, rather than stating that work was carried out in by others within a collaborative effort. On the occasion where others have performed work, this has been stated explicitly both above and in the text.

Chapter 1

Gravitational waves: theory, sources and detectors

I have forgotten much that I thought I knew, and learned again much that I had forgotten.

Gandalf, The Lord of the Rings, J.R.R. Tolkien

Gravitational Waves are ripples in the curvature of spacetime, predicted in 1916 by Einstein as an outcome of the General Theory of Relativity (GR) [40], and the direct detection of gravitational waves is one of the final pieces in the puzzle of evidence that a GR-like mechanism exists as the progenitor of gravitational forces. Section 1.1 gives an outline of the derivation of gravitational waves from GR and the way in which they interact with matter.

Gravitational waves (GWs) offer an exciting prospect for seeing the Universe in an entirely new way, independent of the electromagnetic and astroparticle observations we have been able to make up to now. The prospects are intriguing, from seeing further back into the very early stages of the universe [16, 87] to providing a test bed for competing theories of gravitation which – up until this point – will have been experimentally indistinguishable from one another [102], probing the inside of exotic compact objects [63] as well as the prospect of making new, unexpected, discoveries similar to the advent of radio astronomy leading to the discovery of pulsars. Section 1.2 describes and reviews the types of sources expected to be seen in the early years of GW astronomy.

The observation of gravitational waves will usher in a new era of multi-messenger astronomy, providing more evidence in support of observations which have already been made, and providing triggers for observations of astronomical events which would have been observed previously only due to extreme luck. Gravitational waves interact only weakly with matter, meaning that they can travel through objects which are opaque to electromagnetic observation and have an advantage in that they remain unretarded by the interstellar medium; rapid sky localisation of a GW source can therefore provide electromagnetic astronomers with a trigger to observe an event.

Gravitational waves' weak interaction with matter also makes them extremely difficult to detect. Section 1.3 details a brief outline of the experiments designed to perform this task and section 1.3.1 specifically gives more detail on the interferometric detectors, including the main noise sources and the ways in which the GW signal is detected and analysed.

1.1 Gravitational waves as a consequence of general relativity

In this section, we outline the derivation of the GW equation from Einstein's field equations in GR, and how GWs interact with matter. This section shall only introduce the concepts which are important to the work in this thesis, and more comprehensive descriptions are available in numerous books and papers (such as [40, 86]).

1.1.1 Obtaining and solving the wave equation

The wave equation for GWs can be found by linearising the spacetime metric, this means that we approximate the metric as a flat metric plus a small, linear perturbation.

The Einstein equation from GR in tensor notation is given by (e.g. [86])

$$G_{ab} = R_{ab} - \frac{1}{2} g_{ab} R = -\frac{8\pi G}{c^4} T_{ab}, \tag{1.1.1}$$

where the Einstein tensor G_{ab} is given in terms of the spacetime metric, g_{ab}, the Ricci tensor R_{ab} and Ricci scalar R. These are defined by the Riemann curvature tensor R^a_{bcd} through the relationships $R_{ab} = R^c_{abc}$ and $R = g^{ab} R_{ab}$ respectively[1]. a and b run from 0 to 3 and denote the different dimensions, 0 for time and 1, 2 and 3 for the spatial dimensions. T_{ab} is the stress-energy-momentum tensor. For simplicity, we ignore the cosmological constant, Λ, in this discussion.

The spacetime metric g_{ab} of linearised gravity comes from a flat spacetime η_{ab} and a small linear perturbation h_{ab},

$$g_{ab} = \eta_{ab} + h_{ab}, \tag{1.1.2}$$

where η_{ab} is defined as

$$\eta_{ab} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{1.1.3}$$

The perturbation h_{ab} and its derivatives are considered to be much smaller than the flat metric (i.e. $|h_{ab}| \ll 1$) such that we can safely ignore non-linear combinations of these. The Riemann curvature tensor R^a_{bcd} is therefore[2]

$$R^a_{bcd} = \frac{1}{2} \eta^{ae} \left[\partial_c \partial_b h_{de} - \partial_c \partial_e h_{bd} - \partial_d \partial_b h_{ce} + \partial_d \partial_e h_{bc} \right], \tag{1.1.4}$$

where $\partial_a = \frac{\partial}{\partial x^a}$, the partial derivative with respect to dimension x^a.

By utilising gauge transformations, we can mould these equations into something more workable. Considering the gauge transformation $h_{ab} \rightarrow h'_{ab} = h_{ab} - \partial_a \xi_b - \partial_b \xi_a$, where ξ is an arbitrary function, but its derivatives are of the same order of magnitude

---

[1] We are using the standard notation in which we sum over repeated indices, amd we can raise or lower indices using the spacetime metric, $x^a = g^{ab} x_b$.

[2] Appendix A contains some of the maths that we skip over in this introduction.

as h_{ab}, the Einstein equation remains valid for h'_{ab} as it did with h_{ab} as all of the terms containing ξ cancel by symmetry[3]. We can use this to set gauge conditions, where a physical result still stands but we use a different frame of reference to simplify the problem.

The first condition set is the harmonic gauge:

$$\partial_a \bar{h}^a_b = 0, \tag{1.1.5}$$

given $\bar{h}^a_b = h^a_b - \frac{1}{2}\eta^a_b h$, the trace reverse of the perturbations, where $h = h^a_a$. When we set this condition on the Riemann curvature tensor, the Ricci tensor becomes[4]

$$R_{ab} = \frac{1}{2}\Box h_{ab}, \tag{1.1.6}$$

where $\Box$ is the *d'Alembertian* or *wave operator*, $\Box = \partial_a \partial^a$. We then use this R_{ab} in equation 1.1.1 to obtain the wave equation

$$\Box \bar{h}_{ab} = -\frac{16\pi G}{c^4} T_{ab}. \tag{1.1.7}$$

In free space ($T_{ab} = 0$), solutions to this have the form

$$\bar{h}_{ab} = A^{ab} \exp\left(i k_d x^d\right), \tag{1.1.8}$$

where A^{ab} is a constant, second order, symmetric amplitude tensor and k_d is the wave vector $k^d = (\frac{\omega}{c}, \vec{k})$. Solving the wave equation given this $\bar{h}_{ab}$ gives

$$k_d k^d A^{ab} \exp\left(i k_d x^d\right) = 0, \tag{1.1.9}$$

the non-trivial solutions to which are $k_d k^d = 0$. This is also the condition for a particle to be light-like ($c^2 \Delta t^2 = \Delta r^2$) which indicates that GWs propagate at the speed of light.

---

[3] See footnote 2.

[4] See footnote 3.

1.1.2 Wave polarisation and the transverse traceless gauge

As a 4×4 symmetric matrix, the amplitude tensor A^{ab} initially has 10 degrees of freedom. Within the harmonic gauge, this number is reduced, as $\partial_a \bar{h}^{ab} = 0$ leads to the non-trivial solution of $k_a A^{ab} = 0$. Using an example of a wave travelling in the z (or x^3) direction, $k_1 = k_2 = 0$, $k_0 = \omega$ and $k_3 = k$, this leads to $A^{a0} = A^{a3}$, and by symmetry $A^{0a} = A^{3a}$. Using gauge freedom we can set $A^{00} = 0$, leading to $A^{a0} = A^{a3} = A^{0a} = A^{3a} = 0$, meaning the perturbation is transverse (i.e. there is no perturbation along the direction of travel). One can also utilise the gauge freedom to demand that $h = 0$, the traceless condition[5], and the perturbation becomes:

$$h_{ab} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{1.1.10}$$

where h_+ and $h_\times$ are the plus and cross polarisations of the gravitational wave respectively. When $h_\times$ is zero, we say that the wave is plus polarised, and when h_+ is zero, we say that it is cross polarised. The exact form of h_+ and $h_\times$ will vary depending upon the source of the gravitational wave.

1.1.3 Interaction with matter

In order to see how a gravitational wave affects matter, one needs to consider the coordinate invariant proper distance between particles

$$L = \int_P \sqrt{g_{ab}\,\mathrm{d}x^a\,\mathrm{d}x^b}. \tag{1.1.11}$$

So for a GW propagating in free space, $T_{ab} = 0$, the metric, $g_{ab} = \eta_{ab} + h_{ab}$ gives

$$L = \int_P \sqrt{-\,\mathrm{d}t^2 + (1 + h_+)\,\mathrm{d}x^2 + (1 - h_+)\,\mathrm{d}y^2 + 2h_\times\,\mathrm{d}x\,\mathrm{d}y + \,\mathrm{d}z^2}. \tag{1.1.12}$$

---

[5] When we set $h = 0$, we can also immediately see that $\bar{h}_{ab} = h_{ab}$.

We see that for a plus polarised wave travelling in the z direction at a given time (i.e. $\mathrm{d}t = 0$, $h_\times = 0$) for two particles lying upon the x axis, $\mathrm{d}z = \mathrm{d}y = 0$, then $\mathrm{d}x$ will stretch by a factor of $\sqrt{1+h_+}$. If we consider two particles on the y-axis ($\mathrm{d}x = \mathrm{d}z = 0$) then $\mathrm{d}y$ is stretched by $\sqrt{1-h_+}$, leading to an orthogonal stretching and squeezing motion.

This motion and the cross terms are visualised in figure 1.1.1 for a sinusoidal GW, which also shows that the plus and cross polarisation cases are identical but with a $\pi/4$ rotation about the direction of propagation. This rotational symmetry can be seen mathematically by applying a rotation matrix with an angle of $\pi/4$ to h_{ab}. A GW will generally be a linear combination of a plus and cross polarised wave.

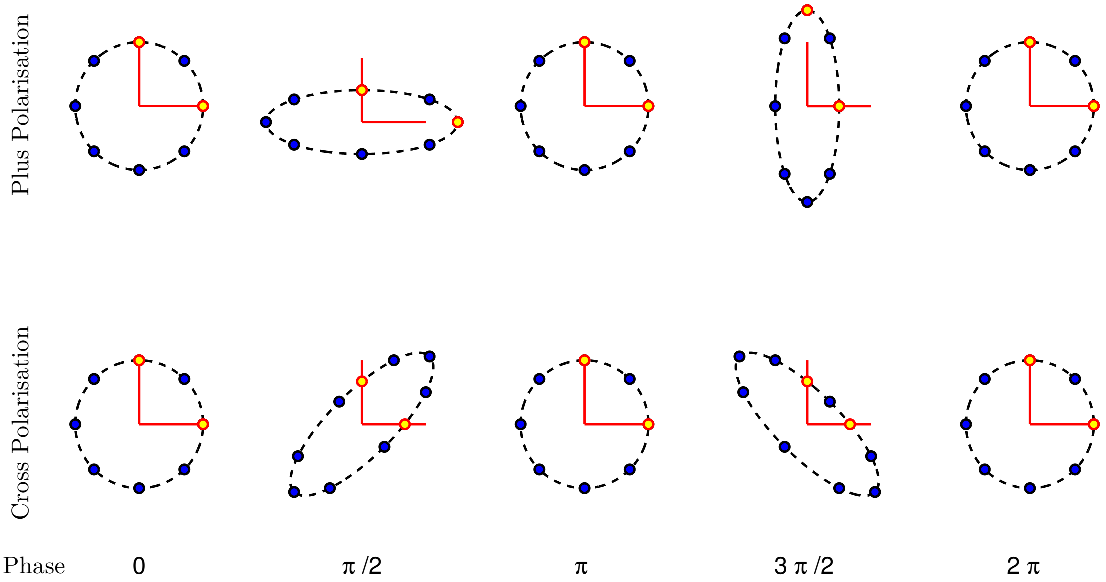

Figure 1.1.1: A series of snapshots of a ring of test particles in the xy plane which is stretched and squeezed by a gravitational wave coming out of the page in the z direction, shown are plus polarised (top) and cross polarised (bottom). The red circles indicate the position of the detector mirrors in an interferometric GW observatory, with the red lines showing the detector positions for when there are no GWs, or the phase is zero.

1.2 Gravitational-wave sources and associated analyses

Gravitational waves are emitted by any motion with a non-zero second derivative of the mass quadrupole term, and an approximation of the amplitude of the wave is [87]

$$h_{ab} = \frac{2}{d}\frac{\mathrm{d}^2 Q_{ab}}{\mathrm{d}t^2}, \tag{1.2.1}$$

where Q_{ab} is the quadrupole moment and d is the distance from the source to the detector. However only the most massive, violent, events generate signals with realistic prospects of detection. Man-made sources are ruled out; two 10^3 kg masses rotating at 10 Hz on each end of a 10 m beam would produce GW strains of $h \sim 5 \times 10^{-43}$ [87]. This strain is almost twenty orders of magnitude less than even the most sensitive current concepts for detectors [19].

Gravitational-wave sources with genuine prospects for detection tend to be split into four distinct categories, each with independent associated data analysis techniques. The following sections describe some of the most promising candidates for GW sources.

Burst sources

Burst sources are any short duration, unmodelled source of astrophysical origin, and are generally found by detecting excesses of power within the detectors, *triggers*.

The main difficulty with this search is that there can be occasional excesses of power due to the noise of the detector, so a strong understanding of the detector and its contributing noise components is required. To estimate the background rate from the noise excesses in detectors, *time slides* are used [14, 99], where the triggers from one detector are shifted by an amount of time greater than the wave travel time between the detectors.

By shifting the detectors in time there will no longer be any correlated signal present, and the correlation technique can be use to estimate a *false alarm rate* for a particular signal-to-noise ratio (SNR) of signal. From this false alarm rate, a significance can be calculated, if a standard new discovery threshold of 5σ is used, this would correspond to a single false alarm in a million gravitational-wave networks running for the observation time considered.

Sky localisation is performed by calculating the difference between the times of arrival in the different detectors to triangulate a location in the sky [41]. For two detectors the best estimate would provide a ring on the sky which would be perpendicular to the straight line between the detectors, three detectors would provide two possible locations (each is a reflection of the other in the plane passing through all of the detector sites) and so four detectors are required to provide a single, definitive, location. Astrophysical event rates are unknown, due to the unknown nature of the source, but burst pipelines may in fact be able to detect compact binary coalescence signals with high enough SNR.

Compact binary coalescence

The type of signal often considered most likely to yield the first detection is gravitational waves coming from a compact binary coalescence (CBC) involving two compact objects inspiralling and merging into one. These will most likely be a combination of neutron stars (NS) and black holes (BH).

Analysis of these signals generally involves comparison of associated waveforms from the different sources with the data using matched filtering techniques. The outcome of this is similar to the burst search, giving a list of triggers, for which the time slide technique is utilised as above to work out false alarm rates and significance estimation [18, 99].

Sky localisation techniques are also similar, and will utilise time of arrival differences in the same way as burst sources mentioned before [43]. Expected event rates vary for each source type, and estimated rates with the four-detector configuration of advanced detectors (LIGO detectors at Hanford and Livingston, Virgo and a new LIGO detector in India) at design sensitivity range from 0.4 to 400 per year for coalescing binary NSs, 0.4 to 1000 per year for binary BHs, and 0.2 to 300 per year for NS-BH binaries [10]. The wide ranges of the expected observed event rates are due to the lack of knowledge of the astrophysical populations and hence event rates, meaning that even the non-detection of the signals can lead to interesting astronomy.

Continuous waves from rapidly rotating neutron stars

Continuous-wave sources and their associated data analysis techniques are to be the main focal point of this thesis, and will therefore be covered much more thoroughly later on, particularly in chapter 2, but in keeping with the theme of this introduction, a brief explanation of this source and analyses follows.

Rapidly rotating neutron stars with a non-axisymmetric distortion are expected to emit continuous waves, with almost monochromatic signals modulated by the detector motion in orbit around the solar system barycentre (SSB) and any binary motion of the source. For signals from a non-precessing triaxial neutron star this is expected to be at twice the rotational frequency.

Continuous-wave source analyses can have two specific advantages over analyses for other gravitational-wave sources: The first is that we may already know many of the source parameters such as frequency and sky location to a high degree of accuracy from electromagnetic observations; this is known as a *targeted* search, and this type of search is the main focus of chapters 2, 3 and 4. The second advantage is that as the signal has a long duration, longer integration times lead to a deeper search into the detector noise.

Continuous-wave searches also include *directed* searches, where the sky position of the source is well constrained but other source parameters such as the frequency, frequency evolution and any binary parameters may be unknown; this generally searches for a NS with unknown parameters or in a direction at which we would expect large numbers of sources – such as the galactic centre [1], a globular cluster [7] or the Orion spur [6]. *All-sky* searches, including the Einstein@Home search [3] can be computationally expensive, and will often take a grid-based approach, matching templates for frequency and sky position. Searches for sources with unknown frequency and sky position parameters will be discussed in more detail in chapter 6.

Sky localisation in all-sky continuous-wave searches generally comes from a fit to the phase and amplitude evolution of the signal being affected by the antenna response throughout the day (see section 1.3.2), so unlike other searches it is possible to get good accuracy in position from only one detector.

Stochastic background

It is unlikely that all gravitational-wave sources will be distinguishable from one another, and as such, there will be what seems like a random gravitational-wave field coming from these as well as signals from cosmological sources such as signatures of inflation, similarly to the cosmic microwave background. From these inseparable sources, there will be a background which can be treated as statistical noise, which can be separated from instrumental noise through methods such as cross correlation between detectors [12].

Stochastic gravitational waves are expected to be seen particularly in space-based and future generation detectors, as the better sensitivity at lower frequencies means that the galactic white dwarf binary population is within the range of detection [16].

1.3 Gravitational-wave detectors

> *Never fire a laser at a mirror*
>
> Larry Niven

There are numerous gravitational-wave detectors in operation around the world, falling into a few distinct categories[6].

Interferometric detectors are currently seen as one of the most promising candidates to yield the first detection and go on to produce astrophysically interesting data. There have been multiple attempts at gravitational-wave detection using this method, including three detectors run by LIGO, as well as GEO 600, Virgo and TAMA 300 [29, 51, 68, 74] and multiple smaller ‘test’ interferometers. At the time of writing the LIGO detectors, Virgo and TAMA 300 have been decommissioned and only GEO 600 is operational in an ‘astrowatch’ capacity, while Advanced LIGO, Advanced Virgo and KAGRA [5, 20, 60] are under construction. Proposals for future for interferometer detectors include ET (the Einstein Telescope [19]), another LIGO detector in India [53], and two space-based detectors, NGO [33, 42] and DECIGO [61].

Another way to detect gravitational waves is using Doppler tracking in pulsar timing arrays [49], which uses radio data from well-known millisecond pulsars to detect

---

[6] [81] contains much more detail on gravitational-wave detectors, their history and their use.

the shifts caused by gravitational waves in the times of pulse arrivals. Finally, resonant mass antennas, including resonant bars [24] and spheres (e.g. MiniGRAIL [35] and Mário Schenberg [22]), have also been used to attempt to detect gravitational waves. These work by the gravitational wave exciting particular resonance modes in the detector, and measuring the way in which the mass responds, though these are intrinsically extremely narrow-band detectors.

So far, all of these methods have been unsuccessful in yielding the first detection of gravitational waves.

The following sections discuss current and future interferometric detectors, particularly their operation and noise sources in section 1.3.1 and the antenna pattern which affects the way in which we see sources in section 1.3.2.

1.3.1 Interferometric gravitational-wave detectors

Work on the use of interferometers as gravitational-wave detectors has been carried out since 1966 [73], and once in operation, the second generation detectors will provide the best prospect of gravitational-wave detection and analysis [72]. The basic form of an interferometer works by using a laser shining onto a beamsplitter which is at an angle[7] with respect to the beam direction.

This beamsplitter causes half of the light to be transmitted towards one mirror, and half to be reflected towards another. Each part of the beam then travels along the respective interferometer arms to a mirror, at which point it is reflected back to the centre. The light is then incident on the beamsplitter again, and the light is again reflected or transmitted. This leads to the beam being recombined and travelling either towards a photodetector or back toward the laser.

At this stage the light travelling toward the photodetector will undergo constructive or destructive interference, depending on the difference in the lengths of the interferometer arms and the laser wavelength. The photodetector then measures the varying amount of incident light, which is indicative of the amount of interference, meaning it is possible to measure the change in the difference in length of the interferometer arms over fractions of a wavelength. In LIGO detectors, the mirrors are positioned such that

---

[7] For initial and advanced detectors this angle is approximately $\pi/4$, see [23].

the interference is *almost* entirely destructive, such that the optimal signal-to-noise ratio is obtained.

Optical layout

There are many optical tricks to be played with the interferometer layout in order to improve the sensitivity, which are illustrated in figure 1.3.1. The first is to introduce Fabry-Pérot cavities [81], these reflect the beam back and forth between the end and input test masses. These mean that the light is kept in the detector for a much longer timescale, increasing *storage time* and giving a stronger gravitational-wave signal, by raising the equivalent laser power.

LIGO[8] also utilises a power recycling mirror in order to enhance the laser power, which is particularly useful when operating at a dark fringe, where the interference is completely destructive. If we consider the conservation of energy; when the interference is almost entirely destructive, then there is a very low power going towards the photodetector, this energy has to go somewhere, and so almost all of it goes back towards the laser. By using the power recycling mirror, this light is sent back into the interferometer and effectively increases the laser power.

Also shown in figure 1.3.1 is the signal recycling mirror (SRM), this sends the signal back into the interferometer, increasing the storage time of the signal in the interferometer and setting up a resonant cavity, leading to the interferometer becoming more sensitive at the cavity's resonant frequency. The resonant frequency can be tuned by changing the position and transmittance of the SRM.

The interference at which the mirrors are held has been described so far as being *almost* entirely destructive, the reason for this 'almost' is the way in which the gravitational-wave data is collected and turned into a strain. A gravitational wave will affect the interferometer readout by introducing a sideband of frequency $f_s = f_{\text{laser}} + f_{\text{gw}}$, where f_{gw} is of order $\sim 10^2\,\text{Hz}$ and f_{laser} is $\sim 10^{14}\,\text{Hz}$. The dark fringe is specifically designed in terms of f_{laser}, and so this change to f_s introduces a change in incident light on the photodiode.

---

[8] LIGO has appeared in many forms, so this footnote is an attempt to clarify the nomenclature. Initial LIGO means the LIGO detector in operation from 2002 to 2007, enhanced LIGO ran from 2009 to 2010, and together these are referred to as LIGO. Advanced LIGO's first observating run (O1) is due to start in September 2015.

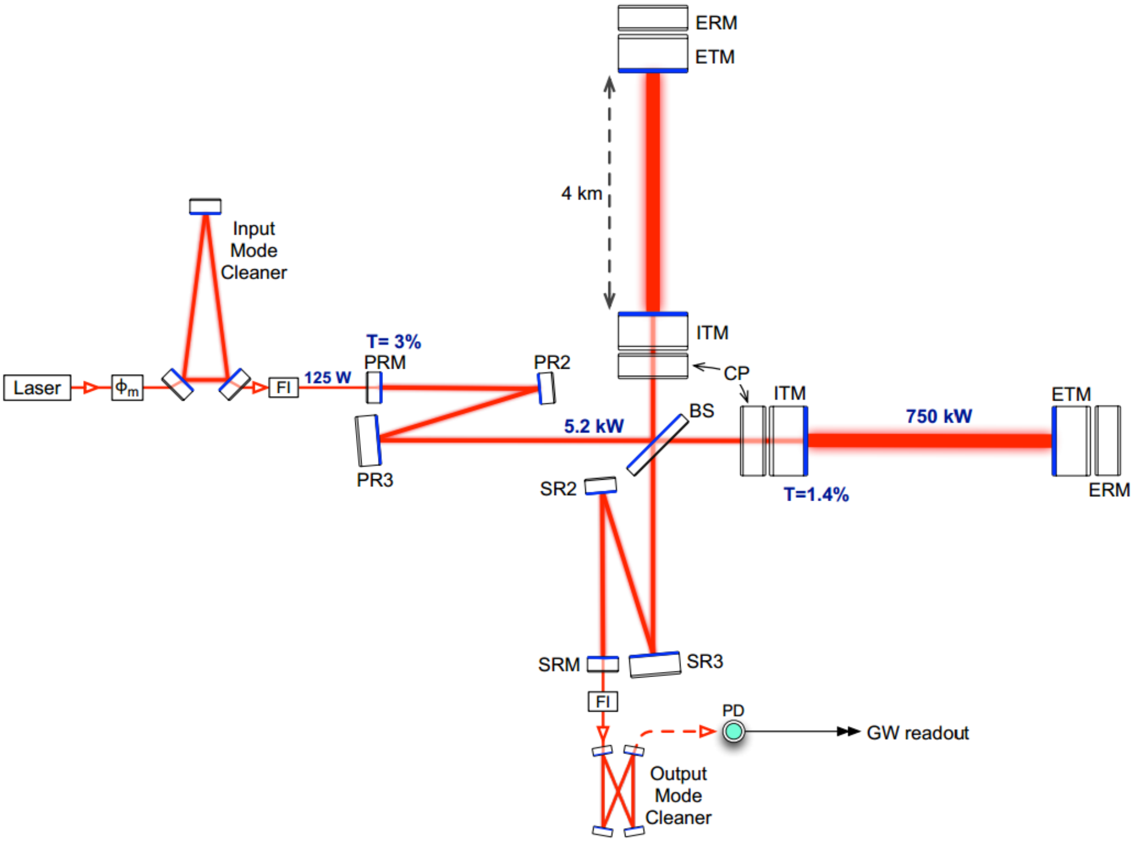

Figure 1.3.1: Simplified diagram of the optical layout of the advanced LIGO interferometers. Shown are the end test masses (ETM), input test masses (ITM), beam splitter (BS), power recycling mirror (PRM) and signal recycling mirror (SRM) as well as the mode cleaners at input and output. Image from [5].

In order to decouple the two components the sideband frequency is compared with the laser light or, as in advanced LIGO, an output mode cleaner is used. In what is known as a DC readout scheme, this is provided through the interferometer by moving the readout to be slightly away from complete destructive interference. This gravitational-wave signal is then converted into a strain $h(t)$ using time-domain filters [48, 92].

Noise sources in interferometric gravitational-wave detectors

The detector will be affected by noise from internal and external sources, which will limit the sensitivity of the detector. Here we discuss four significant sources of noise in interferometric gravitational-wave detectors; seismic noise, gravity gradient noise, thermal noise and quantum noise [81].

Figure 1.3.2 shows the frequency dependence of the noise for ‘enhanced’ LIGO and Virgo, from which the real data discussed in this thesis is derived and figure 1.3.3 shows

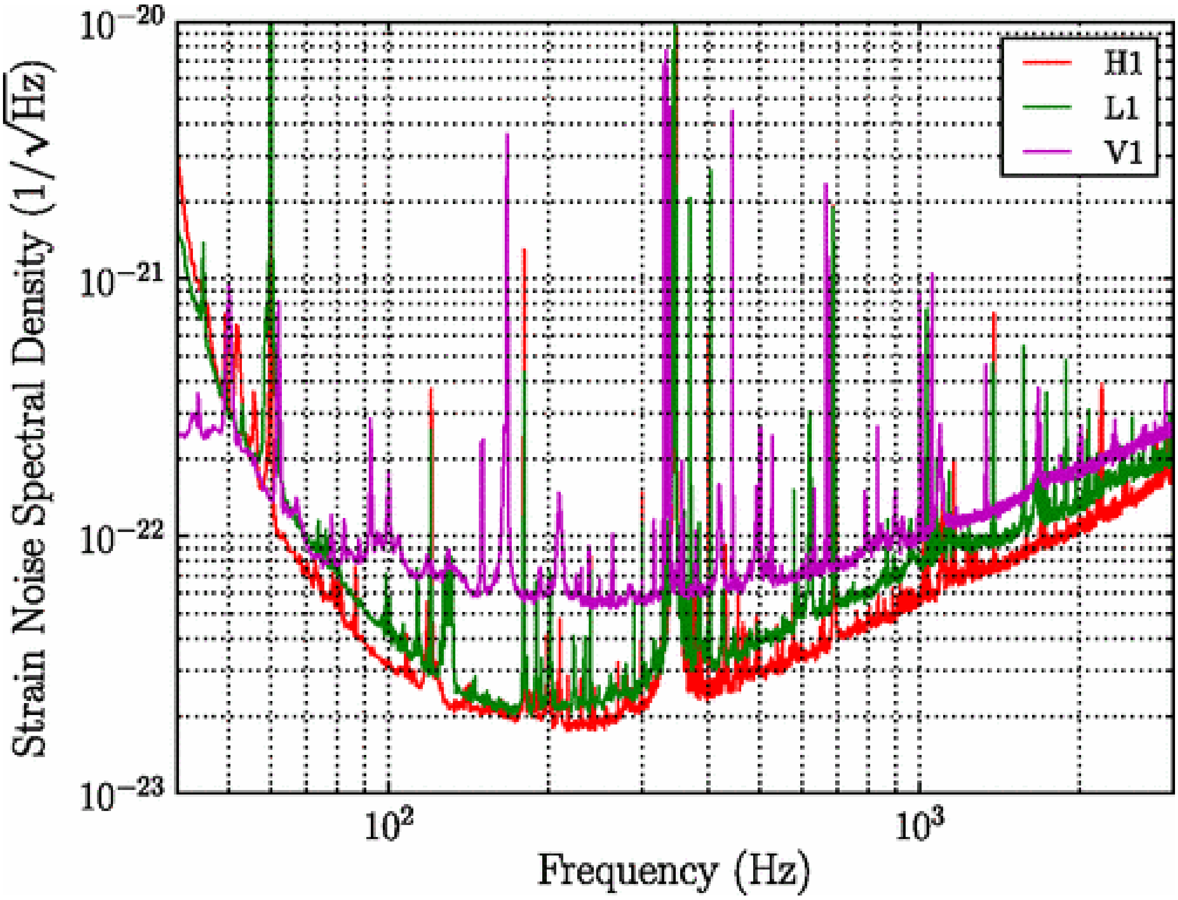

Figure 1.3.2: Plot of typical detector strain noise in enhanced LIGO (H1,L1) and Virgo (V1) noise during LIGO S6 and Virgo VSR2/3. We can see clearly the 60 Hz line from electrical equipment AC current in the LIGO detectors H1 and L1 and the equivalent 50 Hz line in the Virgo detector. We also see the suspension violin modes at around 350 Hz in all detectors [9].

the projected noise budget for advanced LIGO detector and its constituent parts.

Seismic noise comes from the vibration of the Earth around the test masses, and can be reduced by using suspensions. Even using a single-stage pendulum will make the transfer function for horizontal motion fall off as $(1/f)^2$ above the pendulum resonance, and in a similar way the vertical motion can be reduced using a spring. For example, Virgo mirrors are suspended using a seven-stage pendulum arrangement, with six cantilever springs, allowing operation to below 10 Hz [27]. Seismic noise limits sensitivity at low frequencies, below around 40 Hz in LIGO (see figure 1.3.2) and 10 Hz in advanced LIGO (figure 1.3.3).

Gravity gradient noise is caused by nearby mass fluctuations coupling to the mirrors, such as a surface wave. These can be mitigated by physical environment monitoring and subtraction (used in advanced LIGO/Virgo), by going underground (KAGRA), or by the slightly more extreme measure of putting the detector in space (NGO/DECIGO). Gravity gradient noise contributes significantly to the low frequency noise below around

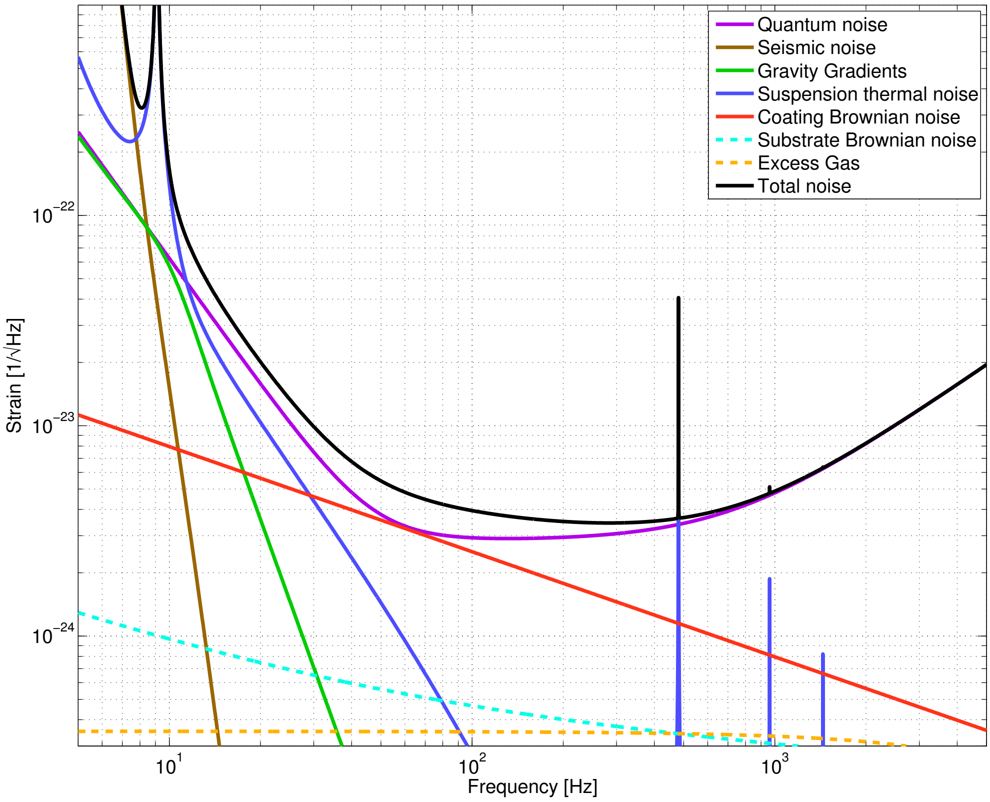

Figure 1.3.3: Plot of the advanced LIGO noise budget. Showing also the constituent parts of the noise discussed in the main text. We also see the noise caused by excess gas in the vacuum system, which is negligible, and noise from the Brownian noise of the substrate, the material onto which the mirror is applied, which is also negligible. Image made using the Gravitational Wave Interferometer Noise Calculator, GWINC [21].

15 Hz in LIGO (figure 1.3.3).

Thermal noise comes from the thermally excited vibration of particles in the mirror coating, the suspension and (a little bit from) the substrate to which the coating is attached. The way to reduce this is to carefully choose the materials for use in things such as mirror coating, mirror substrate and suspension fibres as well as considering bonding methods, the size and mass of the mirrors, suspension thickness and mirror coating methods [98], as well as the possibility of using mirrors at cryogenic temperatures (KAGRA) [71]. Coating Brownian noise is one of the main sources of noise at what would otherwise be advanced LIGO's most sensitive point (~ 70 Hz, see figure 1.3.3). Suspension thermal noise contributes highly at low frequency, as well as causing spikes of noise at the 'violin' modes, where the suspension fibres resonate in the same way as a violin string [88], which we see in figures 1.3.2 and 1.3.3.

Quantum noise comes from photon shot noise and radiation pressure noise. Reduction of one of these noise sources through changes in laser power will increase the other, and the combination of the two noise sources at which the noise is lowest is limited by the standard quantum limit – the Heisenberg uncertainty principle in its position-momentum form. Quantum noise is a broadband limiting noise source in advanced LIGO [81].

One particular noise source of interest to continuous-wave searches in LIGO is the 60 Hz electrical line from AC current in U.S. power lines and electronic equipment. This line affects searches for the Crab pulsar and for J0900$-$3144, as it has occasional broadened wings which overlap the Crab's expected gravitational-wave frequency (60.45 Hz) and a harmonic of this line affects J0900$-$3144's GW frequency (180.023 Hz), these lines can be seen in figure 1.3.2. As the AC current frequency in Europe is 50 Hz, this issue does not affect the search for signals from the Crab in Virgo data.

Many other methods also exist for reducing these and other noise sources, but these are not included here for brevity.

1.3.2 Antenna response

As the detector is not one which can be 'pointed' in the same way that an electromagnetic telescope might be, we need to consider the different ways in which a signal would be affected, given a different source location.

All gravitational-wave signals, regardless of the source will be affected by the detector's response to the different sky position, known as the antenna pattern, or antenna response of the detector. The strain signal detected is

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t), \tag{1.3.1}$$

where h_+ and $h_\times$ are the gravitational-wave signal plus and cross polarisations, as seen in section 1.1.2. F_+ and $F_\times$ are the antenna pattern functions, one for each polarisation, relating to the relative orientation of the interferometer to the direction from which the signal is coming – i.e. the *sky position* of the source.

The antenna patterns arise as a result of the way in which the test masses will respond to a gravitational wave. Considering the ring of test particles in figure 1.1.1, we imagine the mirrors to be at the top (Y-arm) and the right hand side (X-arm) of the initial circle respectively – these are marked as the red and yellow circles on that figure. We see that the cross polarised wave in this configuration does affect the mirrors, but does so equally in both directions. This would lead to no readout in the detector, as the interferometer measures the difference in arm length of the two mirrors. So our sensitivity in a particular direction is also governed by the polarisation of the wave.

The antenna patterns are [55]

$$F_+(t;\psi;\alpha;\delta) = \sin\zeta\left[a(t;\alpha;\delta)\cos 2\psi + b(t;\alpha;\delta)\sin 2\psi\right], \tag{1.3.2}$$

$$F_\times(t;\psi;\alpha;\delta) = \sin\zeta\left[b(t;\alpha;\delta)\cos 2\psi - a(t;\alpha;\delta)\sin 2\psi\right], \tag{1.3.3}$$

where ζ is the angle between the IFO arms, which is $\pi/2$ for most current detectors[9], α and δ are the right ascension and declination of the source, and ψ is the polarisation angle of the gravitational wave. $a(t;\alpha;\delta)$ and $b(t;\alpha;\delta)$ are functions of time which come from the varying detector position and orientation with respect to the sky position of the source, the exact form of these functions is not required here[10] but they are periodic, repeating once per sidereal day. The antenna response to plus and cross

---

[9] For practical reasons, GEO600 actually has $\zeta = 0.524\pi$ [23], and future detectors, such as ET, may decide to use a triangular setup of co-located, non-aligned detectors, for which $\zeta = \pi/3$ [19]

[10] The exact form of $a(t;\alpha;\delta)$ and $b(t;\alpha;\delta)$ can be found in [55]

polarised gravitational waves are shown in figure 1.3.4.

We see that the points directly above and below an interferometer are the most sensitive points, which is due to the transverse nature of the gravitational-wave perturbations on the mirrors. The insensitive points are for waves coming from a direction which bisects the angle between the detector arms.

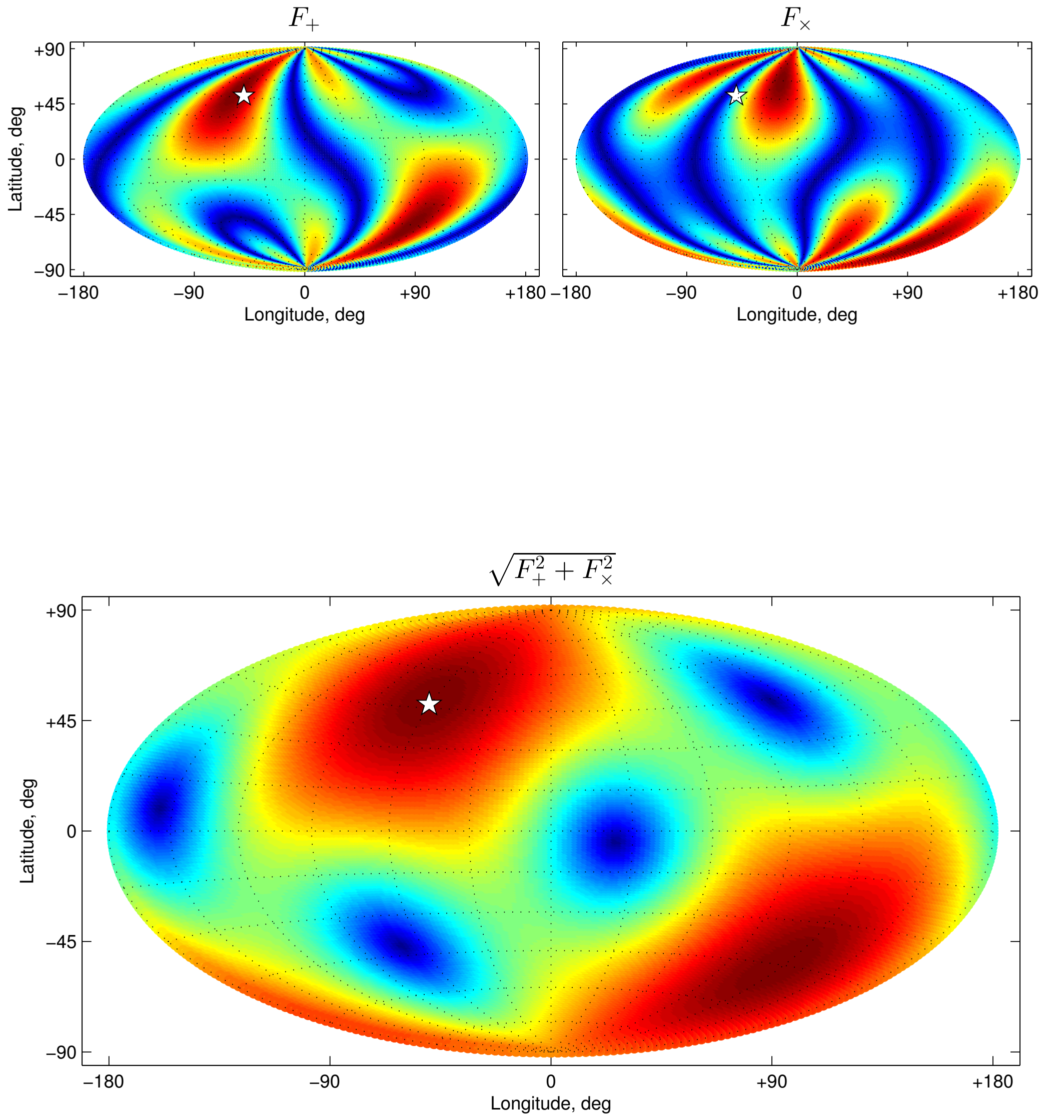

Figure 1.3.4: Sky map of the interferometer antenna response for gravitational-wave plus polarisation (top left), cross polarisation (top right) and combined (bottom), indicated by the star is the location of the detector, for which we used the H1 detector for this example The colour map indicates the sensitivity of the detector to a signal from this direction, from least (blue) to most (red) sensitive.

Chapter 2

Continuous gravitational waves

> Willow: *Carpe diem. You told me that once.*
>
> Buffy: *'Fish of the day'?*
>
> Willow: *Not carp. Carpe. It means 'seize the day.'*
>
> Buffy the Vampire Slayer,1998

Signals from rapidly rotating neutron stars are a promising source of gravitational waves, with the advantage that they are expected to be very long-lived, and as such can be searched for using coherent methods, increasing signal-to-noise ratio (SNR, ρ) proportionally to the square root of observation time. These quasi-sinusoidal signals are analysed using various methods by different groups within the LIGO and Virgo scientific collaborations, whether for unknown sources, [1, 3, 6, 15] or targeted sources [17, 25, 39, 54].

Section 2.1 details the signal model used for continuous-wave searches, as well as an explanation of much of the notation used in the following chapters. This model, or a slight variation thereof [79, 82], is used in all current continuous-wave searches as the expected signal.

Section 2.2 gives an introduction to the basics of Bayesian parameter estimation, a statistical framework used in order to search the parameter space so that we can find the posterior probability distribution on the values of our parameters of interest.

Dupuis, Pitkin and Woan [38, 39, 83] used Bayesian techniques to develop the 'Glasgow pipeline' or 'known pulsar time-domain Bayesian method' in the early-to-mid 2000s. This is an analysis routine to infer gravitational-wave signal parameters of continuous sources, in particular those with known sky position and phase evolution from radio data. This technique is described in section 2.3. This method consists of two parts, the first of which is the heterodyne stage, detailed in 2.3.1. The heterodyne algorithm provides a drastically downsampled time series of interferometer data for a given source which has had rotational phase dependency and relative motion effects removed, this is the B_K data (notation similar to that used in [39]). This is a highly accurate routine, but has the drawback that it can be computationally expensive. The second part of the analysis is a Bayesian parameter estimation, much of which will be covered in section 2.2, but the details of its use are given in section 2.3.2.

2.1 Continuous gravitational-wave signals from a non-precessing triaxial neutron star

In order to search for a long-duration signal in the data, we need to know our expected signal. This will be a combination of the gravitational-wave signal $h_{+/\times}$ and the antenna pattern $F_{+/\times}$ as given in equation 1.3.1; these four time-dependent functions make up our gravitational-wave strain signal, with $F_{+/\times}$ given by equations 1.3.2 and 1.3.3 and $h_{+/\times}$ dependent upon our source.

The source that we will consider is a non-precessing, triaxial neutron star. We make this choice as there is little to no evidence of precessing behaviour in observed neutron stars, and if the star was precessing it would be almost immediately damped by the rotation [59]. A triaxial star is the simplest, most consistent way to see a mass quadrupole in a rotating neutron star. The gravitational waves expected from this type of source will have h_+ and $h_\times$ given by [55]

$$h_+(T) = h_0\left(1+\cos^2\iota\right)\cos\Phi(T;\lambda), \tag{2.1.1}$$

$$h_\times(T) = h_0\cos\iota\sin\Phi(T;\lambda), \tag{2.1.2}$$

where $\Phi(T;\lambda)$ is the phase evolution of the gravitational wave due to source rotation, T is a time frame of reference stationary with respect to the source and ι is the angle between the source rotation axis and the line of sight from the detector to the source. λ is the set of parameters which will affect the phase evolution of the signal, and will consist of the intrinsic rotational motion of the neutron star – characterised by the frequency and its derivatives $f^{(0),(1),(2)\cdots}$, the sky position α and δ, and any binary parameters. λ is entirely known in a targeted search, and so is omitted in the rest of this chapter; however it is important in non-targeted searches, and so is returned to in chapter 6, which discusses follow up of all-sky and directed candidates using Bayesian parameter estimation. The other parameters, h_0, ψ, ι, are often known as the amplitude or intrinsic parameters, as they mainly affect the amplitude of the signal, though ψ affects both phase and amplitude over longer duration. With somewhat counter-intuitive nomenclature, the initial phase ϕ_0 is often included in this set of amplitude parameters, as this is also unknown in the search. ϕ_0 is the gravitational-wave phase at the reference epoch in the frame of the detector. As a point of notation, ϕ_0^{rot} represents the rotational phase at the reference time and throughout this work, $\phi_0 = 2\phi_0^{\mathrm{rot}}$.

The reason for making the clarification between gravitational and rotational phase is that in this work we consider only the deformation of the neutron star with spherical harmonic $l = m = 2$ [11, 17] which is a triaxial ellipsoid rotating about its axis of inertia, and this harmonic is associated with gravitational-wave emission at twice the rotation frequency, such that the gravitational-wave frequency, f, is twice the rotational frequency, ν, $f = 2\nu$. We make this assumption though chapters 3 and 4 are ambivalent towards the exact emission mechanism, so long as the frequency evolution is known. This means that it can be used in searches for the $l = 1$ harmonic, such that $f = \nu$ and $\phi_0 = \phi_0^{\mathrm{rot}}$ [45, 57, 79]. The expected gravitational-wave strain amplitude h_0 from the $l = m = 2$ harmonic is[1] [3]

$$h_0 = \frac{4\pi^2 G}{c^4}\frac{I_{zz}{f^{(0)}}^2}{d}\varepsilon, 1 \tag{2.1.3}$$

---

[1] Other papers such as [39] use a factor of 16 instead of 4 in the front of equation 2.1.3, this is down to the use of the gravitational-wave frequency $f^{(0)}$ or $\nu^{(0)}$, differing by a factor of two

where d is the distance to the source, $f^{(0)}$ is gravitational-wave frequency, and I_{zz} is the neutron star's moment of inertia about its principal axis of rotation. Ellipticity ε is defined by its moments of inertia as[2] [39]

$$\varepsilon = \frac{I_{xx} - I_{yy}}{I_{zz}}. \tag{2.1.4}$$

This means that by setting limits on h_0, we can set limits on the ellipticity of the source if we know its rotation frequency and its distance e.g. via $\varepsilon I_{zz} = \sqrt{8\pi/15}Q_{22}$ in the $l = m = 2$ harmonic mode [56, 76].

The time-dependence of h_+ and $h_\times$ is in the phase term $\Phi(T)$; this phase is approximated by a Taylor expansion, found by fitting frequency parameters to radio/X-ray/γ ray pulse times of arrival, (preferably from data spanning the same analysis period as the gravitational-wave search),

$$\Phi(T) \approx \Phi_0 + 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)}(T - T_0)^{l+1}}{(l+1)!}, \tag{2.1.5}$$

where $f^{(l)}$, is the l^{th} time derivative of the gravitational-wave frequency ($f^{(0)} = f$, $f^{(1)} = \dot{f}$, $f^{(2)} = \ddot{f} \ldots$), Λ is the limit to which we utilise the Taylor expansion and T_0 is a reference time, at which $\Phi(T_0) = \Phi_0$. The true frequency evolution is generally well-described by this Taylor expansion for our targets; however inaccuracies or noise in this frequency evolution can be an issue for particular sources [82], but these are not taken into account here, as they only affect a few very young pulsars (less than a few tens of thousand years old).

Table 2.1 shows order-of-magnitude ranges for the rotational frequency and its derivatives of known isolated pulsars, where $\nu^{(l)}$ is the l^{th} derivative of rotational frequency, analogously to $f^{(0)}$ and figure 2.1.1 shows where these pulsars lie in the $f^{(0)}$,$|f^{(1)}|$ plane. There are two distinct populations; one which generally has higher frequencies, containing most of our targets, is a population of *recycled* millisecond pulsars, which have been spun up by accretion from a binary companion, the rest have not. Most LIGO targets are in the recycled millisecond pulsar population as we require

---

[2] Some papers [3] use an ellipticity with an absolute value of the numerator $\varepsilon = \frac{|I_{xx} - I_{yy}|}{I_{zz}}$, which prevents a negative h_0, in our case it is taken into account by ϕ_0, as a change in sign of h_0 is equivalent to a change in ϕ_0 of π.

Frequency derivative	Range	Number of pulsars
$\nu^{(0)}$	0.1 - 650 Hz	2150
$\lvert\nu^{(1)}\rvert$	$\sim 10^{-18} - 10^{-10}\,\mathrm{Hz\,s^{-1}}$	1853
$\lvert\nu^{(2)}\rvert$	$\sim 10^{-29} - 10^{-20}\,\mathrm{Hz\,s^{-2}}$	400
$\lvert\nu^{(3)}\rvert$	$\sim 10^{-33} - 10^{-27}\,\mathrm{Hz\,s^{-3}}$	8

Table 2.1: Magnitude ranges of rotational frequency and absolute values of its frequency derivatives for known isolated pulsars, and the number of pulsars for which this information is known (information from [26, 67]).

our targets to be above a certain frequency in order to be within the sensitive band of the detectors. For the initial LIGO and Virgo detectors this minimum frequency was $\sim 40\,\mathrm{Hz}$ and $\sim 20\,\mathrm{Hz}$ respectively [2]; the difference this makes to the number of sources available for targeted searches is highlighted in figure 2.1.1. For the advanced LIGO and Virgo detectors [5, 20], their wider frequency band will open up a larger population of sources to be considered. There are some pulsars within our frequency range which are not used as targets,as there is no up-to-date timing information.

In table 2.1, $\nu^{(1)}$, $\nu^{(2)}$ and $\nu^{(3)}$ are given as absolute values. $\nu^{(1)}$ will usually be negative; only 34 out of 1847 pulsars with known $\nu^{(1)}$ have positive values listed, however this is usually from acceleration effects due to the source being within a globular cluster, these cases are discussed more in chapter ??. $\nu^{(2)}$ and $\nu^{(3)}$ are more equally split between positive and negative values. $\nu^{(3)}$ is only known for a handful of pulsars, and the values are so small that they barely affect our searches.

Our detector is not inertial with respect to the source, due to relative motion effects including the orbital motion of the Earth, Earth rotation [31] and any motion of the source in a binary system. The inertial reference frame in our solar system is at the Solar System Barycentre (SSB), and the transformation between the time of arrival of the gravitational wave at the SSB, T, and detector arrival time, t, is

$$T = t + \delta(t), \tag{2.1.6}$$

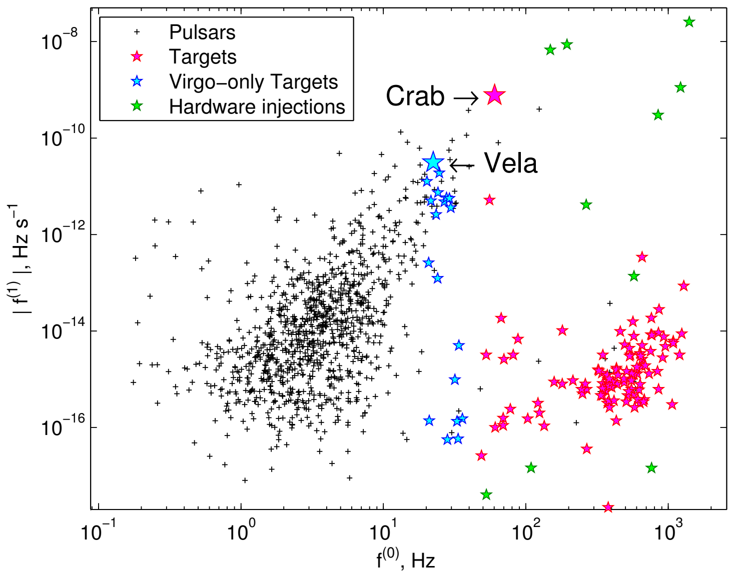

Figure 2.1.1: Locations of isolated pulsars in the $f^{(0)}$, $|f^{(1)}|$ plane, showing the two distinct populations of recycled millisecond pulsars and younger, more slowly rotating pulsars with higher spindowns. Targets which have had $f^{(1)}$ calculated from a characteristic spin-down age rather than their observed spin-down are not included [2, 26, 67]. Also shown are the locations of hardware injected signals, as discussed in section 4.1.6.

where $\delta(t)$ is made up of four terms[3]

$$\delta(t) = \Delta_{\mathrm{R}\odot} + \Delta_{\mathrm{S}\odot} + \Delta_{\mathrm{E}\odot} + \Delta_{\mathrm{Binary}}. \tag{2.1.7}$$

The Roemer delay $\Delta_{\mathrm{R}\odot}$ is the dominant term; it is the classical difference in time taken for the signal to arrive at the detector and at the SSB. $\Delta_{\mathrm{S}\odot}$ is the Shapiro delay, caused by the bending of space-time near to massive bodies, which in the case of an Earth-based detector is dominated by the Sun's contribution. The Einstein delay $\Delta_{\mathrm{E}\odot}$ is the time dilation due to redshifts, including relative motion and gravitational redshift caused by the Sun and other planets. Δ_{Binary} is an all-encompassing term considering the Roemer, Shapiro and Einstein delays caused by the source's own binary orbital system; for isolated sources, $\Delta_{\mathrm{Binary}} = 0$. Considering these time effects, the phase in the time frame of the detector is

$$\Phi(t) = \phi_0 + 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)}(t + \delta(t) - t_0)^{l+1}}{(l+1)!}, \tag{2.1.8}$$

---

[3] This time conversion will look familiar to pulsar timing astronomers, but with the notable exception of Δ_{DM}, the delay caused by a dispersion medium. As has been noted in section 1.1.3, the gravitational wave interacts very weakly with matter, this means that the wave is almost unaffected by the interstellar medium, which would slow down an electromagnetic counterpart.

where t_0 is the time at which the frequency parameters are defined, and $\Phi(t_0 - \delta t) = \phi_0$.

2.2 Bayesian parameter estimation

A wise man proportions his belief to the evidence[4].

David Hume, Scottish philosopher, 1711-1776

Bayes theorem is a cornerstone of statistical analysis, and we will use Bayes theorem to calculate the distributions of our gravitational-wave source parameters given the data. Using a dataset $\{B_K\}$, the provision of which we shall discuss later, we find the posterior (i.e. *after* the experiment) probability distributions on some parameters. In this section, we will use the example of the continuous-wave amplitude parameters $\mathbf{a} = [h_0, \psi, \phi_0, \iota]$. Bayes theorem states that

$$\mathrm{p}(\,\mathbf{a} \mid \{B_K\}\,, I\,) = \frac{\mathrm{p}(\,\mathbf{a} \mid I\,)\,\mathrm{p}(\{B_K\} \mid \mathbf{a}, I\,)}{\mathrm{p}(\{B_K\} \mid I\,)}, \tag{2.2.1}$$

where $\mathrm{p}(A|B)$ is probability of A given B, $\mathbf{a}$ is our set of parameters, I is background information which involves our assumed model (though this is dropped from further equations for ease of reading) and $\{B_K\}$ is our dataset. Each probability term in the above expression is as follows:

- $\mathrm{p}(\,\mathbf{a} \mid \{B_K\})$ (also denoted $P(\mathbf{a})$), the posterior: the degree of belief (or probability) that the data $\{B_K\}$ indicates these values of $\mathbf{a}$ – the distribution of this or its maximum value are generally the outcome of an analysis.
- $\mathrm{p}(\mathbf{a})$ (also denoted $\pi(\mathbf{a})$), the prior: our prior knowledge of the parameters, encoded here represents previous knowledge of the values of $\mathbf{a}$. For example, as an angle ϕ_0 must be within $[0, 2\pi]$.
- $\mathrm{p}(\{B_K\} \mid \mathbf{a}\,)$ (also denoted $\mathcal{L}(\mathbf{a})$), the likelihood: this is the probability of obtaining the data given the parameters specified.
- $\mathrm{p}(\{B_K\})$ (also denoted Z), the evidence: the probability of getting these values of B_K for any $\mathbf{a}$. The evidence is often unknown, and as such Bayes theorem

---

[4] A Bayesian statistician, however, proportions it to $\frac{1}{\text{the evidence}}$.

is often stated in the form $\mathrm{p}(\,\mathbf{a}\,|\,\{B_K\})\propto \mathrm{p}(\mathbf{a})\,\mathrm{p}(\{B_K\}\,|\,\mathbf{a}\,)$. The evidence is very useful in Bayesian model selection, as it is the total probability that the data comes from a signal with the defined information I, e.g. a particular model which has been assumed in the analysis.

As the detectors have so far been unable to detect continuous-wave signals, our figure of merit has been to set a 95% upper limit on h_0, $h_0^{95\%}$, meaning that 95% of our probability is contained within the range zero to $h_0^{95\%}$, defined as

$$0.95 = \int_{h_0=0}^{h_0^{95\%}} \mathrm{p}(h_0\,|\,\{B_K\})\,\mathrm{d}h_0. \tag{2.2.2}$$

This was used, for example, in [17], where the gravitational-wave energy emission from the Crab pulsar was limited using $h_0^{95\%}$ to be less than $\sim$ 2% of its spin-down limit. The spin-down limit is the total amount of rotational energy being lost, and hence an upper limit on the gravitational-wave energy being emitted. This energy is calculated by working out the energy output required to obtain the observed decrease in rotation frequency, or 'spin down'. In [11], the gravitational-wave energy from the Vela pulsar was limited to less than $\sim$ 41% of its spin-down limit. The spin-down limit is discussed more in chapter 5.

In searches in O1 – the first observing science run of advanced LIGO – there is a prospect of detection of continuous waves, meaning that the 95% upper limit will therefore not tell the entire story, and thus an SNR or a Bayes factor will be stated. The Bayes factor B_{ij} for competing models M_i and M_j is defined by

$$B_{ij} := \frac{Z(\{B_K\}\,|M_i)}{Z(\{B_K\}\,|M_j)}, \tag{2.2.3}$$

which intrinsically assumes equal prior weighting for the two models. This will be used in this thesis to compare the evidence that a signal is present against a signal not being present, i.e.

$$B = \frac{p(\{B_K\}\mid 0 < h_0 \le h_{\max})}{p(\{B_K\}\mid h_0 = 0)} = \frac{Z}{Z_{\text{noise}}}, \tag{2.2.4}$$

where $h_{\max}$ is the maximum allowed value of h_0 according to the prior.

The choice of priors for all of our parameters is important, and the least informative priors can be the best choice, particularly when we have no prior knowledge of the system. This choice means that our prior on the initial phase, $\mathrm{p}(\phi_0)$ is uniform over $[0, 2\pi]$, leading to a prior on ϕ_0^{rot} of $[0, \pi]$. The polarisation angle, $\mathrm{p}(\psi)$ is uniform over $[-\pi/4, \pi/4]$; we note that $\mathrm{p}(\psi)$ does not run over a total cycle of 2π, as a change in polarisation angle ψ of $\pi/2$ is equivalent to a change in ϕ_0 of π, and we want to avoid degeneracies. The prior on ι is uniform in $\cos\iota$ in the range $[-1, 1]$, giving a uniform probability in solid angle for the spin axis orientation.

Ideally the prior on gravitational-wave amplitude, h_0, would be either astrophysically motivated or a scale-invariant Jeffreys prior, however, when considering what we want our experiment to tell us, we decide against these as follows [79]:

An astrophysically motivated prior could set limits on the value of h_0 far below the sensitivity of all detectors so far, in this case the analysis would provide little or no evidence for gravitational waves in the data, so utilising a non-astrophysical prior may give a more informative result, as we can gain a scientifically interesting result which is independent of other methods. There is a possibility of using astrophysically set priors as the detection becomes more credible in the advanced detector era, but this is not discussed here as we only discuss initial/enhanced detector results. Given the fact that we want an interesting outcome to our analysis, we set our sights on a 95% upper limit on h_0 which is entirely independent of electromagnetic observation. A scale-invariant Jeffreys prior would be the standard uninformative prior to use. As h_0 is a multiplicative factor the Jeffreys prior would be $\mathrm{p}(h_0) \propto 1/h_0$ for $h_0 > 0$, $\mathrm{p}(h_0) = 0$ otherwise. In order to get conservative estimates for $h_0^{95\%}$, for which the data genuinely tells us that there is a 5% chance that the gravitational-wave amplitude is larger than this, we use a prior which is uniform in h_0 up to a sensibly chosen maximum value. We also do not want to exclude the not implausible value of $h_0 = 0$, which would imply the non-existence of gravitational waves from that pulsar.

When we are focussed only on a single parameter, the other parameters are often considered as *nuisance parameters*, to be removed from the posterior probability. The process to remove nuisance parameters is called marginalisation, and involves integrating over the likelihood with respect to these parameters; this means that the parameter is allowed to take any value according the prior range. Doing this means that we do not mistakenly ignore the value of the parameter, but use our prior knowledge of the system. To make our h_0 results independent of other parameters, we must marginalise over the nuisance parameters; in a targeted search, this would be

$$\mathrm{p}(\ h_0 \mid \{B_K\}) \propto \int\int\int \mathrm{p}(\mathbf{a})\,\mathrm{p}(\{B_K\}\,|\mathbf{a})\,\mathrm{d}\psi\,\mathrm{d}\phi_0\,\mathrm{d}\cos\iota. \tag{2.2.5}$$

In practice this integral is usually calculated numerically.

One way of combining data from separate analyses is to use the posterior results of one analysis as the prior probability for a new one, this is numerically the same as using a joint likelihood for the two analyses. One can also do this in an incoherent way, using only the posteriors on h_0 and $\cos\iota$ from another analysis as the prior, with uniform priors on ϕ_0 and ψ – as was done in [2].

2.2.1 Posterior distribution sampling algorithms

A posterior distribution sampling algorithm such as Markov Chain Monte Carlo [30] (MCMC) or nested sampling [45, 80, 97] is used in order to obtain the posterior distribution given the likelihood and our chosen priors. Each will be briefly described here:

Markov Chain Monte Carlo

The Markov Chain Monte Carlo method, or in particular the Metropolis-Hastings algorithm of posterior sampling, has been used in the past for Bayesian parameter estimation of continuous-wave signals [30]. The Metropolis-Hastings algorithm sets a number of points randomly in the prior space, and from each of these points comes a chain of points, for which a particular point in the chain, $\mathbf{a}'$, is based only upon the previous point in the chain, $\mathbf{a}_n$, with a probability of acceptance given by

$$\alpha(\mathbf{a}'|\mathbf{a}_n) = \min\left\{\frac{\mathrm{p}(\mathbf{a}')\,\mathrm{p}(\{B_K\}\,|\mathbf{a}')\mathrm{q}(\mathbf{a}_n|\mathbf{a}')}{\mathrm{p}(\mathbf{a}_n)\,\mathrm{p}(\{B_K\}\,|\mathbf{a}_n)\mathrm{q}(\mathbf{a}'|\mathbf{a}_n)}, 1\right\}, \tag{2.2.6}$$

where $\mathrm{q}(\mathbf{a}_n|\mathbf{a}')$ is the candidate generating pdf, for which a normal distribution gives good efficiency. This means that if the posterior probability of the new point is larger than that of the previous point it is always accepted, but if the new point has a lower posterior probability it is accepted with a probability given by the ratio of the two posterior points. From these chains, we should then climb any 'hills' in the posterior, as well as having a good probability of finding multi-modal peaks due to the probability of accepting new values with lower posteriors.

Nested sampling

The nested sampling algorithm is the posterior sampling method which will be used in this work. Nested sampling was first developed by Skilling in 2004 [94][5], intended primarily for use in evidence calculation, however, it was found to have the useful by-product of posterior sampling. The method involves setting N 'live' points in the parameter space according to the priors, and working out the likelihood $\mathcal{L}$ of each point. Each point is then ranked in order of likelihood, and the least likely point is assigned likelihood λ_ϑ, and removed, where ϑ denotes the iteration. This point will enclose prior mass, $\xi(\lambda)$, defined as

$$\xi_\vartheta := \underset{\mathcal{L}(\mathbf{a})>\lambda_\vartheta}{\int \cdots \int} \pi(\mathbf{a})\,\mathrm{d}\mathbf{a}, \tag{2.2.7}$$

however $\mathcal{L}(\mathbf{a})$ is not known during the execution of our algorithm as we have a finite number of points. Provided the placing of new points is done in a sensible manner, we can assume that the remaining prior mass enclosed is distributed equally to each live point,

$$\xi_\vartheta \approx \frac{\xi_0 - \sum\limits_{b=1}^{\vartheta-1}(1-\xi_b)}{N}. \tag{2.2.8}$$

We also assume that given a high enough number of live points, almost all of the prior mass is enclosed, and so $\xi_0 \approx 1$. This removed point is replaced by a randomly placed point, given the proviso that it has a higher likelihood value than the previous point. As the likelihood samples get more tightly distributed, it may no longer be viable to use the priors as the limits on the placing of this point. The method of choosing what limits to set on the random placing of this next point is defined by the proposal

---

[5] A detailed explanation is also available in [93].

distribution, and the setup of this proposal distribution is one of the most difficult and subtle points of nested sampling, and a few different methods can be used. Codes such as MultiNest use a clustering algorithm [44] and it is possible to use the covariance matrix of the current datapoints to create a proposal distribution [97]. In our case, we use an ensemble sampler [47].

As each iteration is performed, the likelihood of each new point will therefore increase and the associated prior mass for removed points will decrease.

The evidence in the nested sampling algorithm is calculated by using the prior mass and likelihood from the discarded points as an addition to the total evidence. As point ϑ with $\mathcal{L}_\vartheta = \lambda_\vartheta$ is removed, the evidence contained within the removed prior mass is added to the overall evidence total. This contained evidence ΔZ_ϑ is then $\Delta Z_\vartheta = \lambda_\vartheta \xi_\vartheta$, and the total evidence is the sum of all these points,

$$Z = \sum_\vartheta \Delta Z_\vartheta. \tag{2.2.9}$$

The stopping point of the algorithm is decided based upon the amount of evidence that is believed to remain in the parameter space, the *tolerance*. This remaining evidence is approximated by the multiplying the maximum likelihood by the remaining prior mass, $\mathcal{L}_{\max}\xi_\vartheta N$, where N is included as we remember that ξ_ϑ is the prior mass for the given live point.

The posterior sampling part of the nested sampling method comes from using all of the nested samples, discarded or not, and noting their posterior value. This is ΔZ_ϑ for the discarded points and for the remaining points, the leftover prior mass is divided equally between each point and multiplied by its likelihood. Each point is then accepted as a posterior sample with a probability given by its posterior value.

2.3 Known pulsar time-domain Bayesian method

The known pulsar time-domain Bayesian method is a well studied, and well used algorithm [2, 11, 17, 83] and consists of heterodyning the data and then performing parameter estimation with nominally unknown variance at each datapoint. A brief explanation follows, with a reminder that [38, 39] contain a lot more detail.

2.3.1 The heterodyne algorithm

To perform a complex heterodyne, we take the real detector output, $s(t)$, from file *frames*, which contain detector data and header information such as channel information and timestamps. $s(t)$ is a combination of the gravitational-wave signal $h(t)$ and the frequency-dependent noise $n(t) = \mathcal{N}\left[\sigma_T(f)\right]$, such that $s(t) = h(t) + n(t)$. In the heterodyne, each datapoint is multiplied by a time-varying complex exponential according to its phase, calculated from equation 2.1.8. $h(t)$ is of the form

$$h(t) = A(t)\mathrm{e}^{i\Phi(t)} + A^*(t)\mathrm{e}^{-i\Phi(t)}, \tag{2.3.1}$$

where $A(t)$ is a combination of the antenna pattern and source amplitude parameters [39]

$$A(t) = \frac{1}{4}F_+(t;\psi)h_0(1+\cos^2\iota) - \frac{i}{2}F_\times(t;\psi)h_0\cos\iota, \tag{2.3.2}$$

and $A^*(t)$ is its complex conjugate. We then multiply the time series by a time-varying complex phase correction term in order to remove the time dependence of the signal. To do this we use the known Doppler parameters to calculate the way in which the phase evolves, and multiply by its complex exponent. Because we do not know ϕ_0, we are unable to use $\Phi(t)$, so we define here $\phi(t) = \Phi(t) - \phi_0$. Once heterodyned (multiplied by the time-varying phase correction) the signal becomes

$$\begin{aligned} s_{\mathrm{het}}(t) &= s(t)\exp\left[-i\phi(t)\right] \\ &= A(t)\mathrm{e}^{i\phi_0} + A^*(t)\mathrm{e}^{-i\phi_0-2i\phi(t)} + n_{\mathrm{het}}(t). \end{aligned} \tag{2.3.3}$$

This multiplication removes the rotational phase evolution from the first term and gives us a rapidly varying term at twice the gravitational-wave frequency, the upper sideband. The noise $n_{\mathrm{het}}(t) = n(t)\mathrm{e}^{-i\phi(t)}$ has also been modulated by the phase evolution, however by the central limit theorem, this will remain approximately Gaussian. The data is then passed through a band-pass filter, $s_{\mathrm{het}} \to s'_{\mathrm{het}}$, in order to remove the upper sideband and most of the frequency-dependent noise (see section 1.3.1). After this procedure, the signal should be completely independent of the rotational phase evolution, and the noise contained will remain Gaussian

$$s'_{\rm het}(t) = A(t){\rm e}^{i\phi_0} + n'_{\rm het}(t). \tag{2.3.4}$$

The mean value of these heterodyned datapoints is taken over a time window Δt, $t_K - \Delta t/2 < t < t_K + \Delta t/2$ in order to downsample the data to a more manageable, but still informative amount. t_K is the centre of our time window, and Δt is the separation of these points. This mean value is the B_K dataset used in the parameter estimation [39],

$$B_K = \frac{1}{M}\sum_{i=1}^{M} s'_{\rm het}(t_i), \tag{2.3.5}$$

where K is the index of the datapoint, and M is the number of datapoints within the time scope of each B_K value.

We need to consider the limitations on the sample rate of the B_K values, i.e. how large can M (or directly related, Δt) be? Once the B_K values have been obtained, the time-varying component of $h_{+/\times}$ has been removed, and the time variance is only that of the antenna pattern functions from equations 1.3.2 and 1.3.3. $F_{+/\times}$ cycles over timescales of a day and the new sample rate cannot be reduced to be comparable to the time variation of these functions. This limits the separation of B_K values to be less than a couple of hours.

Computationally there is an advantage to performing two separate heterodyne procedures; the first, *coarse*, heterodyne does not take the time delay $\delta(t)$ into account, and downsamples to a much lower rate. In practice this is taken as downsampling from 16384 samples per second to one per second. The second, *fine*, heterodyne uses this downsampled coarse data, applies detector motion corrections and downsamples again to one per minute. This advantage comes from not having to calculate $\delta(t)$ so many times over the course of the coarse calculation. As with our argument with regards to the antenna pattern, $\delta(t)$ varies over time scales of a day[6] and is therefore treated as constant during this coarse procedure.

---

[6] See figure 4.1.15.

B_K then fits into our wider picture of the gravitational-wave signal as

$$B_K = \frac{1}{4}F_+(t_K;\psi)h_0(1+\cos^2\iota)\mathrm{e}^{i\phi_0} - \frac{i}{2}F_\times(t_K;\psi)h_0\cos\iota\mathrm{e}^{i\phi_0} + n_K, \tag{2.3.6}$$

or

$$B_K = y_K(\mathbf{a}) + n_K, \tag{2.3.7}$$

where n_K is the noise on B_K, $\mathbf{a}$ is again the set of amplitude parameters $[h_0, \phi_0, \iota, \psi]$ and y_K is the signal we expect at time t_K,

$$y_K(\mathbf{a}) := A(t_K)\mathrm{e}^{i\phi_0} = \frac{1}{4}F_+(t_K;\psi)h_0(1+\cos^2\iota)\mathrm{e}^{i\phi_0} - \frac{i}{2}F_\times(t_K;\psi)h_0\cos\iota\mathrm{e}^{i\phi_0}, \tag{2.3.8}$$

t_K denotes the time at which B_K is calculated, the centre of the time window. The noise n_K remains a zero mean Gaussian, the variance of which, σ_K^2, is related to that of $n(t)$, σ_T^2, by the central limit theorem,

$$\sigma_K^2 = \frac{\sigma_T(f^{(0)})^2}{r\Delta t}, \tag{2.3.9}$$

where r is the sample rate and Δt is the length of the time window, the B_K separation. For current searches, $r = 16384\,\mathrm{Hz}$ and $\Delta t = 60\,\mathrm{s}$.

Our unknown parameters $\mathbf{a}$ are still represented within the B_K values and this reduced dataset can be used as the basis of a Bayesian analysis as follows in section 2.3.2.

The heterodyne technique is an accurate method for calculating B_K from a time series, however it is computationally expensive and so we want to find a way in which to reduce the computational load; the way that we propose to do this follows in chapter 3.

2.3.2 Post-downsampling parameter estimation

Our B_K values are a signal $y_K(\mathbf{a})$ within Gaussian noise (equation 2.3.7). The likelihood that each value of B_k is drawn from a Gaussian distribution with known variance σ_K^2 and for a particular parameter set is

$$\mathrm{p}(\ B_K \mid \mathbf{a}, \sigma_K\ ) = \frac{1}{2\pi\sigma_K^2}\exp\left(-\frac{|B_K - y_K(\mathbf{a})|^2}{2\sigma_K^2}\right). \tag{2.3.10}$$

To combine the data for N datapoints we use $\mathrm{p}\left(\{B_K\} \mid \mathbf{a}, \{\sigma_K\}\ \right)$, the joint likelihood, which in our case is [39][7]

$$\begin{aligned}
\mathrm{p}\left(\{B_K\} \mid \mathbf{a}, \{\sigma_K\}\ \right) &= \prod_{k}^{N} \mathrm{p}\left(B_K \middle|\ \mathbf{a}, \sigma_K\ \right) \\
&= (2\pi)^{-N} \prod_{K}^{N} \frac{1}{\sigma_K^2} \exp\left(-\frac{|B_K - y_K(\mathbf{a})|^2}{2\sigma_K^2} \right). \qquad (2.3.11)
\end{aligned}$$

This likelihood is based on the data being present in stationary Gaussian noise. The noise from a detector is not Gaussian, but we can make this assumption over the bandwidth we are using. This noise however may not be stationary over the duration of our analysis, so we need to find a way to work around this. One way is to use an estimate of σ_K for each datapoint, in which case we can use equation 2.3.11 as the likelihood.

In many cases however we will not know σ_K for each datapoint, so, following [39], we can split B_K into ‘chunks’, each subscripted j, starting at $k_{1(j)}$ and running to $k_{2(j)}$ – a total of $m_j = k_{2(j)} - k_{1(j)} + 1$ points. The chunk is chosen such that the noise can be assumed to be stationary over its course, meaning that σ_K is a constant, $\sigma_{k(j)}$. For each chunk, the likelihood is [39]

$$\mathrm{p}\left(\{B_K\}_j \mid \mathbf{a}, \sigma_{k(j)}\ \right) = \frac{1}{\left(2\pi\sigma_{k(j)}^2\right)^{m_j}} \exp\left(-\frac{1}{2\sigma_{k(j)}^2} \sum_{k=k_{1(j)}}^{k_{2(j)}} |B_K - y_K(\mathbf{a})|^2 \right). \qquad (2.3.12)$$

We can marginalise over the constant unknown $\sigma_{k(j)}$, using a scale invariant Jeffreys prior on $\sigma_{k(j)}$ of $\mathrm{p}(\sigma_{k(j)}) \propto 1/\sigma_{k(j)}$, $\sigma_{k(j)} > 0$ to obtain

$$\mathrm{p}\left(\{B_K\}_j \mid \mathbf{a}\ \right) \propto \left(\sum_{k=k_{1(j)}}^{k_{2(j)}} |B_K - y_K(\mathbf{a})|^2 \right)^{-m_j}. \qquad (2.3.13)$$

This is the Student’s t-distribution with $2m_j - 1$ degrees of freedom, which tends towards the underlying Gaussian as $m_j \to \infty$. The joint likelihood given all chunks is

$$\mathrm{p}\left(\{B_K\} \mid \mathbf{a}\ \right) \propto \prod_{j} \left(\sum_{k=k_{1(j)}}^{k_{2(j)}} |B_K - y_K(\mathbf{a})|^2 \right)^{-m_j}. \qquad (2.3.14)$$

---

[7] This is equation 19 in [39], corrected as it did not account for non-constant σ_K^2.

Now that we have the likelihood, we can perform our parameter estimation using one of the methods in section 2.2.1.

Chapter 3

Spectral Interpolation of fast Fourier transforms

If we knew what it was we were doing, it would not be called research, would it?

Albert Einstein

The heterodyne method used to search for continuous gravitational-wave signals can be computationally expensive, but we can utilise frequency-domain data, which contains much of the information about the signal in only a handful of datapoints. This efficiency in terms of number of datapoints leads to massive increases in efficiency when calculating the B_K of section 2.2.

The method we use is the Spectral Interpolation algorithm (or *SplInter*), and it is intended as a 'black box' replacement for the heterodyne routine detailed in section 2.3.1. For rapid calculation of an equivalent to B_K, SplInter utilises fast Fourier transforms (FFTs), a specific algorithm for discrete Fourier transforms (DFTs). As a point of notation, we shall also refer to short Fourier transforms (SFTs, see [13], section IV C 1), which are relatively short FFTs compared to the duration of the experiment, and the form in which we access the data. We denote the result of the SplInter algorithm B_k, both to clarify the algorithm used, and as the time window used for each B_k value is longer than that used in the heterodyne algorithm.

The use of SFTs gives a remarkable increase in computational speeds in the calculation of B_k with little to no loss of accuracy in most situations. When we perform this calculation, we are able to ignore swathes of data in frequency space which is of little or no relevance to each particular source, needing less than 1 Hz of data for each source once the SFT has been created, rather than the equivalent of 16384 Hz of data used in the heterodyne algorithm. The SFTs have also already been created for other searches, e.g. [69] so use of them is free in terms of added computational expense for the collaboration, though this reuse of SFTs means that we must use 1800 s chunks of data, which causes issues discussed in section 3.4. A flowchart showing the SplInter algorithm is given in figure 3.4.3, but we will discuss each part of the algorithm in turn for clarity.

3.1 Continuous gravitational-wave signals in the frequency domain

Though the method described in this work is independent of the exact form of the signal, we require that the the phase evolution is known. For illustration, we continue to use the signal from a non-precessing, triaxial neutron star. We know the form of the signal in the time domain (equation 2.3.2) and so we also know it in the frequency domain through its Fourier transform. The discrete Fourier transform, by design, is a series of samples of the continuous Fourier transform[1] and the FFT is a different calculation of the DFT. This means the signal we expect to see in the FFT can be analytically calculated using the Fourier transform of $s(t')$

$$S(f) = \mathcal{F}[s(t')] = \int_{-\infty}^{\infty} [s(t')] \exp(-2i\pi f t')\, \mathrm{d}t' \tag{3.1.1}$$

$$= \int_{-\infty}^{\infty} \left[A(t')\mathrm{e}^{i\Phi(t')} + A^*(t')\mathrm{e}^{-i\Phi(t')} + n(t')\right] \exp(-2i\pi f t')\, \mathrm{d}t'. \tag{3.1.2}$$

---

[1] Appendix B shows for one of the later approximations that for sufficiently high sample rates this approximation is the same whether calculated from the discrete or continuous Fourier transform.

As addition and integration commute, $S(f) = H(f) + N(f)$, where $H(f) = \mathcal{F}[h(t')]$ and $N(f) = \mathcal{F}[n(t')]$.

The data has not been collected for an infinite amount of time, so we need to consider the window we are applying to the data in the time domain. If we take the data with equal weighting from within a certain time interval, this is known as a *rectangular* window,

$$W_R = \begin{cases} 1, & t_S \leq t' < t_E \\ 0, & \text{elsewhere} \end{cases}, \tag{3.1.3}$$

where t_S and t_E are the time at the start and end of the window respectively, and $\Delta t = t_E - t_S$ is the FFT length – the t_k separation, similar to t_K as used before. Other windows are in use for other searches within the collaboration, but as the standard FFTs used in continuous-wave searches are *almost* rectangular, this window is the only one considered here[2].

Given the rectangular window, the Fourier Transform of the windowed signal will be

$$H_k(f) = \int\limits_{t_k - t_0 - \Delta t/2}^{t_k - t_0 + \Delta t/2} \left[A(t') \mathrm{e}^{i\Phi(t')} + A^*(t') \mathrm{e}^{-i\Phi(t')} \right] \exp\left[-2i\pi f(t' - t_S)\right] \, \mathrm{d}t', \tag{3.1.4}$$

where $t_k - t_0$ is the difference in time between the middle of the FFT, the time at which B_k is calculated, and the reference epoch of the parameters, t_0. $H_k(f)$, $S_k(f)$ and $N_k(f)$ denote that we are using the SFT with the window around t_k and hence associated with B_k. The exponent in the Fourier transform term becomes $-2i\pi f(t' - t_S)$ as the FFT algorithm is performed in a time frame where the argument of the exponent is zero at the start of the transform.

We now consider the expected signal and its phase evolution $\phi(t')$, the time-varying part of $\Phi(t')$, which we remember from equation 2.1.8 to be

$$\Phi(t') = \phi_0 + 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)}(t' + \delta(t') - t_0)^{l+1}}{(l+1)!},$$

---

[2] Section 4.1.4 contains more about different windows and the tests to show whether the approximation still holds when Tukey windows are used in the FFT. Although rectangular windows are notorious for spectral power leakage, our method takes this into account in the calculations.

and move into a new time frame which is zero at the centre of the FFT, $t' \to t = t' - t_k + t_0$, so $\mathrm{d}t = \mathrm{d}t'$, and consider $\delta(t)$ as a linear expression

$$\delta(t) \approx \delta_k + \dot{\delta}_k t,$$

where $\delta_k = \delta(t_k - t_0)$, and $\dot{\delta}_k = \frac{\mathrm{d}}{\mathrm{d}t}\delta(t)$, $\delta(t)$'s first derivative at $t = 0$, the centre of the FFT.

$$\begin{aligned}\phi(t) &= 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)} \left[t + \delta(t) + t_k - t_0\right]^{l+1}}{(l+1)!} \\ &= 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)} \left[t(1 + \dot{\delta}_k) + t_k - t_0 + \delta_k\right]^{l+1}}{(l+1)!}.\end{aligned}$$

Expanding out the terms using the binomial expansion, this becomes

$$\phi(t) = 2\pi \sum_{l=0}^{\Lambda} f^{(l)} \sum_{\kappa=0}^{l+1} \left[t(1 + \dot{\delta}_k)\right]^{\kappa} (t_k - t_0 + \delta_k)^{l+1-\kappa} \frac{1}{\kappa!(l+1-\kappa)!}. \tag{3.1.5}$$

If we discard all terms higher than second order in t, $\phi(t)$ approximates as

$$\phi(t) \approx \phi_k + 2\pi f_k t + \pi \dot{f}_k t^2, \tag{3.1.6}$$

where we have defined the following:

$$\phi_k := 2\pi \sum_{l=0}^{\Lambda} \frac{f^{(l)}(t_k - t_0 + \delta_k)^{l+1}}{(l+1)!} \tag{3.1.7}$$

and

$$f_k := (1 + \dot{\delta}_k) \sum_{l=0}^{\Lambda} \frac{f^{(l)}(t_k - t_0 + \delta_k)^{l}}{l!}. \tag{3.1.8}$$

We approximate $\dot{f}_k$ by calculating the frequency at the beginning and end of the SFT window, and calculating the gradient between the two

$$\dot{f}_k := \frac{f_{\mathrm{end}} - f_{\mathrm{start}}}{\Delta t}, \tag{3.1.9}$$

where f_{start} and f_{end} are calculated in the same way as equation 3.1.8, but with t_k

replaced with $t_k - \frac{\Delta t}{2}$ and $t_k + \frac{\Delta t}{2}$ respectively (in the δ_k terms as well as in the explicit calculation). This numerical calculation is used as there would be an additional term in $\dot{f}_k$ introduced by t^2 terms in $\delta(t)$, $\ddot{\delta}_k$, the calculation of which would require a similar numerical approximation and any resulting inaccuracies would be amplified.

By using equation 3.1.6 as the time-dependent phase in equation 3.1.4, we can obtain an approximation for $H_k(f)$,

$$H_k(f) \approx \int\limits_{-\Delta t/2}^{\Delta t/2} \left[A(t)\mathrm{e}^{i\phi(t)+i\phi_0} + A^*(t)\mathrm{e}^{-i\phi(t)-i\phi_0}\right] \mathrm{e}^{-2i\pi f\left(t+\frac{\Delta t}{2}\right)}\,\mathrm{d}t. \tag{3.1.10}$$

We assume the signal amplitude and antenna pattern contributions to $A(t)$ are approximately constant on timescales of a small fraction of a day, so when Δt is small we can replace $A(t)$ with $A(t_k)$. Having defined $y_k := A(t_k)\mathrm{e}^{i\phi_0}$ in equation 2.3.8 we can therefore write

$$H_k(f) \approx \mathrm{e}^{-i\pi f\Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} \left[y_k\mathrm{e}^{i\phi(t)} + y_k^*\mathrm{e}^{-i\phi(t)}\right] \exp\left[-2i\pi f t\right]\,\mathrm{d}t. \tag{3.1.11}$$

This is the Fourier transform which will be calculated in the following approximations.

3.1.1 Different approximations for the integration

The signal is quasi-sinusoidal, with complex amplitude y_k varying slowly due to our source moving though the antenna pattern, changes in the delay and Doppler shifts, and intrinsic variations in the source spin rate. Using $H_k(f)$ from equation 3.1.11 and $\phi(t)$ from equation 3.1.6 we can say

$$\begin{aligned} H_k(f) \approx{}& y_k\mathrm{e}^{-i\pi f\Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} \exp\left[i\phi_k + 2i\pi(f_k - f)t + i\pi\dot{f}_k t^2\right]\,\mathrm{d}t \\ &+ y_k^*\mathrm{e}^{-i\pi f\Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} \exp\left[-i\phi_k - 2i\pi(f_k + f)t - i\pi\dot{f}_k t^2\right]\,\mathrm{d}t. \end{aligned} \tag{3.1.12}$$

These expressions are not strictly analytic due to the t^2 phase-dependency of the exponent, but are forms of the familiar Fresnel integral. The limiting form, when $\dot{f}_k$ is small, is the Fourier transform of a time-limited sinusoid, so we will consider this as a special case.

Sinc approximation, assuming $\dot{f}_k = 0$

The intrinsic $f^{(1)}$ of a source is generally very small, and over the course of an SFT, the change in frequency for most sources is small compared to the frequency resolution of the SFT. For example the Crab pulsar, which has an unusually large spin-down of $f^{(1)} = 7.4 \times 10^{-10}\,\text{Hz s}^{-1}$, will change in frequency by $\Delta f = 1.3 \times 10^{-6}\,\text{Hz}$ over the course of a half-hour SFT, or 0.12% of the width of a frequency bin, and a phase change of 1.3×10^{-3}. Instead, $\dot{f}_k$ will usually be dominated by the $\dot{\delta}_k$ terms from source and observer orbital motion. For an isolated source[3] this is usually negligible over $\sim 1\,\text{h}$, so the first approximation is that $\dot{f}_k = 0$, for the duration of the integral; this means that equation 3.1.12 is now

$$\begin{aligned} H_k(f) \approx y_k \mathrm{e}^{-i\pi f \Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} \exp\left[i\phi_k + 2i\pi(f_k - f)t\right]\, \mathrm{d}t \\ + y_k^* \mathrm{e}^{-i\pi f \Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} \exp\left[-i\phi_k - 2i\pi(f_k + f)t\right]\, \mathrm{d}t. \quad (3.1.13) \end{aligned}$$

Solving this integral gives

$$\begin{aligned} H_k(f) \approx y_k \Delta t \exp\left[i\phi_k - i\pi f \Delta t\right] \operatorname{sinc}\left[\pi(f_k - f)\Delta t\right] \\ + y_k^* \Delta t \exp\left[-i\phi_k - i\pi f \Delta t\right] \operatorname{sinc}\left[\pi(f_k + f)\Delta t\right], \quad (3.1.14) \end{aligned}$$

where we have used the *sinc* function[4]

$$\operatorname{sinc}(x) := \frac{\sin(x)}{(x)}. \quad (3.1.15)$$

---

[3] For binary sources, Δ_{Binary} often varies enough to give significant values to $\dot{f}_k$.

[4] This is the convention used in this work. The sinc function is sometimes defined elsewhere as $\operatorname{sinc}(x) := \frac{\sin(\pi x)}{\pi(x)}$.

We notice here that when we are around the signal frequency $f \approx f_k$, the first term will dominate the frequency-domain signal, as $\operatorname{sinc}\left[\pi(f_k - f)\Delta t\right] \gg \operatorname{sinc}\left[\pi(f_k + f)\Delta t\right]$. This means that we can safely remove this term, so equation 3.1.14 becomes

$$H_k(f) \approx y_k \Delta t \exp\left[i\phi_k - i\pi f \Delta t\right] \operatorname{sinc}\left[\pi(f_k - f)\Delta t\right]. \tag{3.1.16}$$

This is what we shall refer to as the *sinc approximation.*

When f_k is ‘on a bin’ (equal to the central frequency of the bin) the sinc function will be unity at f_k, but zero on all other bin frequencies. This can be seen by setting $f_k - f$ equal to a non-zero integer multiple of $1/\Delta t$, and by using the small angle approximation $\sin\left[\pi(f_k - f)\Delta t\right] \approx \pi(f_k - f)\Delta t$ for $\pi(f_k - f)\Delta t \approx 0$

$$H_k(f) \approx \begin{cases} y_k \Delta t \exp\left[i\phi_k - i\pi f \Delta t\right], & f = f_k \\ 0, & \text{on all other bins.} \end{cases} \tag{3.1.17}$$

We consider this as a special case and refer to it as the *on-bin approximation*, though it is mainly used for preventing ‘divide by zero’ errors in the algorithm. This is applied to signals which are close to a frequency bin, within 0.1% of the bin separation. Using a frequency evolution which is not exactly matched to that of a signal for the interpolation will reduce the SNR, as the data would not match the correct frequency evolution template; this SNR drop is particularly evident in a fully coherent search. We present tests of the signal deterioration given an incorrect frequency in section 4.2.2, though we stress at this point that the on-bin approximation does account for the error in phase caused by incorrect frequencies, so the actual effect is an insignificant reduction in the signal amplitude.

Fresnel approximation, assuming $\dot{f}_k \neq 0$

If $\dot{f}_k$ is not negligible, then we can approximate $H_k(f)$ with a numerical integration. Such circumstances would occur if the Doppler-shifted frequency is evolving significantly on timescales of Δt due to relative motion between the source and observer. If the rate of change of the signal frequency is a constant, i.e. $\ddot{f}_k = 0$, we would expect the signal to appear as a ‘Fresnel’ pattern in the spectrum, characterised by $\dot{f}_k$. Fresnel integrals have been studied fairly extensively, and algorithms exist for fairly rapid calculation [84]. The Fresnel integrals comprise a pair of functions defined as [101]

$$\mathcal{C}[w] := \int_0^w \cos\left(\frac{\pi x^2}{2}\right) \mathrm{d}x \tag{3.1.18}$$

and

$$\mathcal{S}[w] := \int_0^w \sin\left(\frac{\pi x^2}{2}\right) \mathrm{d}x. \tag{3.1.19}$$

In order to get $H_k(f)$ in terms of these integrals, we utilise a change of variable. Working from equation 3.1.12 we change the variable of integration to x, through the transformation

$$t \to x, \quad x^2 = 4(f_k - f)\frac{|\dot{f}_k|}{\dot{f}_k}t + 2|\dot{f}_k|t^2, \tag{3.1.20}$$

the factor of $|\dot{f}_k|/\dot{f}_k$ is included due to complications which would otherwise arise from the fact that we require x to be real, this is a factor of ± 1 depending on the sign of $\dot{f}_k$ which can be taken out of the integrals using $\cos(-\theta) = \cos(\theta)$ and $\sin(-\theta) = -\sin(\theta)$. Using Euler's formula and the change of variable above, equation 3.1.12 is now

$$H_k(f) \approx \frac{y_k}{\sqrt{2|\dot{f}_k|}} \mathrm{e}^{i\Delta\phi} \int_{w_{\text{start}}}^{w_{\text{end}}} \left[\cos\left(\frac{\pi x^2}{2}\right) + i\frac{\dot{f}_k}{|\dot{f}_k|}\sin\left(\frac{\pi x^2}{2}\right)\right] \mathrm{d}x, \tag{3.1.21}$$

where

$$\Delta\phi = \phi_k - \pi f \Delta t - \frac{\pi}{\dot{f}_k}\left(f_k - f\right)^2, \tag{3.1.22}$$

$$w_{\text{end}} = \frac{\sqrt{2|\dot{f}_k|}}{\dot{f}_k}\left(f_k - f\right) + \sqrt{\frac{|\dot{f}_k|}{2}}\Delta t, \tag{3.1.23}$$

and

$$w_{\text{start}} = \frac{\sqrt{2|\dot{f}_k|}}{\dot{f}_k}\left(f_k - f\right) - \sqrt{\frac{|\dot{f}_k|}{2}}\Delta t. \tag{3.1.24}$$

Here, we have ignored the second term in equation 3.1.12, as this term again has negligible effect in the region $f \simeq f_k$. Using the Fresnel integrals and the fact that

$\int_a^b f(x)\,\mathrm{d}x = \int_0^b f(x)\,\mathrm{d}x - \int_0^a f(x)\,\mathrm{d}x$, equation 3.1.21 becomes

$$H_k(f) \approx \frac{y_k}{\sqrt{2|\dot{f}_k|}} \exp\left(i\Delta\phi\right) \left\{ \mathcal{C}[w_{\text{end}}] + i\frac{\dot{f}_k}{|\dot{f}_k|}\mathcal{S}[w_{\text{end}}] - \mathcal{C}[w_{\text{start}}] - i\frac{\dot{f}_k}{|\dot{f}_k|}\mathcal{S}[w_{\text{start}}] \right\}. \tag{3.1.25}$$

We shall refer to equation 3.1.25 as the *Fresnel approximation* to the signal spectrum, and calculate the Fresnel integral terms with sufficient numerical precision[5] using the algorithm in [84], which for small $\dot{f}_k$ approximates the sinc approximation. Computationally, this procedure is more expensive than the sinc approximation, however it only needs to be used rarely, during periods of large $\dot{f}_k$. The value of $\dot{f}_k$ we choose to change from one approximation to the other is defined by $|\dot{f}_k|\Delta t^2 = 0.1$, which we discuss in section 4.1.3.

3.2 B_k and σ_k calculation

In order to find B_k, an unbiased estimate of y_k, we use a least squares fit, though we will now briefly go though why this is the best test to use from a Bayesian viewpoint. In order to calculate the best estimate of y_k, we want to marginalise over the noise, and maximise the posterior such that the B_k output is the most likely value of y_k for any noise level. We then find a maximum value of the posterior probability distribution of y_k.

We start by noticing that $H_k(f, \mathbf{a})$, (as calculated from one of equation 3.1.16, 3.1.17 or 3.1.25), can be expressed as a product of the unknown signal amplitude $y_k(\mathbf{a})$ and a known signal shape function, which we shall call a model, μ_k, defined as

$$\mu_k(f) := \frac{H_k(f, \mathbf{a})}{y_k(\mathbf{a})}. \tag{3.2.1}$$

In intrinsically noisy data, we cannot calculate y_k, and so the best estimate is B_k, which we remember from equation 2.3.7 as y_k plus noise n_k. The joint likelihood of obtaining the spectrum of the data, $S_k(f)$, given B_k and the set of model values $\{\mu_k\}$ in Gaussian noise of variance σ_F^2 is

---

[5] In this case, we use 100 iterations of a continued fraction.

$$\mathrm{p}\left(\, \{S_k\} \mid B_k, \sigma_F, \{\mu_k\}\, \right) = \frac{1}{(2\pi\sigma_F^2)^N} \exp\left(-\frac{1}{2\sigma_F^2} \sum_{j=1}^{N} |S_{kj} - B_k\mu_{kj}|^2 \right), \qquad (3.2.2)$$

where j indicates the frequency bin, and for visual clarity we have written $S_{kj} \equiv S_k(f_j)$ and $\mu_{kj} \equiv \mu_k(f_j)$. We can consider σ_F as a nuisance parameter, and marginalise over it. Choosing a Jeffreys prior of $\mathrm{p}(\sigma_F) \propto 1/\sigma_F$, $\sigma_F > 0$ and a uniform prior on B_k, which is constant anywhere in the complex plane, the posterior is

$$\mathrm{p}\left(\, B_k \mid \{S_k\}, \{\mu_k\}\, \right) \propto \int_0^\infty \frac{1}{\sigma_F^{(2N+1)}} \exp\left(-\frac{1}{2\sigma_F^2} \sum_{j=1}^{N} |S_{kj} - B_k\mu_{kj}|^2 \right) \mathrm{d}\sigma_F. \qquad (3.2.3)$$

To solve this integral, we use a change of variables to u, defined by

$$u^2 = \frac{\sum_j |S_{kj} - B_k\mu_{kj}|^2}{2\sigma_F^2}, \qquad (3.2.4)$$

$$\mathrm{d}u = \frac{2\sqrt{\sum_j |S_{kj} - B_k\mu_{kj}|^2}}{\sigma_F^2}\, \mathrm{d}\sigma_F. \qquad (3.2.5)$$

So we have a posterior of

$$\mathrm{p}\left(\, \{S_k\} \mid B_k, \{\mu_k\}\, \right) \propto \left[\sum_j |S_{kj} - B_k\mu_{kj}|^2 \right]^{-N} \int_0^\infty \mathrm{e}^{-u^2} u^{2N+1}\, \mathrm{d}u. \qquad (3.2.6)$$

The integral on the end of this expression is a standard integral, [101]

$$\int_0^\infty \mathrm{e}^{-u^2} u^{2N+1}\, \mathrm{d}u = N!2^{N+1}, \qquad (3.2.7)$$

so we can ignore it as a constant in the proportionality. The logarithm of equation 3.2.6 is then

$$\log\left[\mathrm{p}\left(\, \{S_k\} \mid B_k, \{\mu_k\}\, \right)\right] \propto -N \log\left[\sum_j |S_{kj} - B_k\mu_{kj}|^2 \right]. \qquad (3.2.8)$$

As we want to maximise the posterior and the logarithm function is a monotonically increasing function we can equivalently maximise the log posterior. From equation 3.2.8 we see the log posterior is maximised when $\sum_j |S_{kj} - B_k\mu_{kj}|^2$ is minimised.

In order to find the minimum, we differentiate $\sum_j |S_{kj} - B_k\mu_{kj}|^2$ with respect to B_k^* and set this derivative to be zero to find a stationary point. We differentiate with respect to B_k^* as it has the simultaneous properties that

1. $B_k \to B_k^*$ is conjugate conformal, leading to $\frac{\mathrm{d}B_k^*}{\mathrm{d}B_k} = 0$
2. B_k^* and B_k are mutually defined, so the most likely value of B_k^* defines the most likely value of B_k.

The following result can also be obtained by finding the minimum values for $\mathfrak{Re}\{B_k\}$ and $\mathfrak{Im}\{B_k\}$ and combining the results, these points are discussed more in appendix C.

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}B_k^*}\left(\sum_j |S_{kj} - B_k\mu_{kj}|^2\right) = \\ \frac{\mathrm{d}}{\mathrm{d}B_k^*}\sum_j \left[S_{kj}S_{kj}^* - B_k\mu_{kj}S_{kj}^* - S_{kj}B_k^*\mu_{kj}^* + B_kB_k^*\mu_{kj}^*\mu_{kj}\right] \\ = \sum_j \left[-S_{kj}\mu_{kj}^* + B_k\mu_{kj}^*\mu_{kj}\right] = 0.\end{aligned}$$

This minimum then gives us an equality which can be solved for B_k to obtain the most probable value,

$$B_k = \frac{\sum\limits_j [S_{kj}\mu_{kj}^*]}{\sum\limits_j [\mu_{kj}^*\mu_{kj}]}, \tag{3.2.9}$$

a result that is familiar from least square analysis. This is the maximum posterior estimator for B_k.

If we take the derivative again, this time with respect to B_k, we see that the second derivative is $\sum_j \mu_{kj}^*\mu_{kj}$, which is always positive, and hence this stationary point is a minimum[6].

As we are using SFTs with $\Delta t = 1800\,\mathrm{s}$, we are not able to use a Student's-t likelihood in the parameter estimation as there are not enough datapoints within a time period for which the noise can be assumed to be stationary. However using the SplInter algorithm, it is possible to estimate the noise on B_k, σ_k, directly from the SFT data.

---

[6] This point can be seen more clearly in appendix C, when dealing with the real and imaginary parts separately.

In order to estimate the variance of n_k, we would ideally follow a similar route, marginalising over B_k in equation 3.2.2 and maximising the posterior for σ_F

$$\mathrm{p}(\ \sigma_k \mid \{S_k\}, \{\mu_k\}\ )_{\max} = \left[\int_{-\infty}^{\infty} \mathrm{p}(B_k)\, \mathrm{p}(\ \sigma_k, B_k \mid \{S_k\}, \{\mu_k\}\ )\, \mathrm{d}B_k \right]_{\max} . \tag{3.2.10}$$

However, this integral is not analytic. We therefore choose to use the calculated value of B_k from equation 3.2.9 to obtain the best estimate of σ_k, this is equivalent to using the Dirac delta function as the prior on B_k in equation 3.2.10,

$$\mathrm{p}(B_k) = \delta_D \left(B_k - \frac{\sum_j [S_{kj} \mu^*_{kj}]}{\sum_j [\mu^*_{kj} \mu_{kj}]} \right) \tag{3.2.11}$$

The application of this is straightforward; we use most probable B_k calculated above to give us back the best estimate of H_{kj}, $H_{kj,\mathrm{best}} = B_k \mu_{jk}$. By taking $H_{kj,\mathrm{best}}$ from S_{kj}, we get the residuals, which are the best estimate of the noise, $N_{kj,\mathrm{best}}$. We use the residuals around the signal frequency and the calculate their variance to give us σ_F^2.

We can relate σ_F^2 to the time domain noise as $\sigma_F^2 = \sigma_T^2 \frac{2}{\Delta t^2 r}$ by Parseval's theorem and so use equation 2.3.9 to convert σ_F^2 to σ_k^2 as

$$\sigma_k^2 = \sigma_F^2 \frac{\Delta t}{2}. \tag{3.2.12}$$

We have now calculated B_k and an estimate of σ_k, meaning that we can use the known Gaussian likelihood, equation 2.3.11, in the posterior sampling algorithm, rather than the Student's-t that has generally been used previously.

3.3 Outlier removal

The noise in LIGO data contains many line features (see figure 1.3.2), which could adversely affect the B_k estimate if they are close to the signal frequency. In order to minimise these effects, we perform three outlier removal routines, with a threshold set by user input:

The first outlier removal, before any calculation has taken place, removes datapoints with real or imaginary values with absolute values above twice the input threshold number of standard deviations from the mean, which is assumed to be zero. This threshold is given by

$$S_{\mathrm{th},1} = 2N_\sigma \sqrt{\frac{\sum_j^N S_{kj}^2}{N-1}}, \tag{3.3.1}$$

where N is the number of datapoints within the bandwidth used for B_k and σ_k calculation and N_σ is the threshold set by the user. To have an effect on the data without removing too much, this will normally be around four or five.

The outlier removal is implemented by setting a variable within the algorithm of whether the datapoint is used, which we shall call U_j [7]

$$U_{j,1} = \begin{cases} 0, & \{(|\mathfrak{Re}[S_{kj}]| > S_{\mathrm{th},1}) \vee (|\mathfrak{Im}[S_{kj}]| > S_{\mathrm{th},1})\} \wedge \left(|f_j - f_k| > \frac{5}{\Delta t}\right) \\ 1, & \text{in all other cases} \end{cases} \tag{3.3.2}$$

This threshold is set to be larger than that used in the other routines, as this routine is designed to only remove very large spikes from the data. This number may seem small but we remember that the spike itself will be involved in the calculation of $S_{\mathrm{th},1}$, artificially inflating the noise estimate. The first ten bins around the signal frequency are not removed, in order to protect the data which will contain the most information about a prospective signal. The type of outlier to be removed by this procedure is illustrated in figure 3.3.1.

The second outlier removal takes place after B_k and σ_k have been calculated. Since we have estimated the residual noise within the FFT, $N_{kj,\,\mathrm{best}}$, in order to estimate the noise on the B_k values, we can use this for outlier removal. The threshold is set by the standard deviation of the residuals,

$$S_{\mathrm{th},2} = N_\sigma \sqrt{\frac{\sum_j^N U_{j,1}\left[N_{kj,\mathrm{best}}\right]^2}{\left(\sum_j^N U_{j,1}\right) - 1}}, \tag{3.3.3}$$

---

[7] We use the standard logical notation of $\wedge$ to mean logical conjunction, ‘and’, and $\vee$ to mean inclusive disjunction, ‘or’, which we stress is *not* exclusive. Exclusive disjunction ‘xor’ would be $\oplus$.

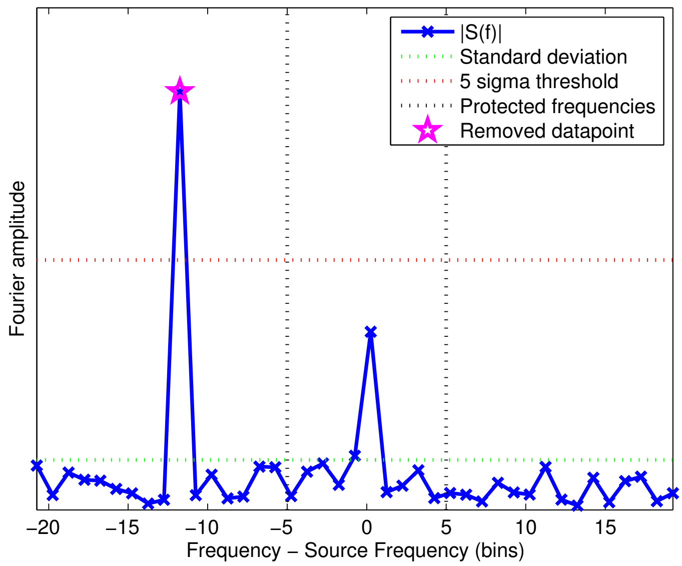

Figure 3.3.1: An illustration of the type of outlier removed by the first outlier removal routine. Shown are the standard deviation of the signal, the threshold for removal, and the protected band around the source frequency. The removed datapoint is indicated by the magenta star.

Using the calculated σ_F and $N_{kj,\,\text{best}} = S_{kj} - H_{kj,\,\text{best}}$ we remove datapoints with residual values above the threshold number of standard deviations.

$$U_{j,2} = \begin{cases} 0, & \{|\mathfrak{Re}[N_{kj,\,\text{best}}]| > S_{\text{th},2} \vee |\mathfrak{Im}[N_{kj,\,\text{best}}]| > S_{\text{th},2}\} \wedge \left(|f_j - f_k| > \frac{2}{\Delta t}\right) \\ 1, & \text{in all other cases} \end{cases} \tag{3.3.4}$$

This means that smaller spikes nearby that may have been missed by the first outlier removal can be removed, or we may remove a larger portion of a wider spike as the first outlier removal is only iterated once. This process removes these spikes without adversely affecting the signal itself; one way we ensure this is that the closest four datapoints are kept immune from this outlier removal. If any datapoints have been removed, then B_k and σ_k are recalculated, and this outlier removal routine repeated, until no more datapoints are being removed. The type of erroneous line that this routine is designed to remove is shown in figure 3.3.2. An example of the first and second outlier removal routines working in tandem with real data is shown later in figure 4.1.19 and explained in the text of section 4.1.5.

Because of these outlier removal routines we must alter the previous estimator for

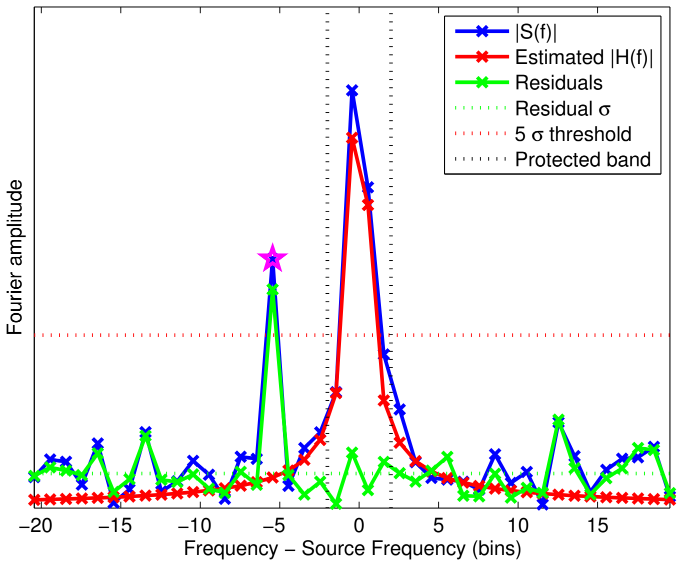

Figure 3.3.2: An illustration of the type of outlier removed by the second outlier removal routine. Shown are the best fit of the data, the standard deviation of the residuals, the threshold for removal calculated from the standard deviation of the residuals, and the protected band around the source frequency. The removed datapoint is indicated by the magenta star.

B_k to take these removed datapoints into account, so equation 3.2.9 becomes

$$B_k = \frac{\sum\limits_j [S_{kj}\mu^*_{kj}U_{j,1}U_{j,2}]}{\sum\limits_j [\mu^*_{kj}\mu_{kj}U_{j,1}U_{j,2}]}. \tag{3.3.5}$$

The third outlier removal occurs after B_k and σ_k values have been calculated for all SFTs. The mean value of all $\{\sigma_k\}$,$\langle\sigma\rangle$ is calculated, and any SFT which has $\mathfrak{Re}[B_k]$, $\mathfrak{Im}[B_k]$ or σ_k above the threshold number times this mean value, i.e.

$$U_k = \begin{cases} 1, & (\sigma_k < N_\sigma\langle\sigma\rangle) \wedge (\mathfrak{Re}[B_k] < N_\sigma\langle\sigma\rangle) \wedge (\mathfrak{Im}[B_k] < N_\sigma\langle\sigma\rangle) \\ 0, & \text{otherwise} \end{cases} \tag{3.3.6}$$

where U_k denotes whether this B_k value is used or not. This outlier removal is designed to remove outliers in the data which have an unusually high noise, but the noise is spread evenly over the frequency range considered by the SplInter algorithm, meaning that the first two outlier removal routines do not see any problem in the data. This outlier could be caused by unusually high low frequency noise, such as in figure 3.3.3.

The low frequency noise is unusually high in this case, and the power bleeds into the spectrum up to around 300 Hz. For sources with signals in this range, the noise from

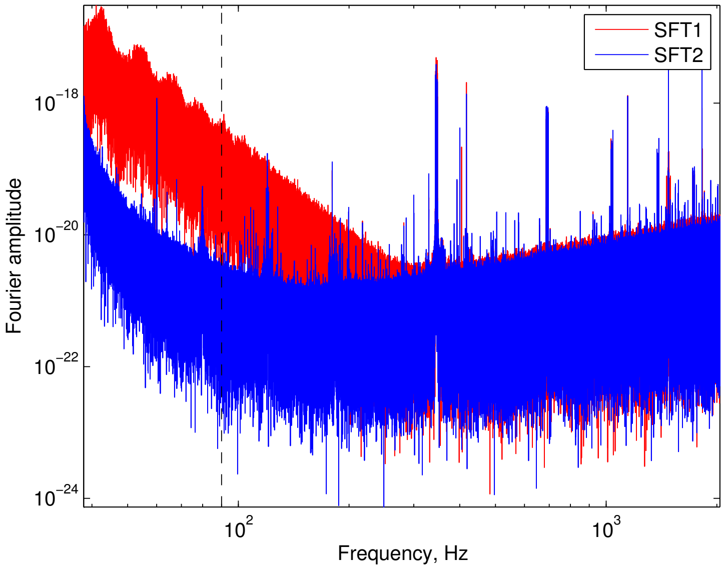

Figure 3.3.3: An illustration of the type of outlier removed by the third outlier removal routine, showing the power spectra of two example SFTs. SFT1 has an unusually high low frequency noise contribution, bleeding power into frequency channels up to around 300 Hz. Indicated by the black dashed line is a source at 90 Hz, for which the noise on the signal would differ by an order of magnitude between the cases.

this SFT would be extremely large compared to that in an SFT with a normal noise profile; for sources above this frequency, the noise is normal, so would be unaffected by this outlier removal.

This procedure removes all the B_k points which would be heavily suppressed in a parameter estimation routine. This means that they do not need to be included in the likelihood calculations, leading to a more computationally efficient parameter estimation stage. An equivalent of this outlier removal is also present in the heterodyne algorithm.

3.4 Algorithm implementation

Rome wasn't built in a day, but I wasn't on that particular job.

Brian Clough

This section will describe the way in which the SplInter algorithm works through the data in order to get the B_k and σ_k estimates.

The data from the detector comes with 'quality flags', and continuous searches tend to use all data in *science mode*, this is data for which it is deemed that the quality is good enough for scientific analysis, but the interferometer may not be at its most sensitive. We use this as all of the data contains the signal, and the noise estimate will decrease the importance of the datapoints (or the outlier removal routine will remove datapoints) in times of poor quality data. During each contiguous *segment* of science mode data, we use data in frequency-domain FFT format, SFTs. The main input files for the algorithm are therefore a segment list, a set of files defining the source parameters and an *SFT cache* of pointers to the Fourier data, made using `ligo_data_find`[8]. Another option at this point is a directory containing an SFT cache for each segment of the data. Using this option drastically reduces the amount of time taken to load the SFT catalogue in long duration searches with many segments (see section 4.2.1), which will be most real analyses of continuous signals.

Figure 3.4.1 shows a histogram of the duration of segments in the LIGO Hanford detector (H1) during science run 6, (S6), a length of time for which the detector was taking data, between Jul 08, 2009 and Oct 20, 2010. The length of these science segments is dependent upon the detector being stable and the interferometer cavities being resonant, this is largely down to transient seismic activity at the detector site. The average segment duration is 2.3 hours and ranges between one second and thirty two hours. We also see the fraction of the total time contained within segments of this length and the *dead time* as a fraction of the total time within the segments. Because we are using 1800 s SFTs there is a significant amount of time lost by requiring data to be in integer multiples of 1800 s; for example, all of the segment time is lost if it is less than 1800 s long as we are not able to produce an SFT in that time, or if the segment

---

[8] This and other routines (usually denoted `lalapps_`) are within the LALsuite software repository `https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html`

is 2700 s long, then 900 s of data will be lost. This dead time is 9.9% of the original amount of data in S6, and leads to a reduction in SNR of around 5%. To compare this to the heterodyne routine, which uses $\Delta t = 60$ s, the dead time here leads to losing 0.35% of the total available time, an SNR reduction of 0.17%. These are lower bounds to the actual values, as the SFTs may not be taken from the very start of the segment, so the more tightly fitting segments may not be available, also the heterodyne routine throws away the first 60 s of any segment due to the impulse response of the filters. The dead time from using Δt of 60 s is not plotted, as it is negligible.

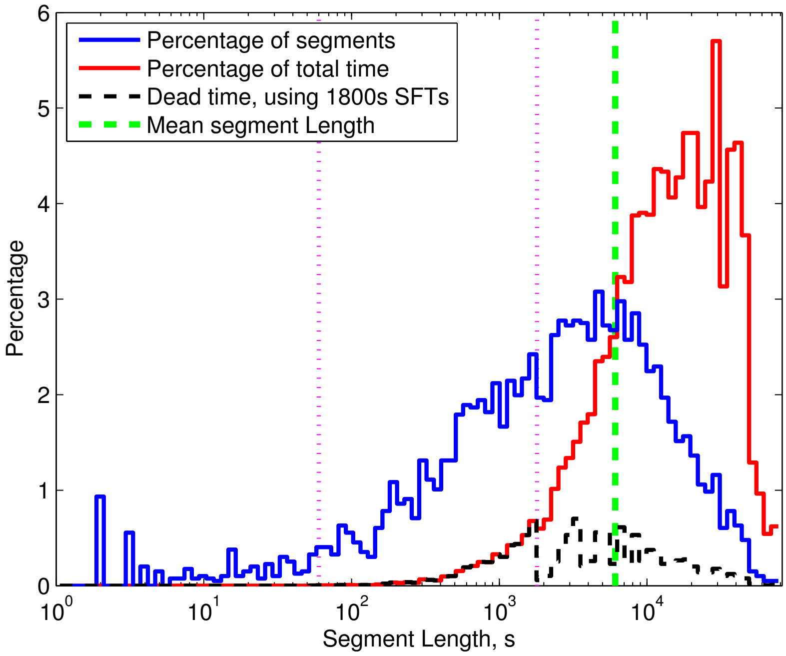

Figure 3.4.1: A histogram of the duration of segments in the H1 detector in S6 (blue). Shown also are the fraction of the total time contained within these segments (red), and the fraction of the time lost by using 1800 s SFTs (black). Indicated by the vertical dotted lines are the 60 s and 1800 s times used in the heterodyne and SplInter algorithms respectively.

Shown in figure 3.4.2 as a function of time, the *duty cycle* is the percentage of time that the detector is in science mode. Overall for S6 this duty cycle was around 51% in H1 and 48% in L1.

The SplInter algorithm loops through each segment, and in each segment processes each SFT in turn. Before going into each segment, we check that the segment falls within the time constraints for which the source parameter files are relevant. We also check that the source frequency is within the frequency range of the SFT, typically 38 to 2038 Hz for LIGO SFTs. If the source frequency is outside of this range, then we do not analyse the source with the SplInter algorithm, and will use the heterodyne

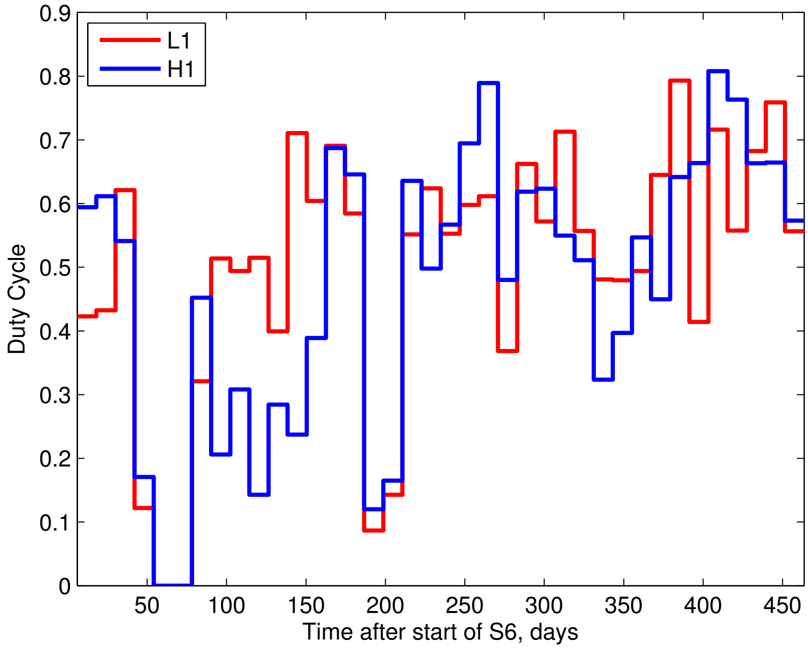

Figure 3.4.2: The duty cycle as a function of time for S6 in H1 (blue) and L1 (red). For this illustration S6 was split into forty equal sections in time, and the duty cycle is the fraction of time for which the detector is in science mode. There are periods where L1 and H1 had low or zero duty cycles, e.g. around days 50 to 80, these are during times of maintenance and commissioning.

algorithm in its search, though sources outside of this range are in regions of very high noise, which may rule out an analysis anyway. We also may use only the Virgo detector to analyse the source, which has a frequency range of 9 to 2038 Hz; we see these sources, which include Vela, in figure 2.1.1. Within each SFT, we also check that the frequency has not drifted out of this range due to orbital motion, if this test fails, then the source is ignored for this SFT, but may be considered again later.

The file output is to a file buffer rather than the file itself; this means that we do not overload the file system with the output, which could be accessing thousands of files a second if not using this buffer. This buffer is emptied to the file whenever it reaches 1 MB of data and when the algorithm finishes. Figure 3.4.3 shows a flowchart of how the SplInter algorithm processes the SFTs within each segment.

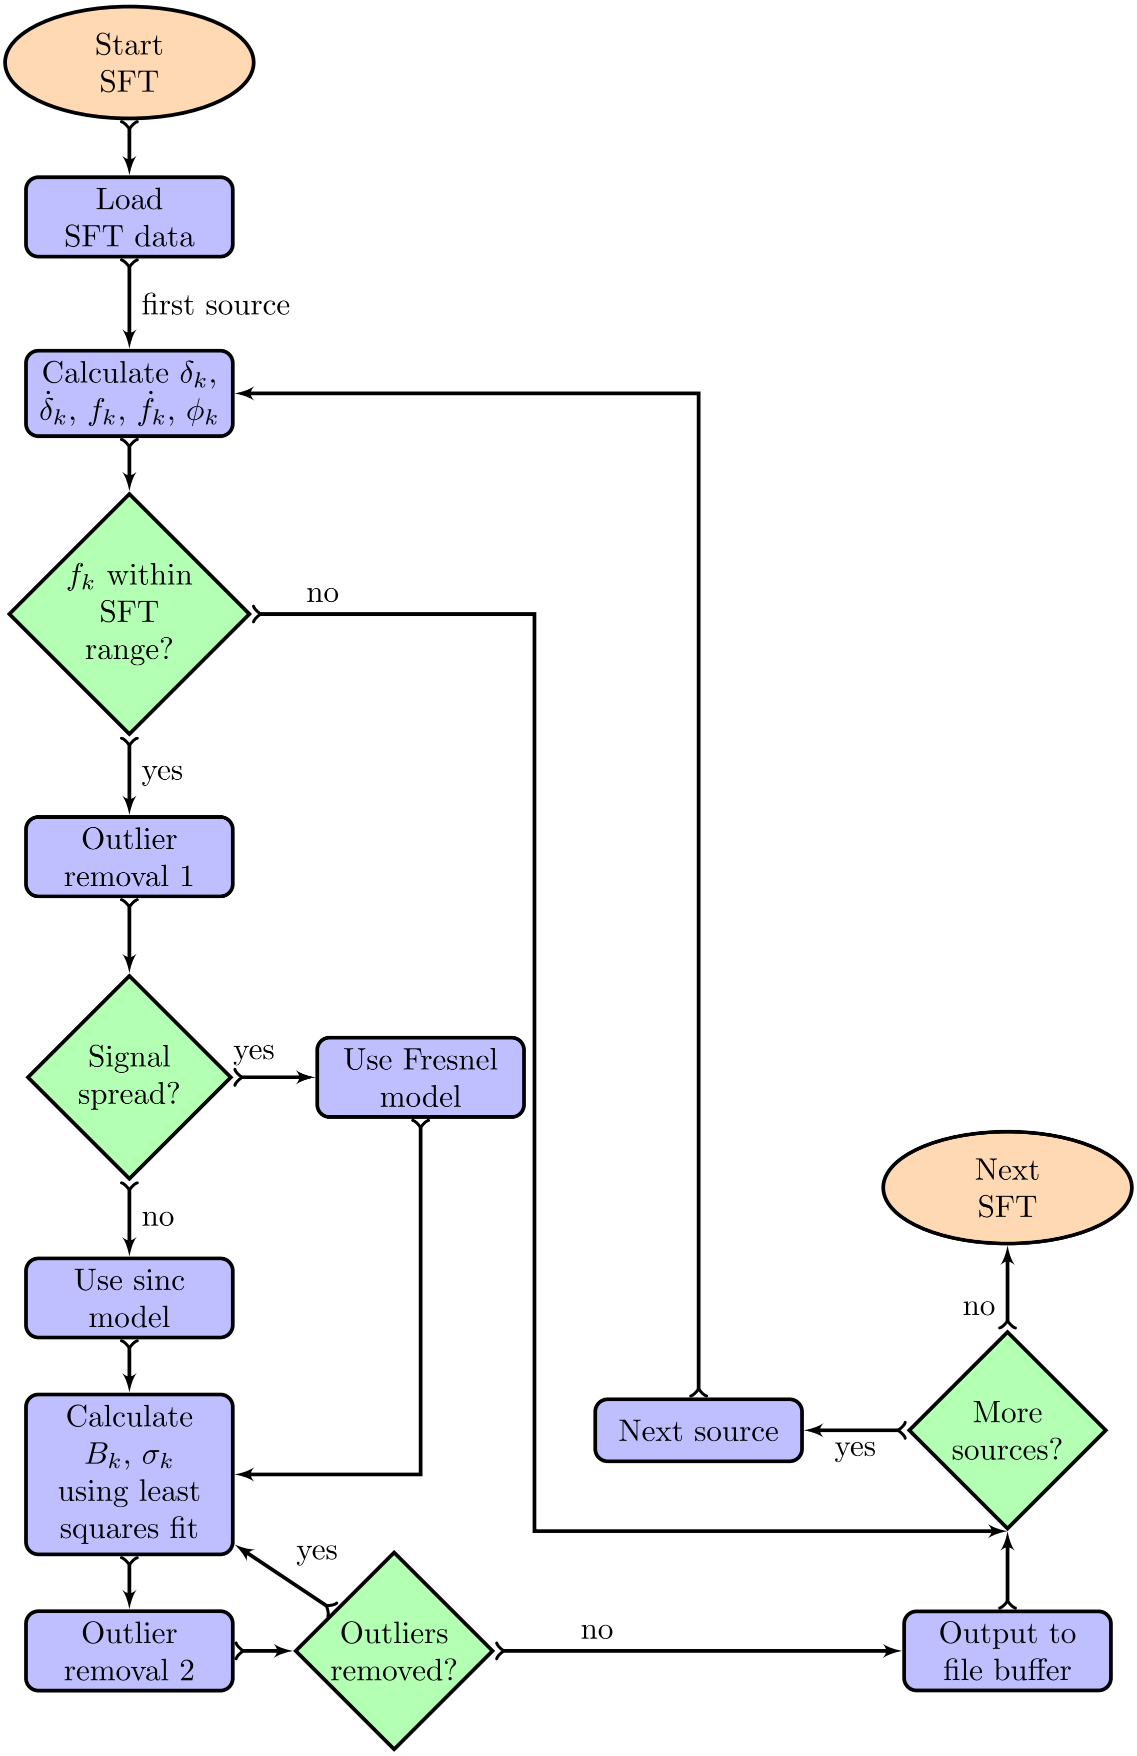

Figure 3.4.3: Flowchart showing the Spectral Interpolation algorithm during each SFT; each step is discussed more within the text.

Chapter 4

Testing of the Spectral Interpolation algorithm

If you try and take a cat apart to see how it works, the first thing you have on your hands is a non-working cat

Douglas Adams

In order to ascertain whether SplInter is a viable black box replacement for the heterodyne routine, we need to test the outputs of the two processes. There are three categories of tests presented here; accuracy, speed and frequency response.

The first accuracy test in section 4.1.1 ensures that we can be confident that the $B_{k/K}$ output from the two routines is recovering the signal equivalently, we perform the two routines in a situation where the output should be exactly the same, using a noiseless signal with known parameters. Section 4.1.2 shows a similar test, but for signals from sources in binary systems; these may not be recoverable by the SplInter algorithm, as the binary orbit may cause the frequency to change in a higher than linear order over the course of the SFT. We look at finding a general way to decide whether we can analyse a signal in a binary system.

The accuracy tests performed in sections 4.1.3 and 4.1.4 check whether the assumptions we have made regarding spin-down and windowing effects respectively are correct. Section 4.1.5 shows tests devised to check the σ_k outputagainst expected or calculated values. Then the most important test in section 4.1.6, we test the end-to-end routine by using short duration hardware injections in real noise.

In section 4.2.1 we then test the algorithm performance and the speed increase of SplInter compared to the heterodyne routine. We also profile the algorithm, giving details on where it is spending most of its time, and hence computational effort, which could (and did, during the work developing the algorithm) lead to improvements in the algorithm efficiency. Section 4.2.2 shows the frequency response of the SplInter algorithm, showing the errors introduced in amplitude when we use incorrect Doppler parameters.

4.1 Accuracy testing

> *There is nothing new to be discovered in physics now. All that remains is more and more precise measurement*
>
> William Thomson, 1st Baron Kelvin, 1824-1907

The SplInter algorithm is designed to replace the heterodyne algorithm, and should only do so if it is found to be sufficiently accurate in comparison. The accuracy tests are implemented to show that the assumptions we have made, and the general SplInter method, are sufficient to give us an adequate replacement.

Although there is no intrinsic reason why a frequency domain analysis in general should be in any way better or worse than its time domain equivalent, the assumptions we have made will have an effect on the overall accuracy of the algorithm. In particuar the assumptions to remove higher order terms than t^2 – in both the phase evolution and the barycentering terms – mean that the SplInter output will differ from the exact calculation of the heterodyne method.

The LIGO and Virgo scientific collaborations use hardware and software injections for testing algorithms, and in section 4.1.6 we present a search for hardware injection signals using both SplInter and heterodyne methods. Section 4.1.1 also uses the parameters of the hardware injections, as in tables 4.1 and 4.2. In hardware injections the test masses of the detector are pushed using electrostatic drives in a way which will mimic a gravitational-wave signal, this is advantageous over software injections, as it is possible to check for calibration errors and any problems in obtaining $s(t)$.

	$f^{(0)}$ (Hz)	$f^{(1)}$ (Hz s^{-1})	α (rad)	δ (rad)
PULSAR0	265.577	-4.15×10^{-12}	1.24882	-0.98118
PULSAR1	849.083	-3.00×10^{-10}	0.65265	-0.51404
PULSAR2	575.164	-1.37×10^{-13}	3.75692	0.06011
PULSAR3	108.857	-1.46×10^{-17}	3.11319	-0.58358
PULSAR4	1403.163	-2.54×10^{-08}	4.88671	-0.21758
PULSAR5	52.808	-4.03×10^{-18}	5.28183	-1.46327
PULSAR6	148.719	-6.73×10^{-09}	6.26139	-1.14184
PULSAR7	1220.980	-1.12×10^{-09}	3.89951	-0.35693
PULSAR8	194.308	-8.65×10^{-09}	6.13291	-0.58326
PULSAR9	763.847	-1.45×10^{-17}	3.47121	1.32103

Table 4.1: Doppler parameters of hardware injections used in testing continuous-waves algorithms (rounded to save space). $f^{(0)}$ is defined at Nov 01, 2003, 00:01:04 GMT, $f^{(2)}$ and higher order terms are zero in these signals.

Parameter	h_0	ϕ_0	ψ	$\cos\iota$
PULSAR0	2.46649×10^{-25}	2.66	0.77009	0.7949
PULSAR1	1.06005×10^{-25}	1.28	0.35603	0.4638
PULSAR2	4.01852×10^{-24}	4.03	-0.22179	-0.9286
PULSAR3	1.62771×10^{-23}	5.53	0.44428	-0.0807
PULSAR4	4.56205×10^{-23}	4.83	-0.64794	0.2773
PULSAR5	4.84996×10^{-24}	2.23	-0.36395	0.4630
PULSAR6	6.92191×10^{-25}	0.97	0.47098	-0.1537
PULSAR7	2.19820×10^{-24}	5.25	0.51232	0.7568
PULSAR8	1.58763×10^{-23}	5.89	0.17047	0.0739
PULSAR9	8.13001×10^{-25}	1.01	-0.00856	-0.6192

Table 4.2: Amplitude parameters of hardware injections used in testing continuous-waves algorithms. ϕ_0 is defined at Nov 01, 2003, 00:01:04 GMT.

4.1.1 Noiseless signal recovery - isolated pulsar signals

In the case of a noiseless signal, the heterodyne and SplInter routines should give exactly the same answer, given the same Δt, as $B_k = y_k$. Figures 4.1.1 to 4.1.8 show the result of using SplInter and heterodyne routines (both with $\Delta t = 1800\,\mathrm{s}$) on signals which have been created without any noise, given an isolated neutron star with parameters as in tables 4.1 and 4.2. The frames were made using `lalapps_create_pulsar_signal_frame` and SFTs were made from these frames using `lalapps_MakeSFTs`.

For the error comparison, we take the mean value of the heterodyned B_K values within the scope of the SFT used to create the SplIntered B_k data. Numerically this is the same as using a 1800 s Δt in the heterodyne algorithm. We denote this averaged heterodyne B_κ,

$$B_\kappa = \sum_K^N \frac{B_K(t_S \leq t_K < t_E)}{N}, \tag{4.1.1}$$

where N is the number of datapoints within $\Delta t/2$ of t_k. The fractional error comparison shows the difference in the absolute values of $B_{k/\varkappa}$ from the two algorithm outputs divided by the heterodyned $B_\varkappa$ output, and the angle error comparison shows the difference between the arguments of the two algorithms.

In the comparisons, we use the *mixed* interpolator, which is a dynamic algorithm, utilising the sinc approximation when appropriate and the Fresnel approximation when it is not. The point at which we switch between the two algorithms is discussed in sections 3.1.1, 3.4 and 4.1.3. We use the sinc interpolation in the case that $|\dot{f}_k|\Delta t^2 < 0.1$, and the Fresnel interpolation otherwise.

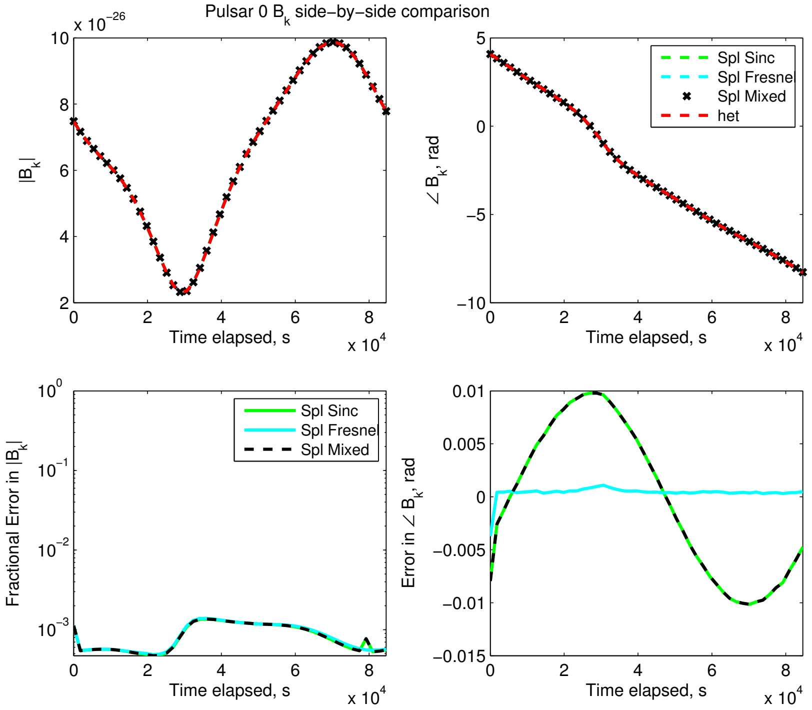

Figure 4.1.1: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR0. We see in this case that there are errors in B_k introduced by the sinc approximation, but that these are not significant enough to cause the switch to the Fresnel approximation.

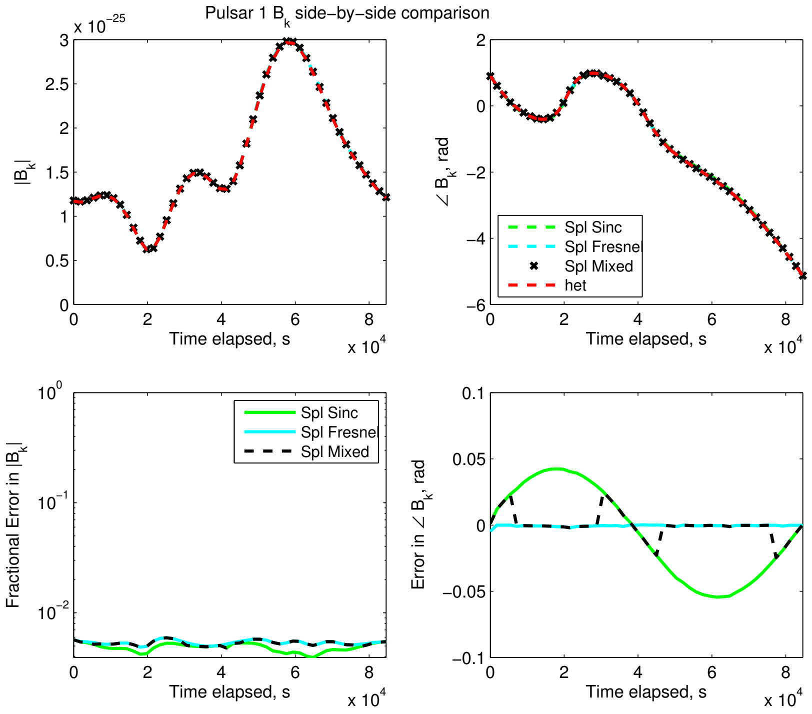

Figure 4.1.2: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR1. We see in this case that there are errors in B_k introduced by the sinc approximation, and that these are occasionally significant enough to cause the switch to the Fresnel approximation, which recovers the signal more accurately.

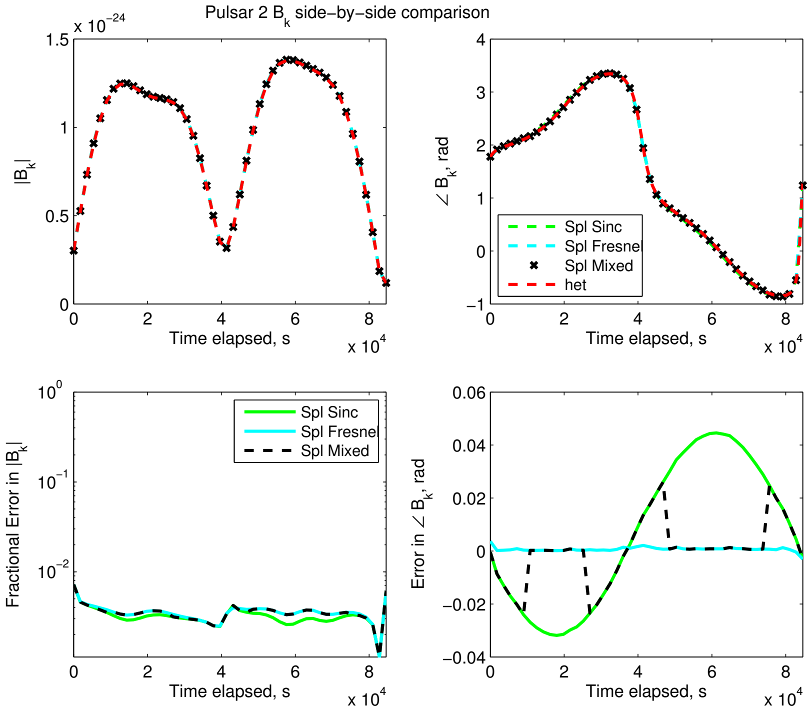

Figure 4.1.3: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR2. We see that there are errors in B_k introduced by the sinc approximation, and that for around half of the time these are significant enough to cause the switch to the Fresnel approximation.

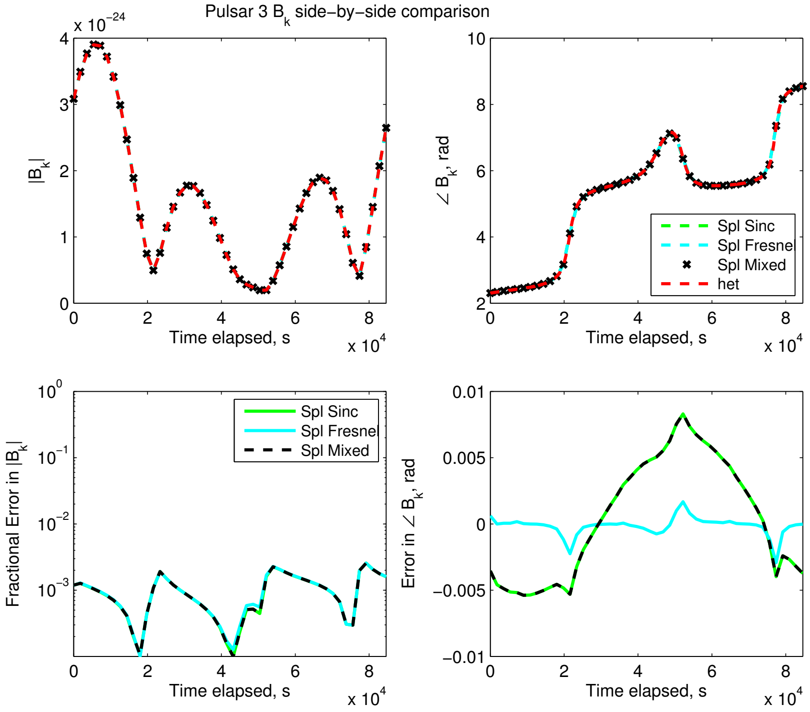

Figure 4.1.4: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR3. We see that there are no significant errors on B_k introduced by the sinc approximation, and we do not need to use the Fresnel approximation for this source.

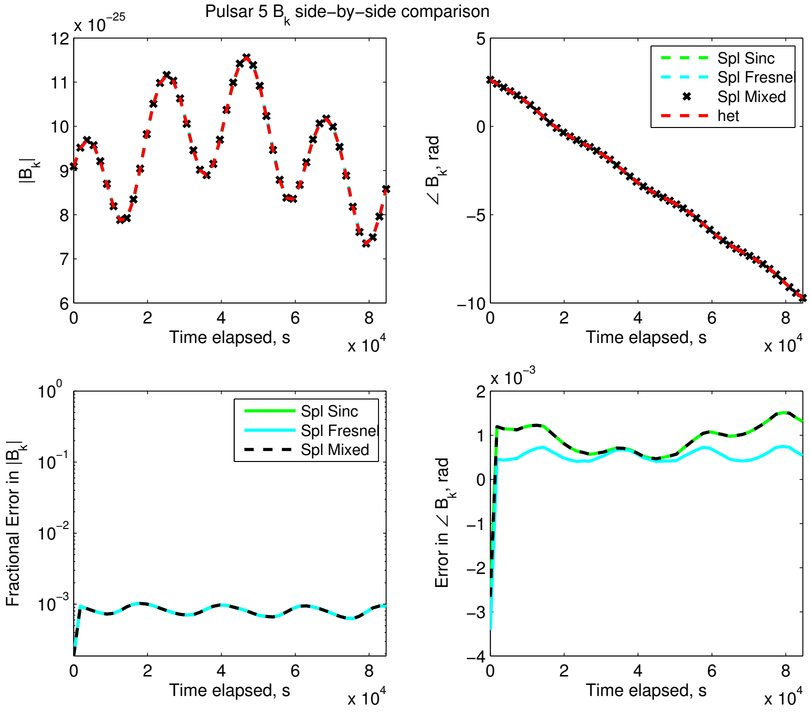

Figure 4.1.5: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR5. We see in this case that there are no significant errors on B_k introduced by the sinc approximation, and that we do not need to use the Fresnel approximation for this source.

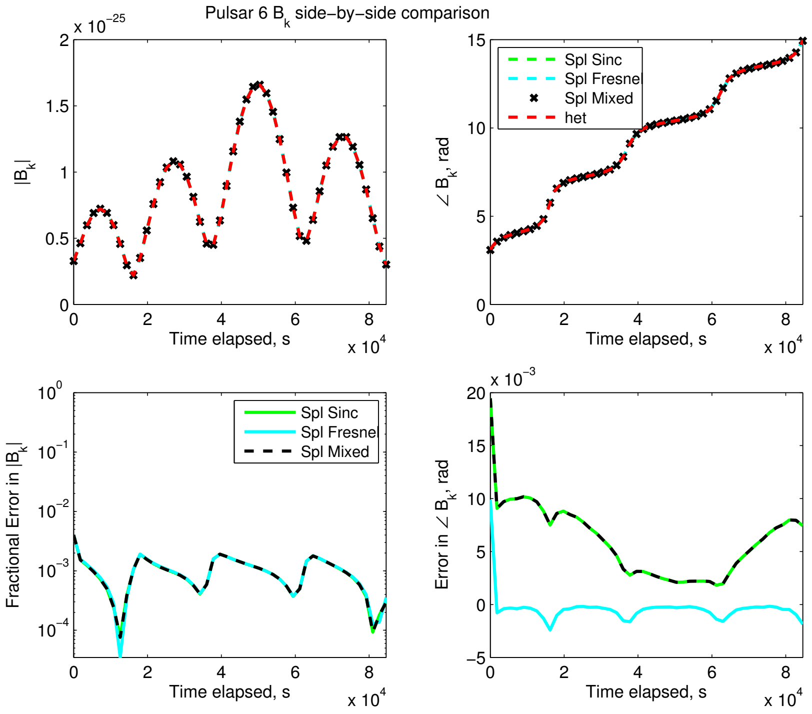

Figure 4.1.6: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR6. We see in this case that there are no significant errors on B_k introduced by the sinc approximation, and that we do not need to use the Fresnel approximation for this source.

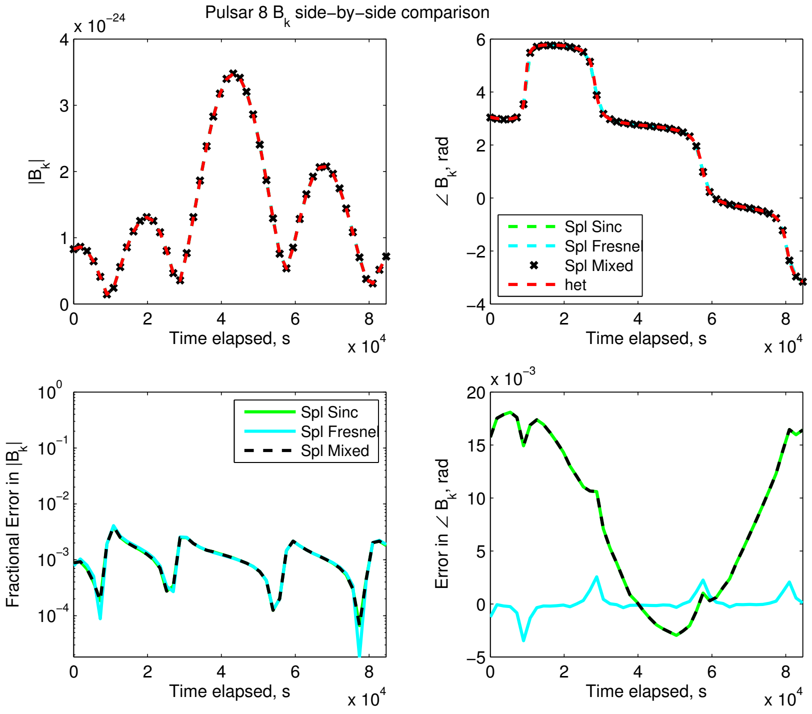

Figure 4.1.7: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR8. We see in this case that there are small errors in B_k introduced by the sinc approximation, but that these are not significant enough to cause the switch to the Fresnel approximation.

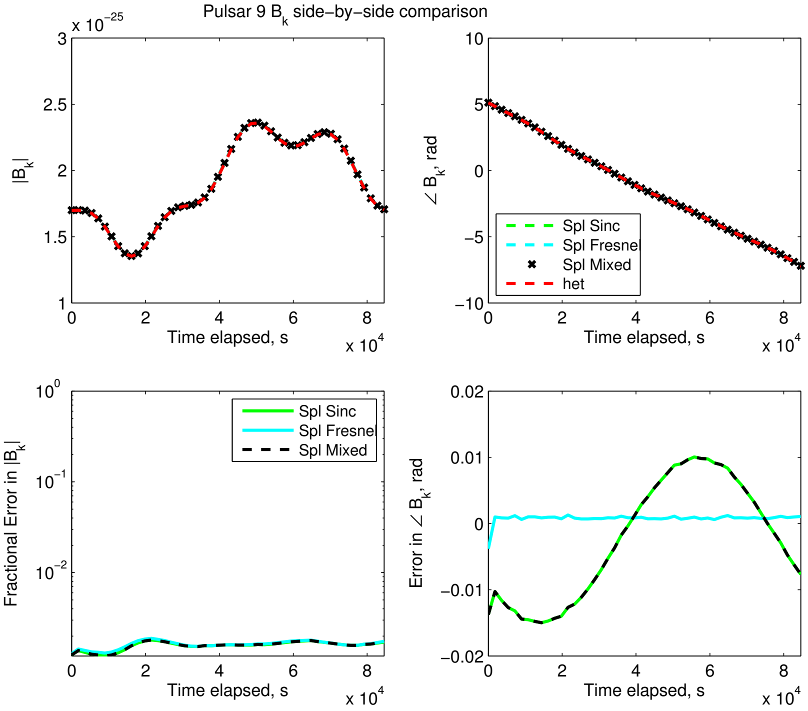

Figure 4.1.8: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR9. We see in this case that there are small errors in $B_{k/\kappa}$ introduced by the sinc approximation, but that these are not significant enough to cause the switch to the Fresnel approximation.

We see that the SplInter and heterodyne $B_{k/\kappa}$ estimates in figures 4.1.1 to 4.1.8 agree very well with one another, and that in the high SNR limit, the SplInter and heterodyne $B_{k/\kappa}$ outputs are almost exactly equal. The Fresnel approximation is used significantly for sources 1 and 2, but not for the others. We also see that the mixed interpolation scheme correctly identifies the points at which it is best to use the Fresnel approximation.

Figures 4.1.9 and 4.1.10 also show a direct comparison between the real and imaginary parts of the SplInter and heterodyne output $B_{k/\kappa}$, as we can directly see in these plots the effect of incorrectly making the assumption that the frequency change during the SFT is zero. These figures show the importance of the Fresnel approximation in comparison to the sinc approximation for sources 4 and 7. These sources have both a high frequency and a relatively low declination, which leads to a large frequency derivative from detector orbital motion.

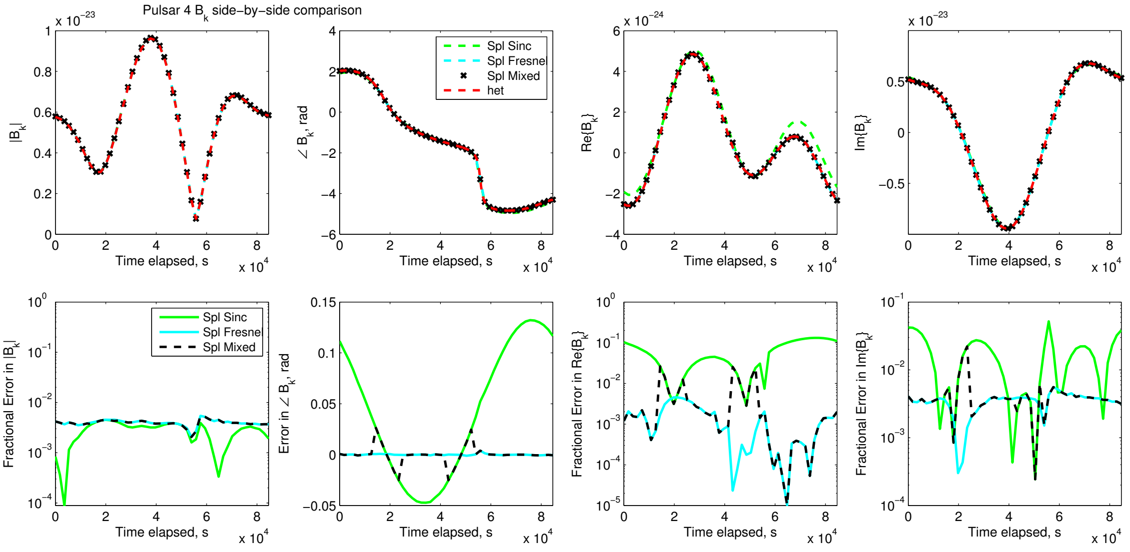

Figure 4.1.9: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR4. This example also includes the real and imaginary parts explicitly so that we can see the errors introduced by the sinc approximation for this source. We see in this case that there are major errors on B_k introduced by the sinc approximation, which are particularly evident in the phase error plot and that we need to use the Fresnel approximation for this source for almost all of the time.

We also see that the Fresnel interpolator is much more accurate than the sinc interpolation in these cases, and that the mixed interpolation scheme correctly identifies points at which it would be best or not to use the sinc interpolator. It should be noted that the algorithm output is an ASCII file containing values of t_k, $\mathfrak{Re}(B_k)$, $\mathfrak{Im}(B_k)$ and σ_k, so an obvious discrepancy in $\mathfrak{Re}(B_k)$, as we see in figure 4.1.9 will cause issues

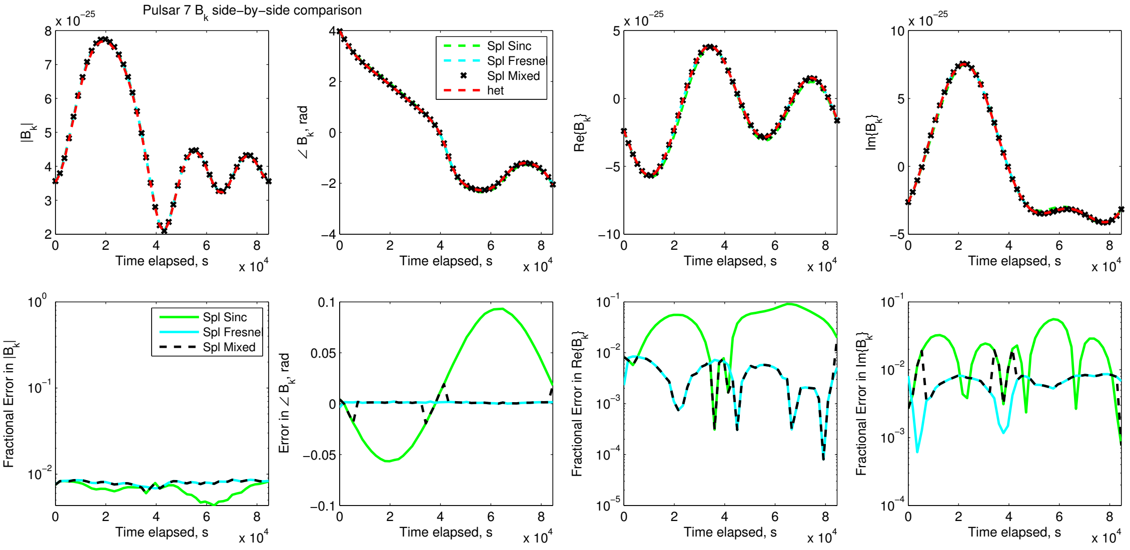

Figure 4.1.10: Comparison of $B_{k/\kappa}$ values from SplInter and heterodyne for a noiseless signal with parameters given by those of hardware injection PULSAR7. We see in this case that there are significant errors on B_k introduced by the sinc approximation, which are particularly evident in the phase error plot and that we need to use the Fresnel approximation for this source for almost all of the time.

in the parameter estimation.

In order to compare the outputs in a more quantitative way, we introduce the concept of mismatch, m, defined by

$$m = \left| 1 - \frac{\sum_k B_{k,\mathrm{Spl}} \cdot B_{\kappa,\mathrm{het}}}{\sum_k B_{\kappa,\mathrm{het}} \cdot B_{\kappa,\mathrm{het}}} \right| . \tag{4.1.2}$$

where $|\ldots|$ denotes absolute value, $B_{k,\mathrm{Spl}}$ is the B_k calculated by the SplInter algorithm, and $B_{\kappa,\text{ het}}$ is the mean value of the B_K calculated from the heterodyne algorithm, within the duration of the SplIntered SFT (the het and Spl subscripts are to clarify the use of each algorithm).

This mismatch shows us the SNR loss caused by using incorrect parameters, or by using approximations in the calculation of B_k. We make the assumption that the heterodyne values are exactly correct in this comparison. For the noiseless injections in figures 4.1.1 to 4.1.10 this mismatch was as given in table 4.3.

Pulsar	0	1	2	3	4	5	6	7	8	9
m_{mixed}	0.0022	0.0028	0.0016	0.0033	0.0040	0.0026	0.0029	0.0024	0.0015	0.0027
m_{Fres}	0.0019	0.0027	0.0017	0.0033	0.0039	0.0026	0.0028	0.0024	0.0015	0.0026
m_{sinc}	0.0022	**0.0052**	0.0018	0.0033	**0.0101**	0.0026	0.0029	**0.0062**	0.0015	0.0027

Table 4.3: Mismatches of noiseless signals using heterodyne and SplInter algorithms with the hardware injection parameters. Marked in bold are the mismatches for which the Fresnel interpolator significantly reduces the sinc interpolator mismatch.

We see that the mismatches are all less than one percent and that the Fresnel and mixed approximations always improve upon the sinc approximation. As such we conclude that the SplInter method is accurate in this recovery.

4.1.2 Noiseless signal recovery - binary pulsar signals

Signals from binary pulsars will contain an extra term in the transformation from the time frame which is inertial with respect to the source to the topocentric time, this term comes from the Roemer, Shapiro and Einstein delays in the source's own orbital system, Δ_{Binary} in equation 2.1.7 above. This manifests as an increased variability in the phase evolution which may not be on timescales such that it can be assumed to be constant or even linear throughout the SFT, with the linear approximation for $\delta(t)$ breaking down. As such, we test the SplInter algorithm using binary signals in comparison to the heterodyne routine, which calculates $\delta(t)$ every second and hence is accurate for all binary pulsars. We present the results of some of these analyses in table 4.4 and figure 4.1.12.

Table 4.4 shows the mismatches over a period of time, for the noiseless signal test as described above. P_{b} is binary period in days, and a_1 is the projected semi-major axis, which is the projection of the orbit onto the line of sight, as in figure 4.1.11.

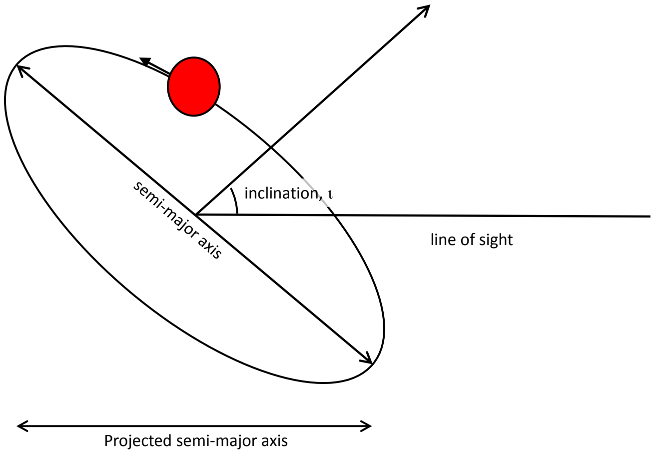

Figure 4.1.11: Diagram explaining projected semi-major axis of a neutron star in a binary system. This binary motion is measured by differences in the time of arrival of pulses from the neutron star, meaning that we are only able to detect line-of sight motion. This means that an orbit on a plane which is inclined with respect to the line of sight can only be seen as the projection of that orbit onto the line of sight, and thus a *projected* semi-major axis is used (binary companion not shown).

The duration of the analysis is chosen to be either a day, if the binary period is less than a day, the binary period, or five days if the binary period is longer than five days,

$$T_{\text{analysis}} = \begin{cases} 1, & P_{\text{b}} \leq 1\,\text{d} \\ P_b, & 1\,\text{d} < P_{\text{b}} \leq 5\,\text{d} \\ 5, & P_{\text{b}} > 5\,\text{d} \end{cases}, \tag{4.1.3}$$

this choice is so that we gain an accurate mismatch over the course of a binary rotation if it has a relatively long period. We choose to analyse the source for the whole of the science run if the mismatch in this noiseless test is below 0.1; this choice is to prevent too large of a drop in SNR, without dismissing too many prospective targets.

The algorithm used to make the signal in this analysis is `lalapps_Makefakedata_v5`, as the previously used `lalapps_create_pulsar_signal_frame` code does not create binary signals.

	Mismatches		Pulsar Parameters				Analyse with
Pulsar	m_{sinc}	m_{mixed}	P_{b}, days	a_1, ls	$\nu^{(0)}$, Hz	$\nu^{(0)} a_1 / P_{\text{b}}^3$	SplInter?
J0023+09	0.74509	0.74535	0.1388	0.0348	327.8470	4.2710×10^{3}	No
J0024−7204E	0.39733	0.39733	2.2568	1.9818	282.7791	48.754	No
J0024−7204H	0.43344	0.43344	2.3577	2.1528	311.4934	51.167	No
J0024−7204I	0.59508	0.59508	0.2298	0.0384	286.9447	909.14	No
J0024−7204Q	0.56030	0.56030	1.1891	1.4622	247.9432	215.64	No
J0024−7204R	0.93394	0.93394	0.0662	0.0334	287.3181	3.2996×10^{4}	No
J0024−7204S	0.52040	0.52040	1.2017	0.7663	353.3062	156.00	No
J0024−7204T	0.38153	0.38153	1.1262	1.3385	131.7787	123.49	No
J0024−7204U	0.78311	0.78311	0.4291	0.5269	230.2648	1.5357×10^{3}	No
J0024−7204Y	0.77959	0.77959	0.5219	0.6686	455.2372	2.1406×10^{3}	No
J0034−0534	0.57837	0.58884	1.5893	1.4378	532.7134	190.80	No
J0218+4232	0.48940	0.48943	2.0288	1.9844	430.4611	102.29	No
J0407+1607	0.00080	0.00080	669.0702	106.4501	38.9079	1.3828×10^{-5}	Yes
J0437−4715	0.01168	0.01168	5.7410	3.3667	173.6879	3.0903	Yes
J0605+3757	0.00166	0.00191	55.6723	18.9487	366.5751	4.0255×10^{-2}	Yes
J0610−2100	0.60421	0.60511	0.2860	0.0735	258.9785	813.42	No
J0613−0200	0.69872	0.69655	1.1985	1.0914	326.6006	207.06	No
J0614−3329	0.00102	0.00202	53.5846	27.6388	317.5945	5.7052×10^{-2}	Yes
J0621+1002	0.00033	0.00033	8.3187	12.0321	34.6574	7.2439×10^{-1}	Yes

Continued

	Mismatches		Pulsar Parameters				Analyse with
Pulsar	m_{sinc}	m_{mixed}	P_{b}, days	a_1, ls	$\nu^{(0)}$, Hz	$\nu^{(0)} a_1 / P_{\mathrm{b}}^3$	SplInter?
J0737−3039A	0.98665	0.98665	0.1023	1.4147	44.0541	5.8296×10^{4}	No
J0751+1807	0.83539	0.83557	0.2631	0.3966	287.4579	6.2569×10^{3}	No
J0900−3144	0.00024	0.00024	18.7376	17.2488	90.0118	2.3600×10^{-1}	Yes
J1012+5307	0.63814	0.63814	0.6047	0.5818	190.2678	500.72	No
J1017−7156	0.08106	0.08106	6.5119	4.8300	427.6219	7.4798	Yes
J1022+1001	0.00780	0.00780	7.8051	16.7654	60.7794	2.1430	Yes
J1045−4509	0.01869	0.01869	4.0835	3.0151	133.7931	5.9242	Yes
J1231−1411	0.32977	0.33020	1.8601	2.0426	271.4530	86.148	No
J1300+1240	0.00193	0.00193	66.3366	0.0013	160.8097	7.4262×10^{-7}	Yes
J1455−3330	0.00138	0.00138	76.1746	32.3623	125.2002	9.1667×10^{-3}	Yes
J1518+4904	0.00048	0.00048	8.6340	20.0440	24.4290	7.6077×10^{-1}	Yes
J1537+1155	0.71289	0.71289	0.4207	3.7305	26.3821	1.3214×10^{3}	No
J1600−3053	0.00501	0.00532	14.3485	8.8017	277.9377	8.2813×10^{-1}	Yes
J1603−7202	0.00412	0.00412	6.3086	6.8807	67.3766	1.8464	Yes
J1614−2230	0.06804	0.06675	8.6866	11.2912	317.3789	5.4672	Yes
J1623−2631	0.00089	0.00089	191.4428	64.8094	90.2873	8.3396×10^{-4}	Yes
J1630+3734	0.00781	0.00989	12.5250	9.0393	301.3762	1.3865	Yes
J1640+2224	0.00134	0.00206	175.4607	55.3297	316.1240	3.2380×10^{-3}	Yes
J1643−1224	0.00082	0.00092	147.0174	25.0726	216.3733	1.7073×10^{-3}	Yes
J1701−3006A	0.06797	0.06795	3.8059	3.4837	190.7827	12.056	Yes
J1709+2313	0.00067	0.00043	22.7119	15.2885	215.9269	2.8178×10^{-1}	Yes
J1713+0747	0.00095	0.00114	67.8251	32.3424	218.8118	2.2681×10^{-2}	Yes
J1719−1438	0.04856	0.04834	0.0907	0.0018	172.7070	423.30	Yes
J1731−1847	0.71455	0.71501	0.3111	0.1202	426.5193	1.7016×10^{3}	No
J1732−5049	0.02627	0.02627	5.2630	3.9829	188.2335	5.1427	Yes
J1738+0333	0.73157	0.73156	0.3548	0.3434	170.9374	1.3145×10^{3}	No
J1741+1351	0.00286	0.00417	16.3353	11.0033	266.8692	6.7365×10^{-1}	Yes
J1745−0952	0.00010	0.00010	4.9435	2.3786	51.6094	1.0162	Yes
J1748−2446A	0.93057	0.93057	0.0756	0.1197	86.4816	2.3915×10^{4}	No
J1748−2446E	0.00021	0.00116	60.0598	23.5959	455.0004	4.9556×10^{-2}	Yes
J1748−2446I	0.33139	0.33139	1.3278	1.8188	104.4911	81.184	No
J1748−2446M	0.78392	0.78388	0.4431	0.5964	280.1457	1.9200×10^{3}	No
J1748−2446N	0.82689	0.82689	0.3855	1.6192	115.3815	3.2619×10^{3}	No
J1748−2446O	0.81612	0.81725	0.2595	0.1118	596.4354	3.8171×10^{3}	No
J1748−2446Q	0.00082	0.00008	30.2954	28.6420	355.6447	3.6635×10^{-1}	Yes

Continued

	Mismatches		Pulsar Parameters				Analyse with
Pulsar	m_{sinc}	m_{mixed}	P_{b}, days	a_1, ls	$\nu^{(0)}$, Hz	$\nu^{(0)}a_1/P_{\mathrm{b}}^3$	SplInter?
J1748−2446U	0.43528	0.43455	3.5703	5.9725	304.0301	39.900	No
J1748−2446V	0.80522	0.80281	0.5036	0.5669	482.5079	2.1412×10^{3}	No
J1748−2446W	0.11067	0.11065	4.8768	5.8837	237.8018	12.063	No
J1748−2446X	0.14808	0.14765	4.9985	5.1072	333.4158	13.635	No
J1748−2446Y	0.74824	0.74844	1.1644	1.1785	488.2434	364.44	No
J1748−2446Z	0.33386	0.33175	3.4881	3.5304	406.0764	33.782	No
J1748−2446ae	0.73708	0.73703	0.1707	0.0406	273.3295	2.2311×10^{3}	No
J1748−2446ai	0.65995	0.65995	0.8509	2.8089	47.1067	214.79	No
J1751−2857	0.00061	0.00067	110.7465	32.5282	255.4361	6.1172×10^{-3}	Yes
J1756−2251	0.81752	0.81752	0.3196	2.7576	35.1351	2.9669×10^{3}	No
J1801−3210	0.00048	0.00048	20.7717	7.8093	134.1636	1.1690×10^{-1}	Yes
J1802−2124	0.72593	0.72593	0.6989	3.7189	79.0664	861.34	No
J1804−0735	0.01689	0.01689	2.6168	3.9205	43.2884	9.4716	Yes
J1804−2717	0.00097	0.00097	11.1287	7.2815	107.0317	5.6545×10^{-1}	Yes
J1807−2459A	0.82456	0.82586	0.0711	0.0122	326.8563	1.1119×10^{4}	No
J1810+17	0.90949	0.90984	0.1482	0.0954	601.4115	1.7637×10^{4}	No
J1810−2005	0.00068	0.00068	15.0120	11.9779	30.4671	1.0787×10^{-1}	Yes
J1811−2405	0.09974	0.09831	6.2723	5.7057	375.8560	8.6905	Yes
J1841+0130	0.00075	0.00075	10.4716	3.5041	33.5877	1.0250×10^{-1}	Yes
J1853+1303	0.00027	0.00059	115.6538	40.7695	244.3914	6.4408×10^{-3}	Yes
J1857+0943	0.00377	0.00414	12.3272	9.2308	186.4941	9.1900×10^{-1}	Yes
J1903+0327	0.00134	0.00321	95.1741	105.5935	465.1352	5.6972×10^{-2}	Yes
J1909-3744	0.70448	0.70446	1.5334	1.8980	339.3157	178.60	No
J1910+1256	0.00050	0.00061	58.4667	21.1291	200.6588	2.1214×10^{-2}	Yes
J1910−5959A	0.72597	0.72613	0.8371	1.2060	306.1674	629.46	No
J1911−1114	0.11986	0.11957	2.7166	1.7629	275.8053	24.253	No
J1918−0642	0.00261	0.00261	10.9132	8.3505	130.7895	8.4029×10^{-1}	Yes
J1955+2908	0.00076	0.00076	117.3491	31.4126	163.0479	3.1694×10^{-3}	Yes
J1959+2048	0.67574	0.68094	0.3820	0.0892	622.1220	995.84	No
J2017+0603	0.46199	0.46392	2.1985	2.1929	345.2781	71.256	No
J2019+2425	0.00058	0.00080	76.5116	38.7677	254.1603	2.1999×10^{-2}	Yes
J2033+17	0.00089	0.00089	56.3078	20.1631	168.0967	1.8985×10^{-2}	Yes
J2043+1711	0.63781	0.64082	1.4823	1.6240	420.1894	209.52	No
J2051−0827	0.85634	0.85636	0.0991	0.0451	221.7963	1.0270×10^{4}	No
J2129−5721	0.01439	0.01439	6.6255	3.5006	268.3592	3.2300	Yes

Continued

	Mismatches		Pulsar Parameters				Analyse with
Pulsar	m_{sinc}	m_{mixed}	P_{b}, days	a_1, ls	$\nu^{(0)}$, Hz	$\nu^{(0)}a_1/P_{\text{b}}^3$	SplInter?
J2140−2310A	0.78213	0.78213	0.1740	0.2350	90.7496	4.0495×10^{3}	No
J2145−0750	0.00523	0.00523	6.8389	10.1641	62.2959	1.9795	Yes
J2214+3000	0.29874	0.30458	0.4166	0.0591	320.5923	261.90	No
J2215+5135	0.94816	0.94819	0.1725	0.4681	383.1976	3.4948×10^{4}	No
J2229+2643	0.00350	0.00402	93.0159	18.9125	335.8162	7.8919×10^{-3}	Yes
J2241−5236	0.75590	0.75649	0.1457	0.0258	457.3101	3.8161×10^{3}	No
J2302+4442	0.00104	0.00104	125.9353	51.4300	192.5920	4.9592×10^{-3}	Yes
J2317+1439	0.30028	0.30065	2.4593	2.3139	290.2546	45.152	No

Table 4.4: Mismatches of noiseless signals from targeted sources in binary systems, using the sinc approximation, m_{sinc}, and the mixed interpolation scheme, m_{mixed}. Also shown is binary period, P_{b}, projected semi-major axis, a_1, frequency, $\nu^{(0)}$, and a value indicating the extent of the $\ddot{f}_k$ from the binary motion, $\nu^{(0)}a_1/P_{\text{b}}^3$. We also state whether this source can be analysed with the SplInter algorithm, based on the mismatch from the mixed interpolation scheme.

In table 4.4, from an analysis of 97 signals[1] with parameters from real pulsars, 50 were sources for which we are unable to use the SplInter algorithm (51.5%), and of the remaining 47, all could be analysed using the sinc approximation as the combination of source frequency, binary period, and projected semi-major axis did not cause significant frequency change during the SFT. We see from these results that the use of the Fresnel or sinc interpolator in this analysis does not produce particularly different results. We shall see the reason for this in section 4.1.3.

From the results of table 4.4, we plot figure 4.1.12, in which we see that there is a cut-off point between the analysed and not analysed pulsars. This is not a sharp, distinct cut-off as the ability to analyse the data is dependent upon $\ddot{f}_k$, which will be dependent upon frequency. By doing a back-of-the-envelope calculation of the time delay caused by the binary motion, we find that

$$\ddot{f}_k \propto \frac{\nu^{(0)}a_1}{P_{\text{b}}^3}, \tag{4.1.4}$$

which is included in table 4.4. Through the examples of the pulsars used in this test,

---

[1] These are the 97 binary pulsars in [2].

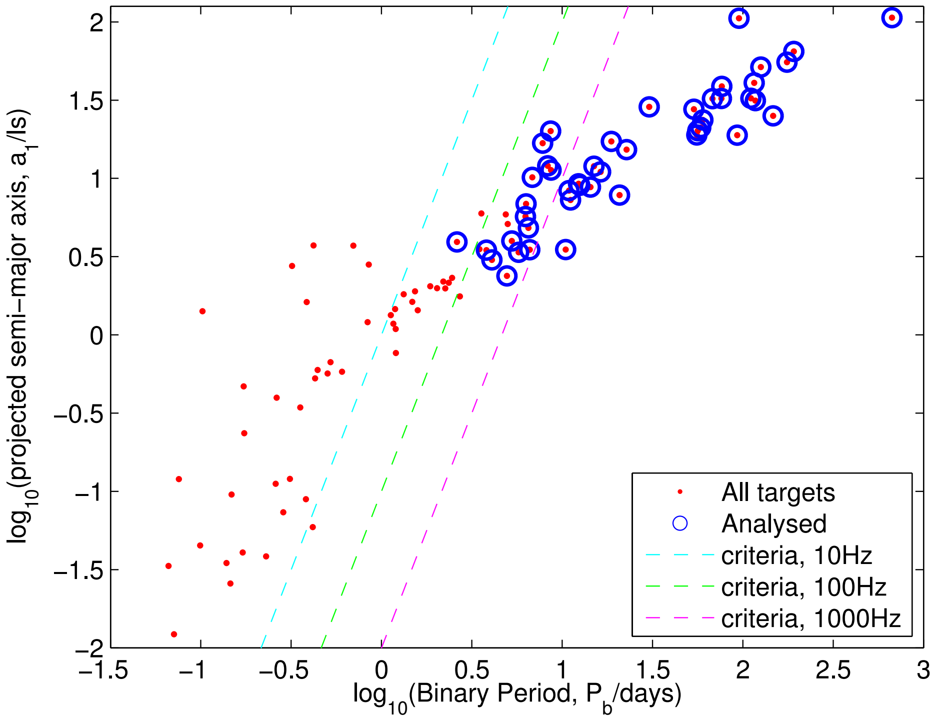

Figure 4.1.12: Binary period vs projected semi-major axis for targeted binary pulsars, indicating which binary pulsars can be analysed using SplInter and which cannot. We include an indication of the empirical criteria we set for analysis of a target in a binary system, given in equation 4.1.5, for pulsars with source frequency of 10, 100 and 1000 Hz.

we can empirically find that the cut-off (in the units as given in table 4.4) is

$$\frac{\nu^{(0)}}{1\,\mathrm{Hz}}\frac{a_1}{1\,\mathrm{ls}}\left(\frac{P_\mathrm{b}}{1\,\mathrm{day}}\right)^{-3} \lesssim 10. \tag{4.1.5}$$

Note that pulsar J1701−3006A, with $\nu^{(0)}a_1/P_\mathrm{b}^3 = 12.056\,\mathrm{Hz\,ls\,day}^{-3}$, does not pass this criterion but can be analysed based on the results of table 4.4, so we choose a cut-off of 10 to be conservative.

We have a list of targets which can and cannot be used in table 4.4. Though this is useful for these targets, setting an approximate criterion such as in equation 4.1.5 can then be used if we are to analyse candidate continuous-wave signals from pulsars in binary systems, such as from the TwoSpect analysis [46]. The TwoSpect minimum binary period of $P_\mathrm{b} > 2\,\mathrm{h}$ is well below the search limits in this work. This is because the TwoSpect search is able to dynamically choose the SFT length based on the binary period, something we do not do in this search.

4.1.3 Frequency derivative assumptions

We have said before in section 3.1.1 that for an isolated source, $\dot{f}_k$ is expected to be very small and, as such, we can assume $\dot{f}_k = 0$, the sinc approximation. This section looks at values of $\dot{f}_k$ for which this assumption is valid, and the maximum values of $\ddot{f}_k$ which can be ignored. We will also discuss what can be done in the cases where the values of $\dot{f}_k$ and $\ddot{f}_k$ are above these thresholds and whether these maximum values are physical.

This test will use a model for the tests where the intrinsic spindown parameters $f^{(1)}$ and $f^{(2)}$ are kept as constants and $\delta(t)$ is kept as zero; meaning that we have better control over the value of the effective spindown parameters $\dot{f}_k = f^{(1)} + f^{(2)}(t - t_k)$ and $\ddot{f}_k = f^{(2)}$. We do not need to compare the analysis to the heterodyned output in this case, as we can use the known input value.

Assuming linear phase evolution during the SFT, $\dot{\mathbf{f}}_{\mathbf{k}} = \mathbf{0}$.

Firstly we consider the $\dot{f}_k = 0$ assumption, we set $f^{(2)} = 0$ so that $\dot{f}_k = f^{(1)}$ and see how the output is affected by increasing the value of $f^{(1)}$. Shown in figure 4.1.13 is the mismatch introduced by this assumption, testing signals with $\dot{f}_k$ ranging from 2×10^{-9} to $2 \times 10^{-4}\,\mathrm{Hz\,s^{-1}}$.

We see that the Fresnel interpolation model accurately recovers the signal for all input $\dot{f}_k$ and that the sinc function interpolator loses accuracy for signals with $\dot{f}_k \gtrapprox 10^{-8}\,\mathrm{Hz\,s^{-1}}$ when using 1800 s SFTs, this means that the signal is spreading too much over the course of the SFT. The important factor to consider is the signal spread in terms of number of bins, which we see as the signals within the 60 s SFTs, with much wider frequency bins and less time for the frequency to change, are being recovered using the sinc interpolation many orders of magnitude higher than the half hour SFTs. We therefore use the signal spread per bin as the switching point from one algorithm to the other.

The vertical lines of figure 4.1.13 indicate where this switch happens – this was chosen to be approximately the point where the phase error introduced by the sinc approximation is 0.05 radians, corresponding to a mismatch of 0.01, or a signal frequency spreading more than 10% of a bin during the SFT, so the sinc interpolator is used for

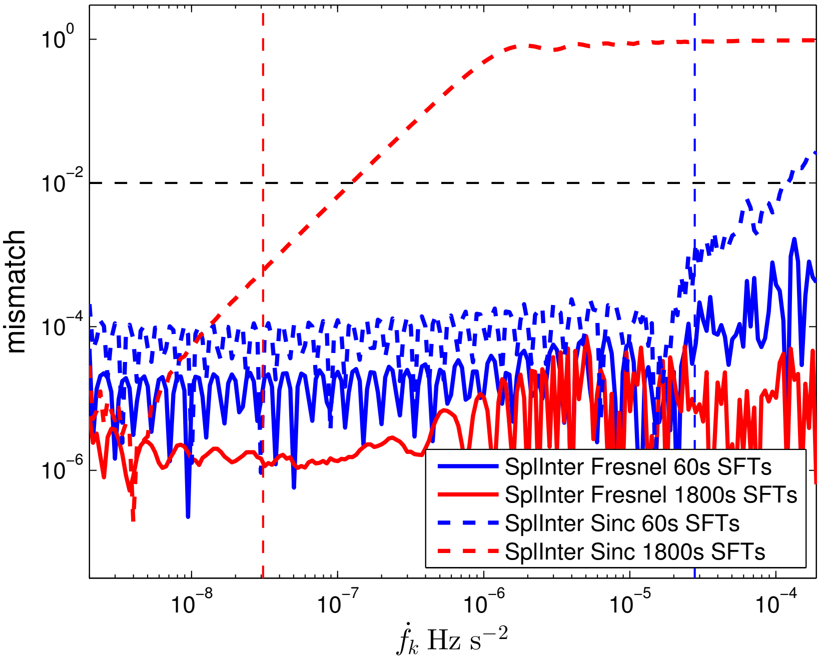

Figure 4.1.13: Plot showing the mismatch when using each interpolation model in the maximum posterior estimator with increasing values of $\dot{f}_k$ for 60 s and half hour SFTs. Also shown are vertical lines corresponding to the switching point of $\dot{f}_k = 0.1/\Delta t^2$ and a horizontal line showing the maximum allowed mismatch.

$|\dot{f}_k|\Delta t^2 < 0.1$.

The reason for not using the Fresnel interpolation model for all cases is computational efficiency – it takes more computational effort to calculate the Fresnel integrals (one numerical integration per frequency bin) than it does to calculate the sinc function, and a mismatch of 0.01 is the point at which we decide that the accuracy error overcomes the speed increase. Another reason is that for $\dot{f}_k = 0$, the calculation of the Fresnel model would lead to divide-by-zero errors.

Assuming linear frequency evolution during the SFT, $\ddot{\mathbf{f}}_\mathbf{k} = \mathbf{0}$.

Next we consider $\ddot{f}_k$, we do this by setting $f^{(1)} = 0$ so that $\ddot{f}_k = f^{(2)}$ and $\dot{f}_k = f^{(2)}(t_k - t_0)$, we use t_k equal to the parameter file epoch such that $\dot{f}_k = 0$ at the centre of the SFT. The interpolation assumes $\ddot{f}_k$ to always be zero, and so figure 4.1.14 shows the increased mismatch associated with a non-zero $\ddot{f}_k$ for values ranging from $\ddot{f}_k = 2 \times 10^{-10}$ to 2×10^{-8} Hz s^{-2}.

We see in figure 4.1.14 that the interpolation for this case breaks down at $\ddot{f}_k \approx 0.01/\Delta t^3$, as such we apply a cut-off threshold at this value, above which the algorithm ignores that source in the SFT.

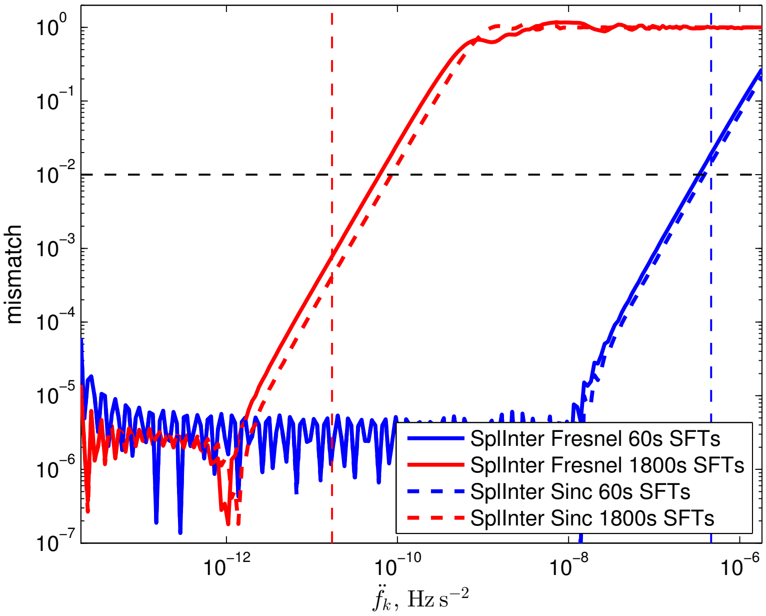

Figure 4.1.14: Plot showing the mismatch when using each interpolation model in the maximum posterior estimator with increasing values of $\ddot{f}_k$ for 60 s and half hour SFTs. Also shown are vertical lines corresponding to the cut-off point of $\ddot{f}_k = 0.01/\Delta t^3$ and a horizontal line showing the maximum allowed mismatch.

If we look at table 2.1, the highest intrinsic value of $\nu^{(1)}$, ($= f^{(1)}/2$ for a non-precessing, triaxial neutron star) for isolated pulsars is around $10^{-10}\,\mathrm{Hz\,s^{-1}}$, and for $\nu^{(2)}$ is $10^{-20}\,\mathrm{Hz\,s^{-2}}$, both of which are still well within the range of the sinc interpolator, but we need to consider the relative motion effects on the signal phase evolution.

The orbital motion of the detector can introduce significant values of $\dot{f}_k$ and $\ddot{f}_k$. For isolated sources, the $\dot{\delta}_k$ term is dominated by the Earth's orbit around the Sun, the 'worst case scenario' for $\dot{f}_k$ is therefore a source on the ecliptic (the plane of orbit of the Earth around the Sun). The $\ddot{\delta}_k$ term is dominated by the diurnal motion, and the 'worst case scenario' for $\ddot{f}_k$ is therefore a source on the celestial equator (i.e. zero declination). In these worst case scenarios, the value of $\dot{\delta}_k$ can be up to $\sim 10^{-4}$ and $\ddot{\delta}_k$ can be up to $\sim 10^{-10}\,\mathrm{s^{-1}}$. Depending on the frequency of the source this will introduce significant values of $\dot{f}_k$, and $\ddot{f}_k$. Shown in figure 4.1.15 is the change in time of arrival δ_k, barycentred frequency f_k and introduced $\dot{f}_k$ and $\ddot{f}_k$ for an isolated source with intrinsic frequency of $f^{(0)} = 2\,000\,\mathrm{Hz}$, at the vernal point, so the worst case scenario for Earth orbit and rotation, and around the maximum allowed frequency by the SFTs.

As we see, this worst case scenario will give us $\dot{f}_k$ values up to $\sim 10^{-7}\,\mathrm{Hz\,s^{-1}}$, well above the acceptable range for the sinc approximation, and $\ddot{f}_k$ up to $\sim 10^{-11}\,\mathrm{Hz\,s^{-2}}$,

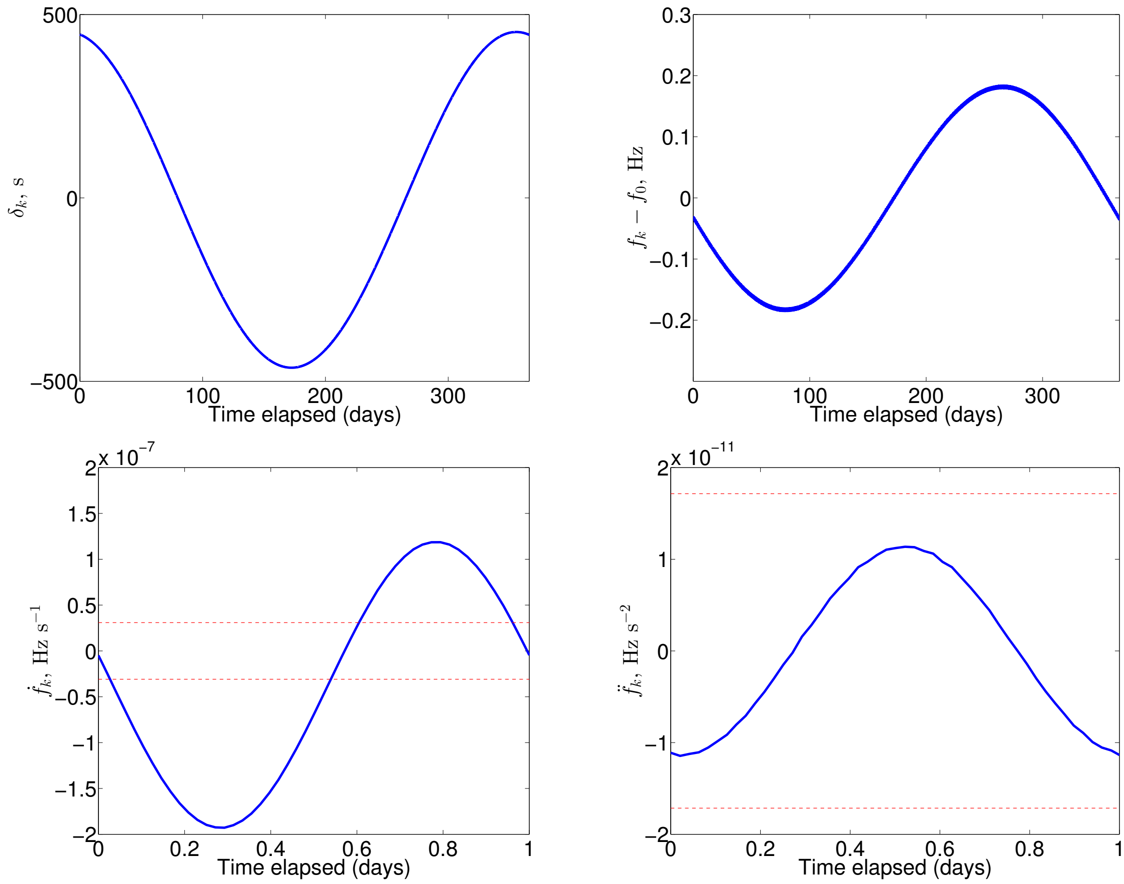

Figure 4.1.15: Change in arrival time (top left), frequency (top right), and its first (bottom left) and second (bottom right) derivatives introduced by the effects of orbital motion of the detector for a source with $f^{(0)} = 2\,000\,\mathrm{Hz}$ at the vernal point. The change in arrival time and frequency are dominated by the Earth orbital motion, and so are plotted over the course of a year, while the frequency derivatives are dominated by the Earth rotation, so are plotted over the course of a day. Included in the bottom two figures of $\dot{f}_k$ and $\ddot{f}_k$ are horizontal lines showing the limits from figures 4.1.13 and 4.1.14 respectively.

which is within the range of assuming $\ddot{f}_k = 0$ with half hour SFTs. This means that for half-hour SFTs, the Fresnel approximation should be sufficient for all isolated sources, and the sinc approximation should be sufficient in most cases.

4.1.4 Choice of windowing function

This work has so far considered only a rectangular (equal weighting) window applied to the time-domain data during the creation of the FFT. However the commonly used window within the LIGO Scientific Collaboration continuous-waves group is the Tukey window. The Tukey window [13, 95] is used in order to suppress the first and last few samples of the time series, which will have been affected by the high-pass filters used to remove low frequency noise in the creation of the SFTs. Very few samples are affected by this, and so the amount of windowing is small – the fraction of the window affected in standard half-hour SFTs is $R = 0.001$, which corresponds to the first 1.8 s. The Tukey window also generally provides faster sidelobe reduction than a rectangular window, but for values of R as small as those used here this is negligible.

The Tukey window takes the form

$$W_T = \begin{cases} 1, & t_S + \frac{R\Delta t}{2} \leq t < t_E - \frac{R\Delta t}{2} \\ \frac{1}{2}\left\{1 + \cos\left[2\pi\left(\frac{t-t_S}{R\Delta t} - \frac{1}{2}\right)\right]\right\}, & t_S \leq t < t_S + \frac{R\Delta t}{2} \\ \frac{1}{2}\left\{1 + \cos\left[2\pi\left(\frac{t-t_E}{R\Delta t} + \frac{1}{2}\right)\right]\right\}, & t_E - \frac{R\Delta t}{2} \leq t < t_E \\ 0, & \text{elsewhere} \end{cases}, \qquad (4.1.6)$$

for which the special case $R = 0$ is a rectangular window and $R = 1$ is known as a Hann window.

The testing here considers using the rectangular window when a Tukey window has been applied, and seeing the value of R at which we must make changes to the estimation algorithm.

To test this, we made SFTs containing signal with $f^{(1)} = 0$, $f^{(2)} = 0$ with an imagined detector at the solar system barycentre with uniform antenna pattern. We altered R in the range $[10^{-6}, 1]$. We then calculated B_k and compared this output value to the injected y_k. Figure 4.1.16 shows the increasing mismatch of the signal in the B_k

when assuming $R = 0$ for increasing values of R.

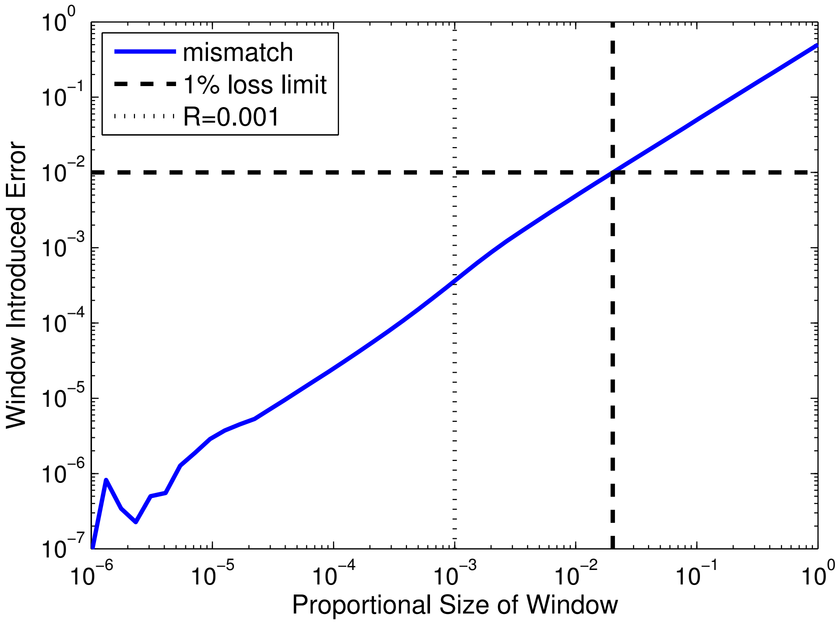

Figure 4.1.16: Plot showing the increase in mismatch when using a rectangular window model for a given windowing amount of a Tukey window. Also indicated are the $R = 0.001$ line, which is the standardly used Tukey window for the LIGO continuous-waves SFTs, and the acceptable mismatch limit, here defined as a 10^{-2}, corresponding to an upper limit of $R = 0.02$. Below a value of $R \sim 10^{-5}$ the mismatch is dominated by numerical precision noise, but is small enough to be safely ignored.

We see that mismatch increases significantly with large amounts of windowing, and that if we define an acceptable mismatch of 1%, this corresponds to an upper limit of $R = 0.02$. Considering this maximum value and that the typically used value of $R = 0.001$ is well below this limit, we conclude that windowing effects will not significantly affect the results.

4.1.5 Noise estimation tests

To test the noise estimation in the SplInter algorithm, we used two tests; the first was to make SFTs and frames with known white noise and no signal. After running the SplInter and heterodyne algorithms, we would check that the $B_{k/K}$ noise was consistent with the injected value, and also with each other. We compare noise estimates from the SplInter routine, σ_k, with noise estimates from B_K values from the heterodyne routine for a signal with zero amplitude with time-domain variance of $\sigma_T^2 = 1$. The estimate of the noise on the heterodyne, σ_H, is made by using an average of the standard deviations from the real and imaginary heterodyne parts which fall within the SFT time range. To then convert into the equivalent noise value for a half hour separation of B_k, we divide by a factor of the square root of the number of datapoints used, as is consistent with equation 2.3.9, which will be thirty if all B_K are present. Algebraically, this is

$$\sigma_H = \frac{1}{2(N_K - 1)}\left\{\sqrt{\sum_K \mathfrak{Re}[B_K(t_S < t_K < t_E)]^2} + \sqrt{\sum_K \mathfrak{Im}[B_K(t_S < t_K < t_E)]^2}\right\}, \quad (4.1.7)$$

where N_K is the number of heterodyne datapoints within the range of the SFT. We see in figure 4.1.17 that the heterodyne and SplInter noise estimates agree with each other, and that these both agree with the expected distribution around the injected value of the noise. The expected distribution is a χ^2 distribution with $n - 1$ degrees of freedom, where n is the number of datapoints used in the noise estimation,

$$\frac{\sigma_k^2}{\sigma_{\mathrm{True}}^2} \sim \frac{\chi^2(n-1)}{n-1}. \quad (4.1.8)$$

and σ_{True} is the true injected noise. The heterodyne noise estimate used thirty B_K datapoints from each of the real and imaginary parts of the data, leading to an expected χ^2 distribution with 59 degrees of freedom, shown in the figure by the red dotted line. In this test, we used the Spectral Interpolation algorithm with a bandwidth of 0.3 Hz around the signal frequency, this meant that one thousand and eighty datapoints were used, ($0.3 \times \frac{1}{1800}$ from each of the real and imaginary parts), leading to a χ^2 distribution with 1079 degrees of freedom (shown on the figure as a blue dotted line). This tells us that in order to get better noise estimates, we could use a wider band of data. However, the frequency dependence of the noise leads to a limiting bandwidth, as well as computational efficiency considerations. A frequency bandwidth of 0.3 Hz is a good balancing point between these considerations.

For the second test software injections are used, selecting the lowest SNRs, so that the signal should not affect the noise levels estimated for the heterodyne routine too much. We compare the σ_H and σ_k data obtained from the two routines by taking the maximum likelihood for a scale factor denoting the difference between the noise

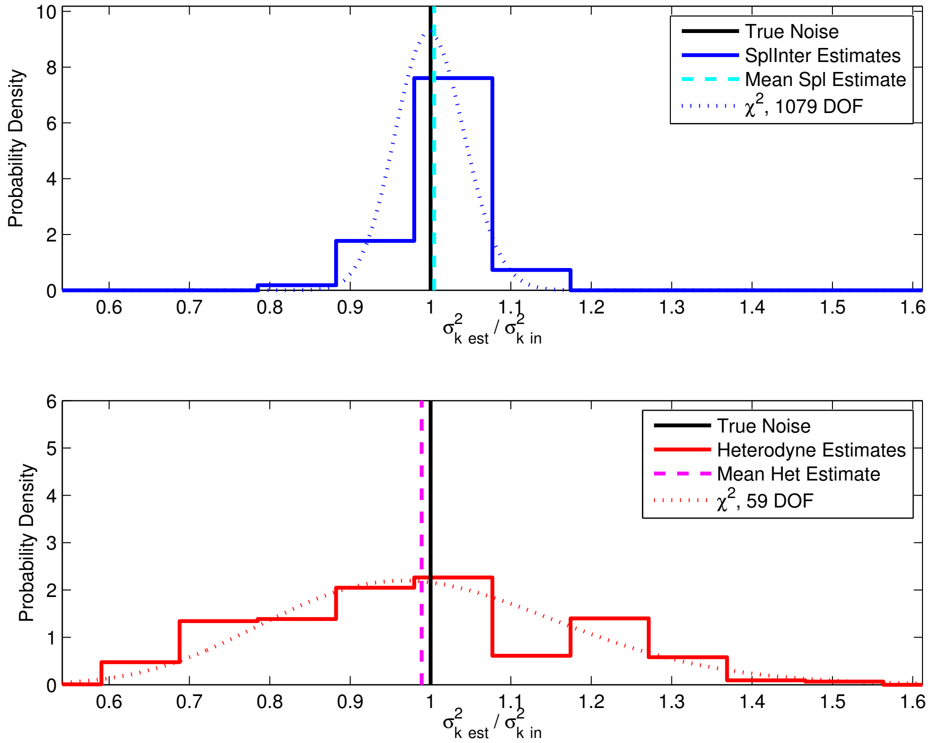

Figure 4.1.17: A histogram of standard deviation estimates of white noise from SplInter (top) and heterodyne (bottom) output, with the mean estimated values (cyan/magenta dashed lines). Also shown is the true value of the noise (black vertical line) and the expected distributions of the noise estimates (dotted lines).

estimates, β, defined by

$$\mathrm{p}(\beta) \propto \exp\left\{-\sum_k \left[\frac{1}{2}\frac{(\sigma_H\beta - \sigma_k)^2}{\sigma_\sigma^2}\right]\right\}, \tag{4.1.9}$$

σ_σ is an estimate of the noise on the standard deviation, calculated from B_K. Figure 4.1.18 shows the most probable values of β obtained from the software injections, $\beta_{\max}$, and we see that the noise estimation data from the Spectral Interpolation algorithm is consistent with the noise on the B_K values calculated from the heterodyne routine, almost entirely lying within a $\sim 4\%$ band around the desired value of unity, with unity lying within the 1σ uncertainty on this value.

One software injection (J0154+4819) showed a value of $\beta = 0.61 \pm 0.05$, which is significantly lower than unity, meaning that the SplInter noise estimate was a long way below the noise estimate from the heterodyne. The injection frequency of ~ 347.7 Hz meant that the signal was within the range of strong lines caused by resonant violin modes of the suspension (see section 1.3.1). Figure 4.1.19 shows that the outlier removal routine in SplInter is working to actively remove the noisy lines, leading to a reduction

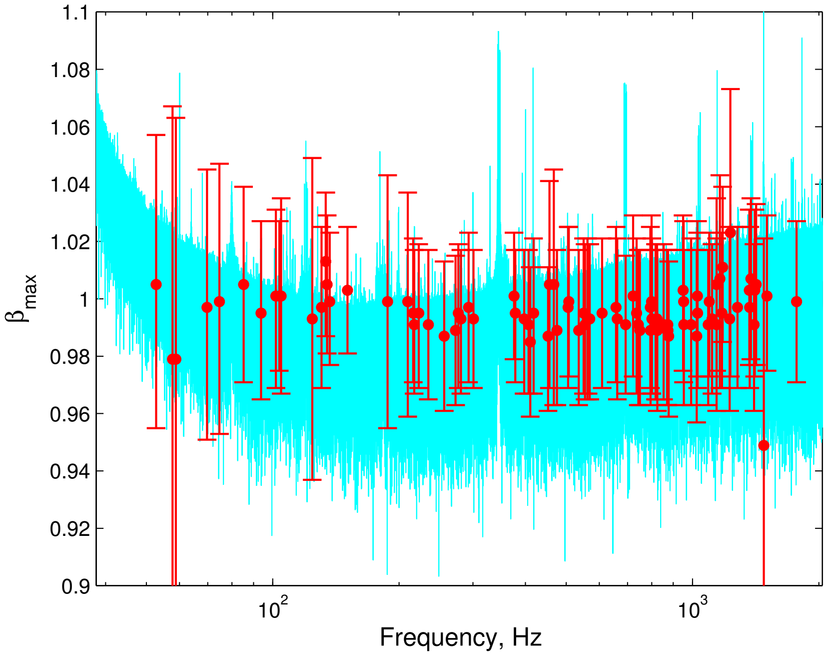

Figure 4.1.18: Estimates, with errors, of $\beta_{\max}$, the most likely ratio between the SplInter noise estimate and the heterodyne noise, estimated from the standard deviation of the B_K values, with an example SFT in the background for reference. Software injections with very low SNR were chosen so that the signal in $B_{k/\kappa}$ would not affect the noise estimate. Not shown in this figure is software injection J0154+4819 for which $\beta_{\max} = 0.61 \pm 0.05$, the reason for this is explained in the text and is shown in figure 4.1.19.

in the noise of this band in the SplInter output[2], and hence a reduction of β. We see that the SplInter algorithm uses a smaller band of data, 0.1 Hz in this case, rather than 0.5 Hz in the heterodyne case, which excludes one of the resonant lines, the power from which bleeds into the heterodyne B_K, but not the SplInter B_k. The other noise peaks are actively suppressed through the outlier removal routines.

We conclude that the SplInter noise estimation routine is accurate in both white noise and noise with spectral features, and that we can see the outlier removal routines actively suppressing noise peaks for signals near to the 'violin' resonance modes of the suspension.

4.1.6 Short-duration hardware injections

We show here the results of a search for the hardware injections – with Doppler and amplitude parameters given respectively by tables 4.1 and 4.2 – using heterodyne and SplInter algorithmsto calculate B_k, and nested sampling for parameter estimation. We use just under four months of data between Oct 19 2009 and Feb 12, 2010 which had

---

[2] This situation also occurred in the search for pulsar J1748−2446ac in [79].

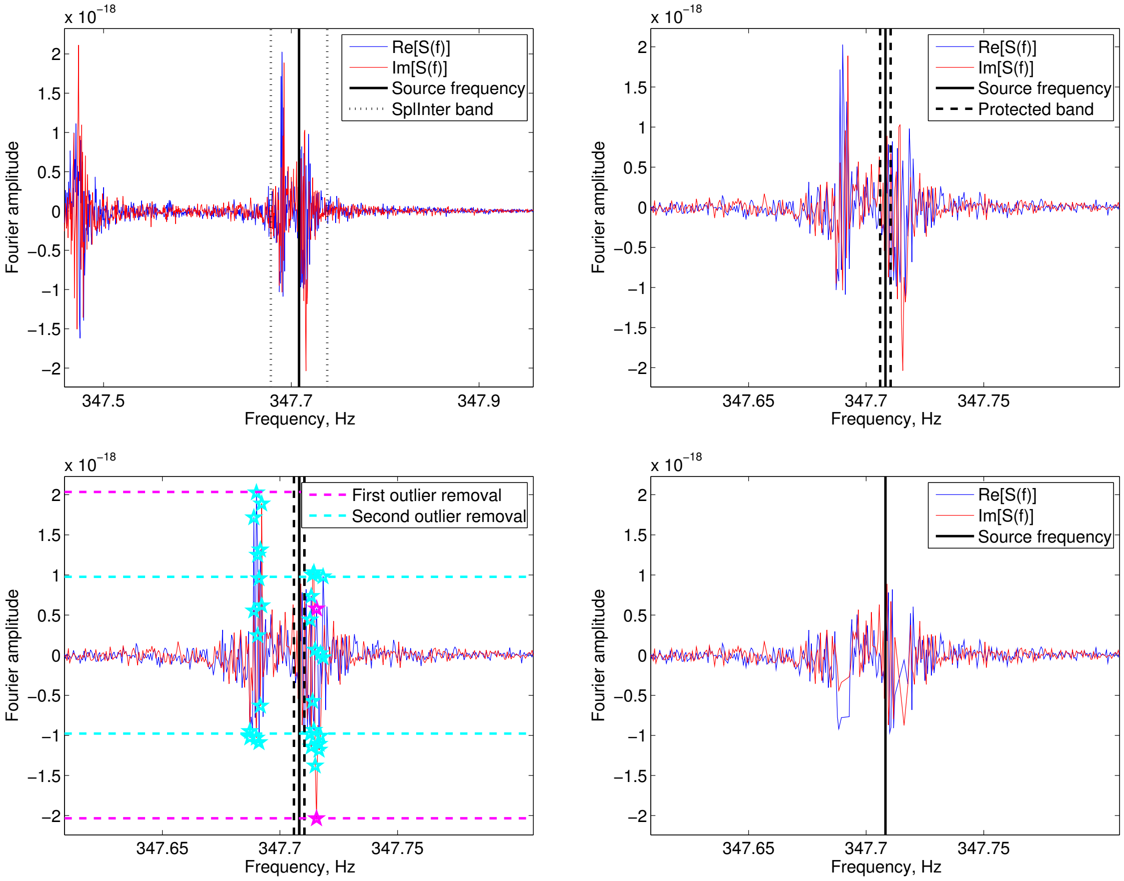

Figure 4.1.19: Noise levels in an example SFT for the 0.5 Hz band around the signal frequency of the outlier software injection J0154+4819 used by the heterodyne routine, compared to 0.1 Hz in SplInter. In the first plot (top left), we see that there are three strong peaks of noise, which correspond to thermally excited violin resonance modes of the detector suspension, and that one of these peaks is ignored by using the smaller bandwidth in the SplInter algorithm. We also see the smaller band used by SplInter, (top right), with an indication of the protected SplInter band. The bottom two plots show the outlier removal process, in which the datapoints removed by each outlier removal routine are highlighted (bottom left). The removal routine by which each point was removed is indicated by the colour. The points within the thresholds which have been removed are ones for which their complex partner is outside of the threshold. The final picture (bottom right) shows the SFT once the outlier removal routines have been performed, we see that the summits of the peaks are removed, leading to a reduction of the noise on the signal, and hence also of its noise estimate and the comparison factor β.

a duty cycle of 47%; this amount of data was used so that the injections would be recovered, but with the posteriors retaining some width. The need for the width in the posteriors is that if the noise calculation is under- or overestimated, it would show in the posterior distribution as a narrowing or widening of the distribution respectively. This will be the test which replicates the end-to-end process of signal recovery most closely, and the one which will indicate with most accuracy whether SplInter is viable as a black box replacement for heterodyne.

Table 4.5 shows the returned SNR values from the nested sampling procedure for the given injections from the SplIntered and heterodyned output with Gaussian likelihood, and the heterodyned output with Student's-t likelihood. The $\Delta t = 1800\,\mathrm{s}$ used in the SplInter routine means that we must use a Gaussian likelihood function as there are not enough B_k values in periods of stationary noise to be able to accurately marginalise over it, and so the heterodyne results are given using both the Gaussian likelihood and the Student's-t distribution likelihoods for fairer comparison.

$B_{k/K}$ algorithm	SplInter	Heterodyne	
Likelihood Distribution	Gaussian		Student's t
PULSAR0	9.217217	10.95006	11.10940
PULSAR1	10.62824	12.63175	12.69116
PULSAR2	80.27499	90.88592	91.04811
PULSAR3	181.3641	205.1734	205.5443
PULSAR4	230.1506	251.8084	253.2708
PULSAR5	10.78574	12.47511	12.33829
PULSAR6	15.83296	17.25627	17.25333
PULSAR7	19.88728	23.38489	23.57995
PULSAR8	197.4373	230.6753	231.9372
PULSAR9	12.57220	13.71983	13.77271

Table 4.5: SNRs of hardware injections in the H1 detector from a four month analysis of S6 data, calculated using the Nested Sampling algorithm `lalapps_pulsar_parameter_estimation_nested`.

We see that the SplInter output is consistent with the heterodyne output, with SNR values generally slightly below those from the heterodyne routine – this is to be expected as the dead time of 1800 s SFTs for this stretch of time is 10.4%, which leads to a drop of 5.3% in SNR. As well as this, the filters applied to the data are different, and there could be some drop in SNR due to spectral leakage from the central peak in the expected sinc function as a result of the almost rectangular windowing. It is also worth considering that in high SNR cases it is more difficult to estimate an underlying noise.

Figures 4.1.20 to 4.1.24 show the posterior distributions of the four parameters which are not fixed in targeted searches, $\mathbf{a} = [h_0, \phi_0, \psi, \cos\iota]$. In all these figures, the blue line shows posteriors made using SplInter for the calculation of B_k, and the red and green lines show use of the heterodyned B_K with Gaussian and Student's t distribution respectively.

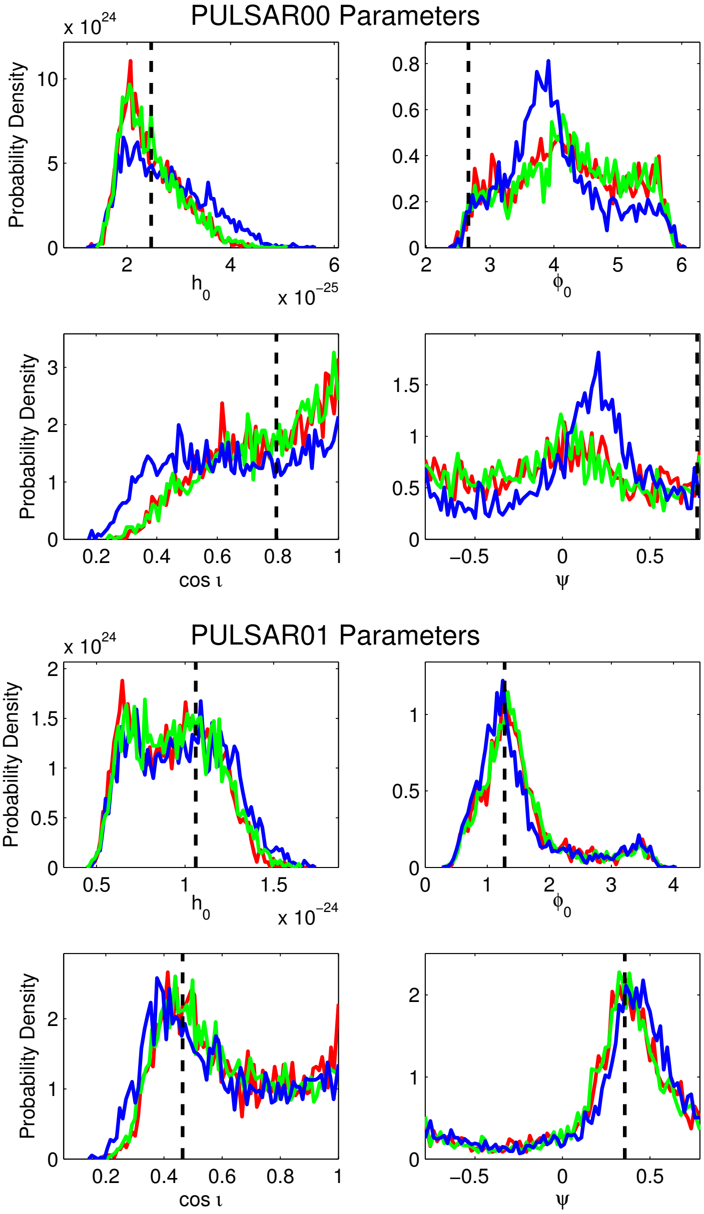

Figure 4.1.20: Posterior distributions for h_0, ϕ_0, $\cos\iota$ and ψ from hardware injections PULSAR0 (top) and PULSAR1 (bottom) using around four months of data from the H1 detector. Posteriors from SplInter with a Gaussian likelihood (blue), heterodyne with a Gaussian likelihood (red) and heterodyne with a Student's t likelihood (green). The vertical black dashed lines show the injected parameters.

The posteriors show that we usually get very good agreement between the experimental outcomes from the heterodyne B_K and SplInter B_k calculation algorithms. The

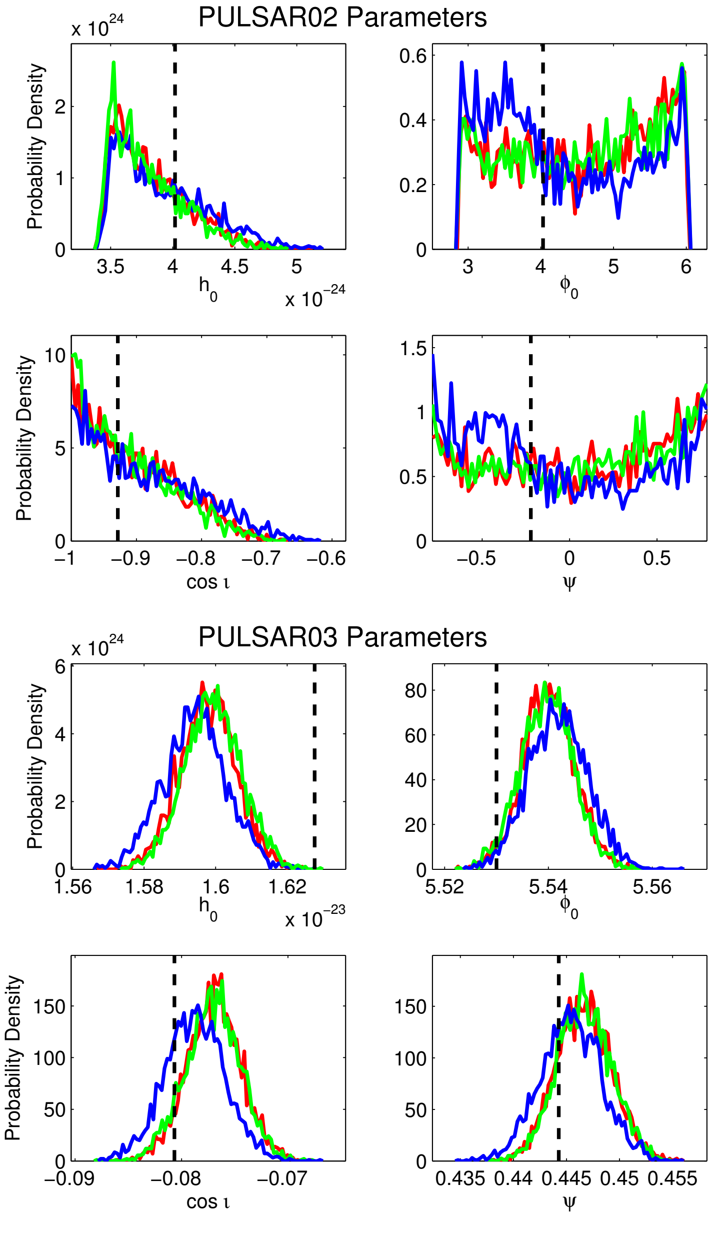

Figure 4.1.21: Posterior distributions for h_0, ϕ_0, $\cos\iota$ and ψ from hardware injections PULSAR2 (top) and PULSAR3 (bottom) using around four months of data from the H1 detector. Posteriors from SplInter with a Gaussian likelihood (blue), heterodyne with a Gaussian likelihood (red) and heterodyne with a Student's t likelihood (green). The vertical black dashed lines show the injected parameters.

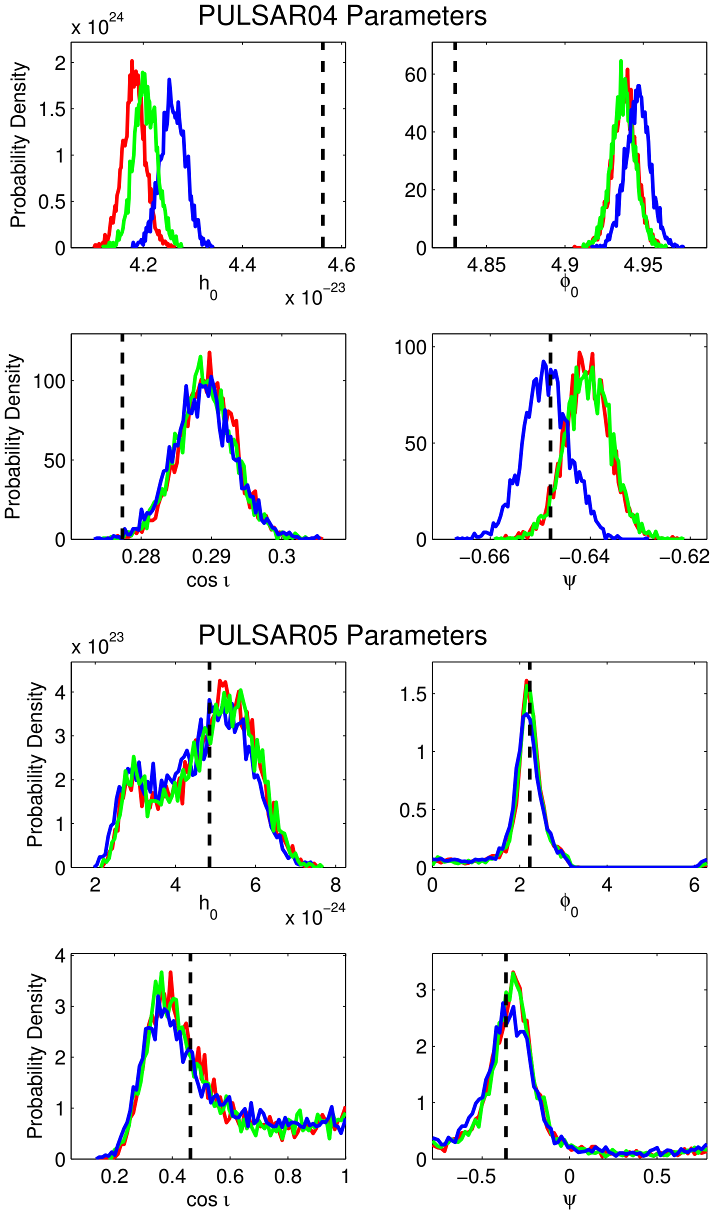

Figure 4.1.22: Posterior distributions for h_0, ϕ_0, $\cos\iota$ and ψ from hardware injections PULSAR4 (top) and PULSAR5 (bottom) using around four months of data from the H1 detector. Posteriors from SplInter with a Gaussian likelihood (blue), heterodyne with a Gaussian likelihood (red) and heterodyne with a Student's t likelihood (green). The vertical black dashed lines show the injected parameters.

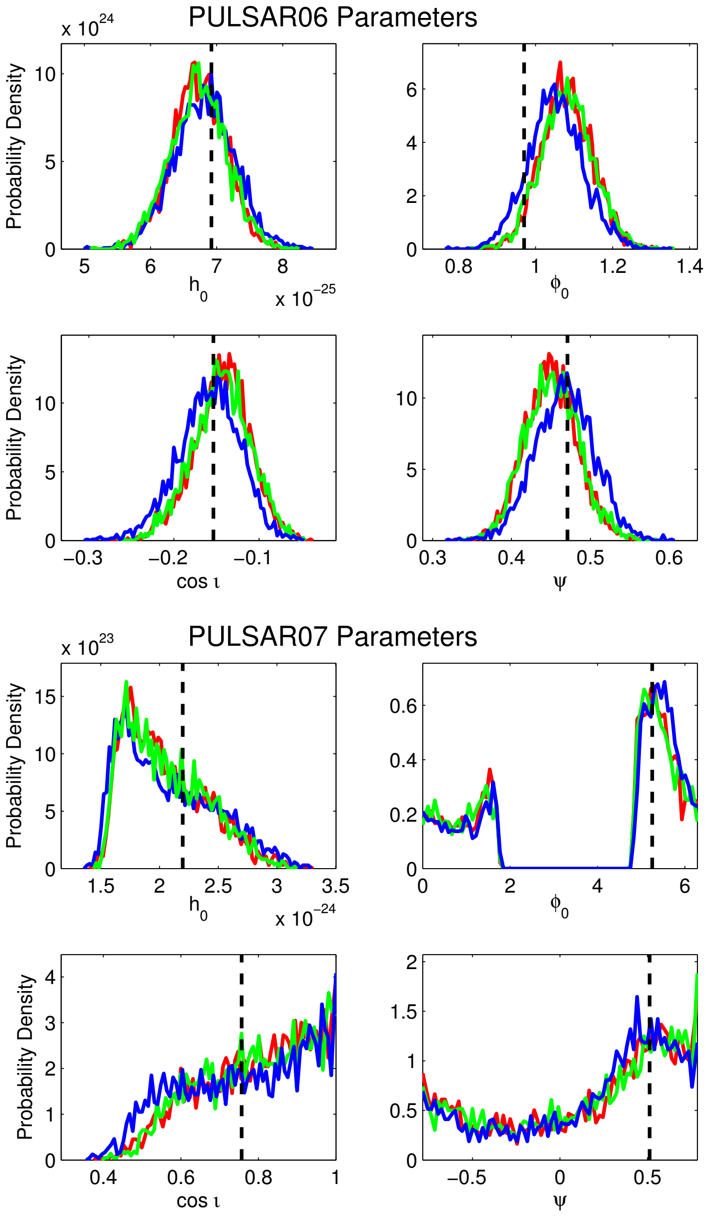

Figure 4.1.23: Posterior distributions for h_0, ϕ_0, $\cos\iota$ and ψ from hardware injections PULSAR6 (top) and PULSAR7 (bottom) using around four months of data from the H1 detector. Posteriors from SplInter with a Gaussian likelihood (blue), heterodyne with a Gaussian likelihood (red) and heterodyne with a Student's t likelihood (green). The vertical black dashed lines show the injected parameters.

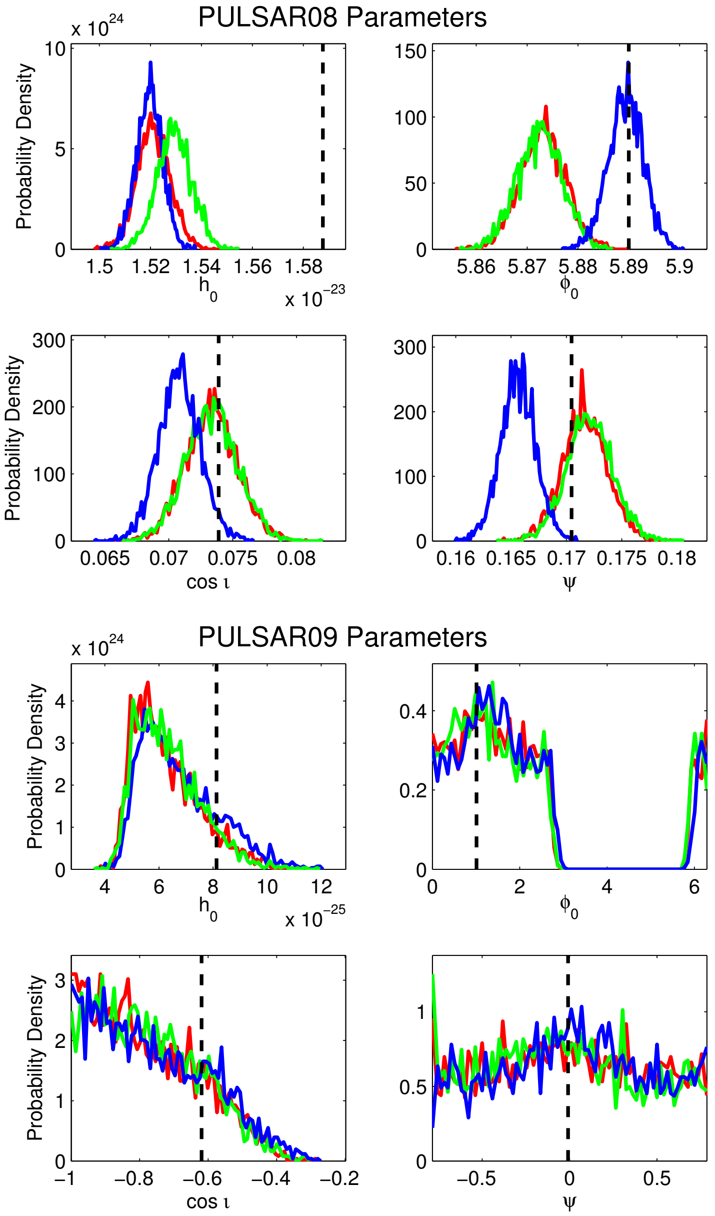

Figure 4.1.24: Posterior distributions for h_0, ϕ_0, $\cos\iota$ and ψ from hardware injections PULSAR8 (top) and PULSAR9 (bottom) using around four months of data from the H1 detector. Posteriors from SplInter with a Gaussian likelihood (blue), heterodyne with a Gaussian likelihood (red) and heterodyne with a Student's t likelihood (green). The vertical black dashed lines show the injected parameters.

first thing to note is that the deviation of parameters from their claimed injection values could be due to different calibrations being used for the injections and the output – it is difficult to know beforehand the calibration that will be used later on in the experiment, but they are similar enough that it is a very small effect on the outcome. There is a small amount of disagreement between the posterior outputs for pulsars 4 and 8, though these are minimal, and insignificant on the scales which affect searches for true signals. We therefore conclude that for all realistic circumstances, the SplInter algorithm is an accurate replacement for the heterodyne routine in these examples.

4.2 Performance testing

The initial aim of developing the SplInter algorithm was to decrease the computational expense of producing the B_K time series, and so we want to be able to quantify how much faster SplInter is compared to the heterodyne algorithm. In doing so, we also profile the times taken in each part of the algorithm, so that we can identify and remedy any bottlenecks.

We also want to know how the algorithm copes with signal parameters which differ from those in the parameter input file, so that we do not have to re-perform the SplInter algorithm for each datapoint in $[f^{(0)}, f^{(1)}]$ space if we wish to search in frequency space, or in another parameter space which will affect frequency, such as sky position.

4.2.1 Speed testing

The main aim of the SplInter algorithm is to greatly reduce computation times, so we want to know the timing statistics. The SplInter algorithm consists of three parts which we will time, as well as the total time taken:

- Within each segment, we will take the SFT *catalogue load* time, during which we use the `XLALSFTdataFind` algorithm to search for which SFTs in the provided list or cache file are within the time constraints given – i.e. the start and end time of the segment.

- The *SFT load* time is the time taken to load the $S_k(f)$ data from the SFT file.

- The *interpolation time* is the amount of time taken to calculate B_k from the SFT data, including obtaining details of barycentric time corrections, and outputting this to a data file buffer.

Table 4.6 shows the mean and median times taken in each part of the algorithm for SplInter and heterodyne. As the heterodyne algorithm analyses the sources one at a time it is fairest to compare the total algorithm time taken per source, and as the SplInter algorithm time per source is not constant for number of sources, we will compare the total time per source for one, ten, one hundred and one thousand sources at a time. The profiling times are not relevant (or not calculated) for the heterodyne routine, and so are not included. The timings calculated use data from a single interferometer, and the time taken will increase linearly with the number of interferometers used, as the algorithms are performed for one interferometer at a time.

	Heterodyne	SplInter			
Sources	1	1	10	100	1000
Total Time, s					
Median	497.4	0.9476	1.6509	2.4710	10.6050
Mean	505.1	0.9619	1.6818	2.5247	10.8586
CPUh/N/h	5.76×10^{-3}	1.10×10^{-5}	1.91×10^{-6}	2.86×10^{-7}	1.23×10^{-7}
Source parameter load time, s					
Median	–	0.0046	0.0447	0.4544	4.6518
Mean	–	0.0050	0.0471	0.4798	4.8651
Catalogue load time, s					
Median	–	0.0823	0.0806	0.0808	0.0852
Mean	–	0.0897	0.0846	0.0900	0.0909
SFT load time, s					
Median	–	0.0774	0.6568	0.6830	0.7325
Mean	–	0.0807	0.6743	0.6921	0.7516
Interpolation time, s					
Median	–	0.0013	0.0110	0.1040	1.0442
Mean	–	0.0013	0.0110	0.1041	1.0445
Other, s					
Median	–	0.7791	0.8544	1.1258	3.9462
Mean	–	0.7852	0.8649	1.1587	4.1065

Table 4.6: Median and mean times in seconds taken to analyse a day of data using heterodyne and SplInter for one, ten, one hundred and one thousand sources. Included is the CPU core hours per source per hour of data, given the median computation time, followed by the amount of time spent in each part of the algorithm; source parameter load time, SFT catalogue load time, SFT load time and interpolation time and other, which includes initialisation of the barycentring routine, loading ephemeris files, segment lists and other calculations.

We also compare the time taken to load the source parameters from the parameter file, so that we fill in as many of the blanks on the time taken to perform the algorithm as possible, this is not included in figure 4.2.1 but it is intuitively proportional to number of sources.

The times used here are from analyses with the sinc interpolation scheme as this is the method which will be used most often in real analyses. Timings for the mixed or Fresnel interpolation schemes will be similar, as we can see that the interpolation itself is a relatively small proportion of the overall time.

We then take the average time taken per day of data for each of the three timings taken, which is shown in figure 4.2.1. Repeated timing values are taken to give their distribution.

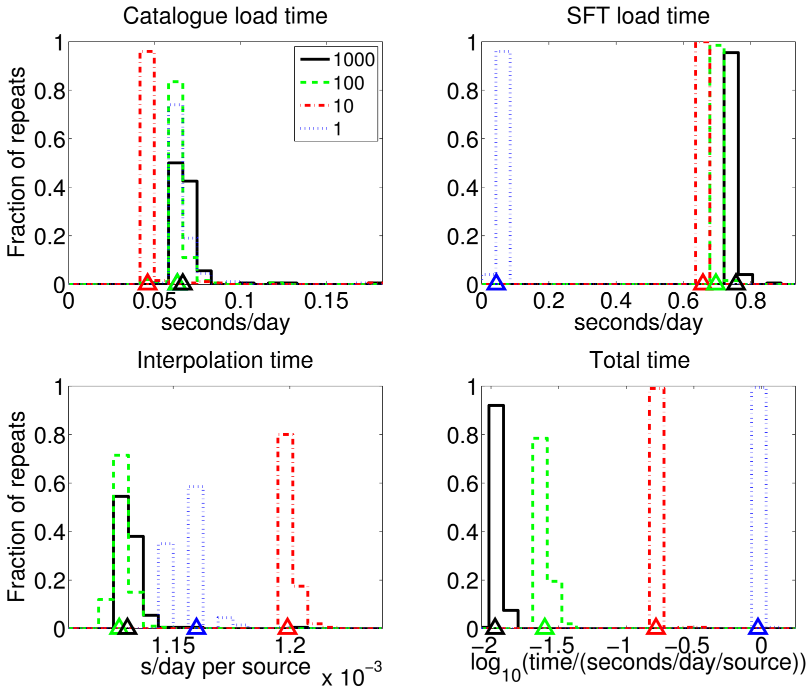

Figure 4.2.1: Histograms of the average time taken to analyse a day of data for various parts of the SplInter algorithm. Times are all in seconds per SFT, except interpolation time and total time which are seconds per SFT per source. The different coloured histograms show the different number of sources used in each analysis. The horizontal axis markers denote the mean values for these times for the different numbers of sources.

We see that the catalogue load time and interpolation time per source are generally independent of the number of sources, though the interpolation time per source has a slight improvement for higher numbers of sources; this is due to the program saving some of the items in RAM for faster access at a later time. The SFT load time is reduced when using a single source by only loading a small ($\sim 3\,\mathrm{Hz}$) band around the approximated signal frequency, but once we are using more than one source, the extra time taken is small enough that it is worthwhile to load a much wider SFT band, we load a band from 1.5 Hz below the lowest source frequency to 1.5 Hz above the highest source frequency.

The analyses for table 4.6 and figure 4.2.1 were run on the `atlas` computer cluster hosted in the Albert Einstein Institute, Hannover.

Use of single or multiple SFT cache files

During this work, it was found that for analyses with long duration and many segments, the SFT catalogue load time was causing a major bottleneck in computational efficiency, with a large amount of time being taken on this part of the algorithm. This was found to be due to using a single SFT cache file, containing paths for all of the SFTs used throughout the duration of the analysis. The problem came from the way the `XLALSFTdataFind` algorithm (for finding SFT data) works, opening the header of each SFT in the cache file to check its timestamp details every time the function is called. In order to work around this problem, we made it so that SplInter could use input of a directory containing an SFT cache for each segment. The mean number of SFTs in each segment's cache file for H1 during S6 was 5.6 with a median of 3 SFTs, and maximum of 63 SFTs, compared to a total of nearly 2500 in S6. A histogram of the number of SFTs contained in each segment is given in figure 4.2.2.

The catalogue load times in table 4.6 and figure 4.2.1 are for a segment of 24 h, (i.e. containing 48 SFTs), so this time is actually toward the higher end of the scale for real analyses if we utilise the one-cache-per-segment configuration.

Timing tests were performed on the stretch of data used in the hardware injection tests in section 4.1.6. These tests were to compare using the single cache input, containing all SFTs found within the time range of the analysis, and the one SFT cache per segment input. This comparison is presented in figure 4.2.3, which shows the time improvement using the multiple SFT cache files rather than the single file.

We see the vast improvement in catalogue load time by using multiple SFT cache files, with the mean catalogue load time per SFT decreasing by two orders of magnitude, from 0.2 s to 0.002 s. This is a major proportion of the overall time taken in long studies.

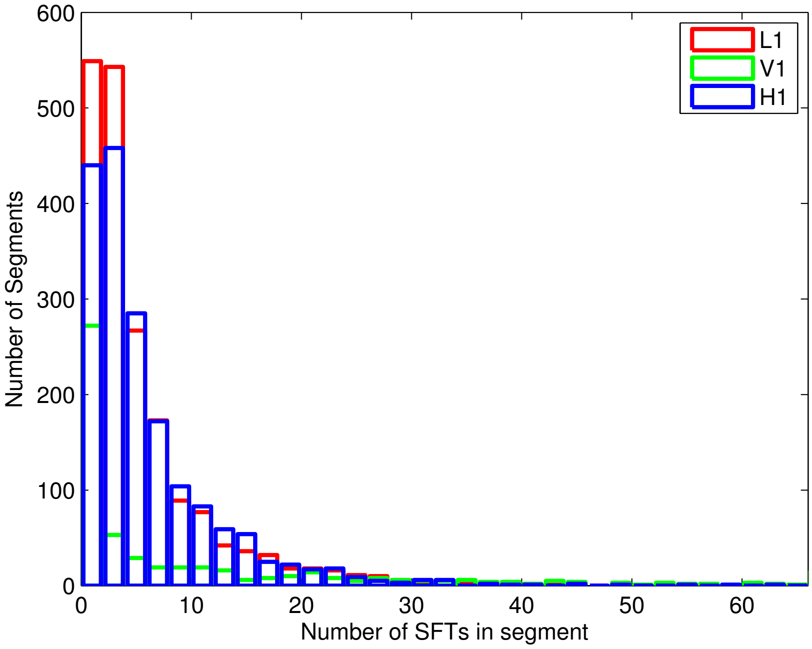

Figure 4.2.2: Histograms of the number of SFTs in each segment for H1 (blue), L1 (red) and V1 (green).

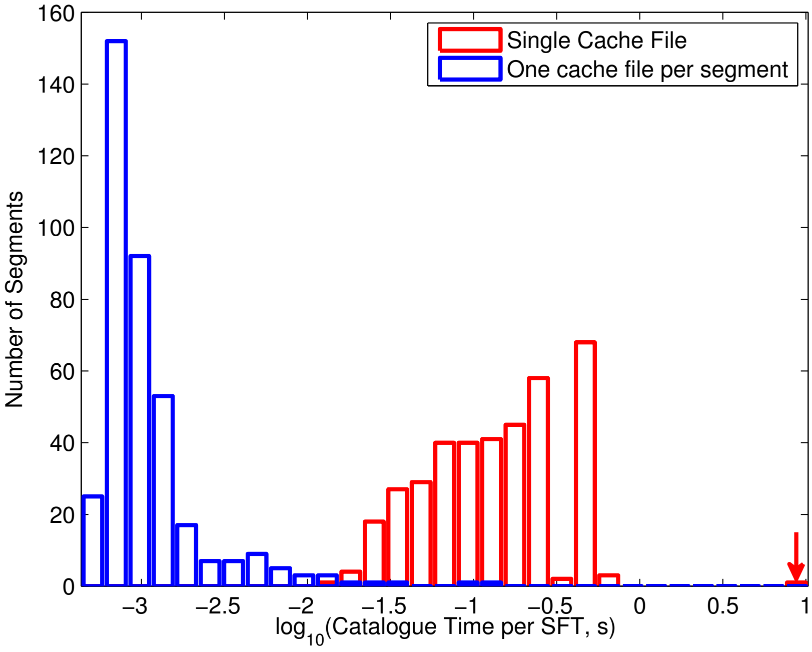

Figure 4.2.3: Histogram of the time taken to load the SFT catalogue using single or multiple cache files. The single large outlier is indicated by the red arrow in the bottom-right hand corner of the plot.

We also see the large outlier in the single cache file histogram, indicated by the arrow, with an SFT catalogue load time of 8.75 s per SFT. This outlier was the first segment in the analysis, and we note that this is not unusual, there is often a very large SFT catalogue load time in the first segment for a large list of SFTs, this is again from the program saving some variables in RAM for faster access at a later point.

4.2.2 Frequency response testing

When performing the Bayesian parameter estimation using the B_k and σ_k values, we may wish to search in frequency space, meaning that the Spectral Interpolation or heterodyne procedure has been performed with incorrect parameters for that parameter space point. So that we can avoid rerunning the SplInter algorithm for each point in the parameter space, we can allow for this in the parameter estimation stage by considering how the algorithms will react to an incorrect $f^{(0)}$, with knowledge of how the other parameters will affect the frequency evolution of the signal. The test in this case is to inject signals into a noiseless frame and SFT, and to use the algorithm to try to recover this signal from a nearby frequency. This test is performed for a source with constant antenna pattern and has had any orbital motion effects removed, so that it is easier to control the output values in comparison to the input.

Figure 4.2.4 shows the result of the frequency response test, and we see that the outcome is very similar to the sinc function, which we can therefore use as the correction factor in frequency space searches, by multiplying by an appropriate amplitude correction factor of $\frac{1}{\mathrm{sinc}(\Delta f)}$. The mismatch is not used in this case as it is dominated by the contributions from the incorrect phase, which can be corrected with a complex multiplication to unwind the phase difference caused by the frequency difference. The frequency resolution of a search will be an extremely small fraction of this central peak in a targeted search, for which the frequency resolution will be up to $\sim 10^{-10}$ Hz. Candidate followup searches for example as in section 6.4, can have frequency resolution of up to 5×10^{-5} Hz, which is approximately 10% of the width of a frequency bin for a 1800 s SFT. The frequency response will change according to the SFT length, and is best represented in terms of the number of frequency bins from the signal frequency. The heterodyne frequency response is also determined by the Δt used, and its effective bin width is 1/60 Hz for standard searches.

For frequency differences which are a more than around half a bin width, the SplInter analysis would be reperformed, as the loss in amplitude would be too significant. As the absolute value of B_k is not expected to change much in these searches, we do not update σ_k for each parameter space point.

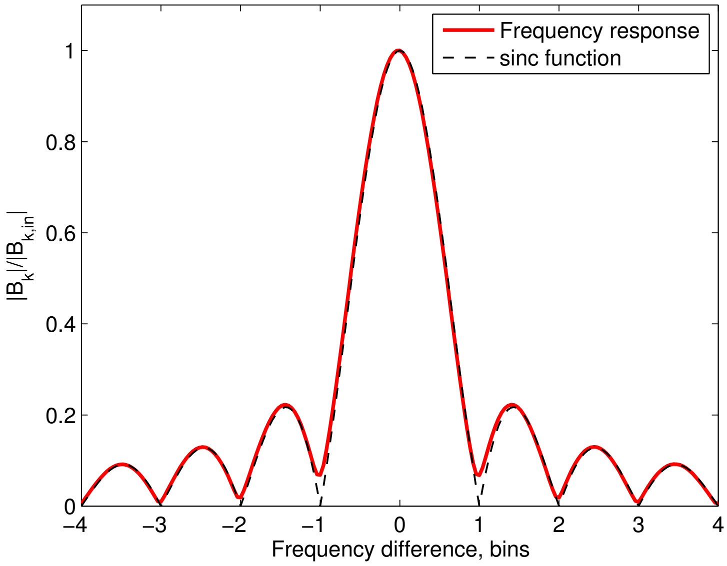

Figure 4.2.4: Loss of amplitude for an incorrect interpolation frequency over the course of a few bins either side of the signal frequency. We see that this closely follows the sinc function, and so we use $\mathrm{sinc}(\Delta f)$ as the correction factor for the amplitude in a frequency search.

Chapter 5

Targeted searches for known pulsars using the spectral interpolation algorithm

> [Hitting the target is] *like trying to hit a puppy by throwing a live bee at it. Which is a weird image and you should all just forget it.*
>
> Willow, Buffy the Vampire Slayer, 2001

The Bayesian analysis pipeline as described in section 2.3 has so far been used exclusively for targeted searches of continuous waves from neutron stars, and almost exclusively for continuous waves with a gravitational-wave frequency at twice the rotational frequency, and so we continue in that tradition in this chapter using the new analysis algorithm.

Section 5.1 presents the results of targeted searches for signals from isolated and binary pulsars. These analyses use the entirety of 'science mode' data from LIGO science run 6 (S6) from the Hanford (H1) and Livingston (L1) LIGO detectors, and Virgo science run 2 (VSR2) of the Virgo detector, (V1). The S6 analysis runs from Jul 08 2009 to Oct 20 2010, and VSR2 is from Jul 08 2009 UTC to Oct 22 2009.

We set the priors on the amplitude parameters to be uniform within a certain range, the reasons for using these priors were discussed in section 2.2. The priors are set as in table 5.1, we conservatively choose an upper limit on h_0 of 10^{-22} as this is an order of magnitude larger than the highest upper limits set by the Einstein@Home all-sky search for S5 data (figure 9 of [3]). The prior on ψ is in the range $[0, \pi/2]$, which is entirely equivalent to the range $[-\pi/4, \pi/4]$ stated previously in section 2.2, as the signal cycles in ψ over a period of $\pi/2$ [58].

parameter	lower limit	upper limit
h_0	0	10^{-22}
ψ	0	$\frac{\pi}{2}$
ϕ_0	0	2π
$\cos\iota$	-1	1

Table 5.1: Uniform prior limits in the targeted search for continuous-wave signals from pulsars in S6.

As discussed previously (also in section 2.2), we shall present the results in the form of 95% upper limits, defined by

$$\int_{h_0=0}^{h_0^{95\%}} \mathrm{p}(h_0 \mid \{B_k\})\,\mathrm{d}h_0 = 0.95, \tag{5.0.1}$$

which means that there is a 95% probability that the value of h_0 is below this value from the data we have. This statement implies that there is a 5% probability that h_0 is higher than this value, but we remember that we are not applying astrophysical priors, and as such we include a comparison to the spin-down limit, h_0^{sd} and show the ratio between the two. This is effectively stating how close we are to be able to make definitive, new, claims about an astrophysical object.

In the following discussion and the rest of this chapter, we use ν and $\dot{\nu}$ to denote rotational frequency (previously $\nu^{(0)}$) and its first derivative ($\nu^{(1)}$) respectively for ease of reading.

The spin-down limit is calculated by equating the loss of kinetic energy seen in the slowing down of the rotation to the gravitational-wave energy emitted by the source, i.e. what would be the gravitational-wave energy if GW emission was the only mechanism responsible for the energy loss causing the slowing down of the rotation. This luminosity

$\dot{E}_{\rm gw}$ is

$$\begin{aligned}\dot{E}_{\rm gw} =& \frac{2048\pi^6}{5}\frac{G}{c^5}\nu^6 I_{zz}^2\varepsilon^2 \\ =& \frac{8\pi^2}{5}\frac{c^3}{G}\nu^2 h_0^2 d^2,\end{aligned} \tag{5.0.2}$$

and the observed spin-down energy – the rotational kinteic energy lost – is

$$\dot{E} = 4\pi^2 I_{zz}\nu\dot{\nu}, \tag{5.0.3}$$

where, as before, I_{zz} is the moment of inertia about the rotational axis and ε is ellipticity, as defined in equation 2.1.4. We therefore have the *canonical* spin-down limit

$$\begin{aligned}h_0^{\rm sd} &= \left(\frac{5}{2}\frac{G}{c^3}\frac{I_{zz}}{d^2}\frac{|\dot{\nu}|}{\nu}\right)^{1/2} \\ &= 8.06\times 10^{-19}\frac{I_{38}^{1/2}}{d_{\rm kpc}}\left(\frac{|\dot{\nu}|}{\nu}\right)^{1/2},\end{aligned} \tag{5.0.4}$$

where I_{38} is the star's principal moment of inertia in units of $10^{38}\,{\rm kg\,m^2}$ and $d_{\rm kpc}$ is the distance to the star in kiloparsecs. In the calculations presented here, we use $I_{38} = 1$, the standard moment of inertia used in most examples of the literature.

In nature, in the vast majority of cases, we see that neutron stars are spinning down, that their rotational frequency is decreasing. However in a few cases we see from electromagnetic observations that the rotational frequency is increasing; in some cases this can be explained through accretion of material onto the neutron star from its companion in a binary system, however for an isolated source, this is not possible. Millisecond pulsars will have been spun-up to high rotation frequencies through this accretion, though this process will have ended if the source is now isolated.

So why is $\dot{\nu}$ positive? The implication from equating gravitational-wave luminosity to the increase in kinetic energy is that this source would somehow be absorbing gravitational waves and using this gained energy to spin faster[1]. The answer is that we do not see positive spin-downs in nature *in the frame of reference of the neutron star*, the positive spin-downs we see are due to the proper motion between the source and the

---

[1] Logically this would be possible, but would require a rotational phase evolution perfectly, negatively, matched to the gravitational-wave phase evolution from another source.

detector, specifically, if the source is accelerating towards us, we will see an apparent spin-up due to the increasing Doppler shifting.

For sources which have significant line-of-sight proper motion, we take this into account, in the following results tables these spin-down limits are marked with a †. For sources which are known to have significant proper motion, i.e. they are in a globular cluster, but the true proper motion is not known, a characteristic 'spin-down age' of $\tau = 10^9$ years is assumed to calculate the used $\dot{\nu}$ [2].

The characteristic age, τ, is the period related to how much rotational energy is in the system, and the rate at which it is being lost. It is defined by

$$\tau = \frac{1}{n-1}\frac{\nu}{\dot{\nu}}, \tag{5.0.5}$$

where n is the braking index, which for a magnetic dipole $n = 3$ and for purely gravitational-wave quadrupole spin-down would be $n = 5$ [2]. In this calculation, $n = 3$ is used, but as we are only really concerned with an order-of-magnitude estimate, the factor of 5/3 is unimportant. The spin-down limits for which we use a characteristic age are indicated by ‡ in the results tables.

From the $h_0^{95\%}$ results of an analysis, we can also constrain the ellipticity ε to an upper limit containing 95% of the probability, $\varepsilon^{95\%}$. To do this, we rearrange equation 2.1.3 to obtain

$$\varepsilon^{95\%} = 0.237\frac{h_0^{95\%}}{10^{-24}}\frac{d_{\mathrm{kpc}}}{I_{38}\nu^2}. \tag{5.0.6}$$

Again, we use the canonical moment of inertia of $I_{zz} = 10^{38}\,\mathrm{kg\,m^2}$, so that $I_{38} = 1$.

5.1 Results from targeted analysis

We present the upper limit results of two searches for known pulsars, with parameters as given in [26, 67]. The difference between the analyses is only in the algorithm utilised, so we expect similar results from the analyses. One of the searches has already been performed by others, utilising the heterodyne routine in section 2.3.1 and an MCMC as in 2.2.1, this work was presented in [2]. The reason for presenting this work is that we perform a similar analysis, but this time utilising the SplInter and nested sampling algorithms, meaning that we want to compare the two routines. We do so with a direct comparison of the 95% upper limits.

The heterodyne/MCMC results for $h_0^{95\%}$ are given in table 7 of [2], though as we use an updated version of the ATNF pulsar database [26, 67] (v1.53 instead of v1.47) for the spin-down limits, we update their comparisons to the spin-down limit $h_0^{95\%}/h_0^{\rm sd}$, as distance estimates can vary by a large amount and these have a strong effect on the spin-down limit value. We do not include results for the heterodyne routine ellipticity for the same reason.

Although this analysis is similar to the one used to produce table 7 in [2], we do not expect the 95% upper limits to be exactly the same, for a few reasons which we shall now discuss.

The 95% upper limits are going to be intrinsically noisy; we are going to be in the tails of the noise distribution, the region containing 95% of the probability will be outside of 2σ for a one-sided normal distribution, which means that the histogram procedure when counting the posterior samples from nested sampling will be counting low numbers, leading to high noise from the counting statistics of around 1 in $N^{1/2}$. Empirically this can be seen to affect the upper limits by as much as 10%, depending upon the exact shape of the tails of the posterior.

We also expect the SplInter upper limits to differ from those given by the heterodyne routine for reasons stated before (e.g. in section 4.1.6) for differences between the two outputs, such as the dead time from only using 1800 s segments of data, and the use of a Gaussian likelihood in the nested sampling compared to the Student's-t likelihood used in the MCMC. As well as this the priors on h_0 and $\cos\iota$ used for the SplInter analysis are uniform rather than utilising the S5 results as in [2]. Not applying these priors could lead to a higher $h_0^{95\%}$ especially in cases where the source is close in frequency to a noisy line which is present in S6 but not S5.

Table 5.2 shows the result of targeted searches for continuous gravitational waves at twice the rotation frequency of known isolated pulsars, and table 5.3 shows the result of targeted searches for known binary pulsars. The analyses have been separated into isolated and binary pulsars, because the results from section 4.1.2 need to be taken into account, and we do so on a case-by-case basis for the pulsars in binary systems presented in table 5.3.

				SplInter & Nested Sampling			heterodyne & MCMC	
Pulsar	ν (Hz)	$\dot{\nu}$ (Hz/s)	$h_0^{\rm sd}$	$h_0^{95\%}$	$\varepsilon^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$	$h_0^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$
J0024−7204C	173.71	1.6×10^{-15}	$6.1\times10^{-28}\dagger$	1.6×10^{-25}	5.0×10^{-6}	260	1.7×10^{-25}	280
J0024−7204D	186.65	1.6×10^{-16}	$1.8\times10^{-28}\dagger$	4.1×10^{-26}	1.1×10^{-6}	220	2.9×10^{-26}	160
J0024−7204F	381.16	-9.3×10^{-15}	$9.9\times10^{-28}\dagger$	7.3×10^{-26}	4.7×10^{-7}	73	6.7×10^{-26}	68
J0024−7204G	247.50	2.6×10^{-15}	$6.6\times10^{-28}\dagger$	6.3×10^{-26}	9.7×10^{-7}	95	8.2×10^{-26}	120
J0024−7204L	230.09	-3.6×10^{-15}	$8.0\times10^{-28}\ddagger$	4.5×10^{-26}	8.1×10^{-7}	56	4.2×10^{-26}	53
J0024−7204M	271.99	-4.3×10^{-15}	$8.0\times10^{-28}\ddagger$	8.2×10^{-26}	1.0×10^{-6}	100	7.0×10^{-26}	88
J0024−7204N	327.44	2.5×10^{-15}	$5.6\times10^{-28}\dagger$	6.5×10^{-26}	5.7×10^{-7}	120	5.1×10^{-26}	91
J0030+0451	205.53	-4.3×10^{-16}	$4.1\times10^{-27}\dagger$	7.8×10^{-26}	1.2×10^{-7}	19	7.2×10^{-26}	18
J0340+41	303.09	-6.5×10^{-16}	4.4×10^{-28}	6.4×10^{-26}	4.4×10^{-7}	150	5.6×10^{-26}	130
J0711−6830	182.12	-2.9×10^{-16}	$9.8\times10^{-28}\dagger$	4.4×10^{-26}	3.2×10^{-7}	45	3.4×10^{-26}	35
J1024−0719	193.72	1.3×10^{-16}	$1.3\times10^{-27}\dagger$	5.3×10^{-26}	1.6×10^{-7}	39	4.6×10^{-26}	35
J1038+0032	34.66	-7.8×10^{-17}	5.1×10^{-28}	1.2×10^{-25}	5.5×10^{-5}	230	1.2×10^{-25}	240
J1453+1902	172.64	-3.2×10^{-16}	$1.2\times10^{-27}\dagger$	1.4×10^{-25}	1.1×10^{-6}	120	1.4×10^{-25}	120
J1518+0204A	180.06	-2.9×10^{-15}	$4.0\times10^{-28}\ddagger$	8.6×10^{-26}	5.0×10^{-6}	210	7.8×10^{-26}	200
J1641+3627A	96.36	-1.5×10^{-15}	$4.9\times10^{-28}\ddagger$	5.8×10^{-26}	9.7×10^{-6}	120	5.0×10^{-26}	100
J1721−2457	285.99	-2.4×10^{-16}	$4.7\times10^{-28}\dagger$	6.2×10^{-26}	2.8×10^{-7}	130	5.5×10^{-26}	120
J1730−2304	123.11	-3.1×10^{-16}	2.5×10^{-27}	3.5×10^{-26}	2.8×10^{-7}	14	4.5×10^{-26}	18
J1744−1134	245.43	-4.3×10^{-16}	$2.5\times10^{-27}\dagger$	4.9×10^{-26}	8.1×10^{-8}	19	6.3×10^{-26}	25
J1748−2446C	118.54	-1.9×10^{-15}	$5.8\times10^{-28}\ddagger$	4.0×10^{-26}	3.7×10^{-6}	68	3.8×10^{-26}	66
J1748−2446D	212.14	-3.4×10^{-15}	$5.8\times10^{-28}\ddagger$	5.4×10^{-26}	1.6×10^{-6}	93	5.3×10^{-26}	91
J1748−2446F	180.50	-2.9×10^{-15}	$5.8\times10^{-28}\ddagger$	7.0×10^{-26}	2.8×10^{-6}	120	7.4×10^{-26}	130
J1748−2446G	46.14	-7.3×10^{-16}	$5.8\times10^{-28}\ddagger$	5.8×10^{-26}	3.5×10^{-5}	99	4.5×10^{-26}	76
J1748−2446H	203.01	-3.2×10^{-15}	$5.8\times10^{-28}\ddagger$	5.8×10^{-26}	1.8×10^{-6}	99	5.8×10^{-26}	100
J1748−2446K	336.74	-5.3×10^{-15}	$5.8\times10^{-28}\ddagger$	7.1×10^{-26}	8.1×10^{-7}	120	5.9×10^{-26}	100
J1748−2446L	445.49	-7.1×10^{-15}	$5.8\times10^{-28}\ddagger$	1.4×10^{-25}	9.0×10^{-7}	240	1.1×10^{-25}	190
J1748−2446R	198.86	-3.2×10^{-15}	$5.8\times10^{-28}\ddagger$	4.7×10^{-26}	1.6×10^{-6}	81	5.1×10^{-26}	87
J1748−2446S	163.49	-2.6×10^{-15}	$5.8\times10^{-28}\ddagger$	7.2×10^{-26}	3.5×10^{-6}	120	5.3×10^{-26}	87
J1748−2446T	141.15	-2.2×10^{-15}	$5.8\times10^{-28}\ddagger$	3.9×10^{-26}	2.5×10^{-6}	67	3.0×10^{-26}	52
J1748−2446aa	172.77	-2.7×10^{-15}	$5.8\times10^{-28}\ddagger$	1.5×10^{-25}	6.6×10^{-6}	260	1.5×10^{-25}	260
J1748−2446ab	195.32	-3.1×10^{-15}	$5.8\times10^{-28}\ddagger$	7.7×10^{-26}	2.6×10^{-6}	130	4.0×10^{-26}	69
J1748−2446ac	196.58	-3.1×10^{-15}	$5.8\times10^{-28}\ddagger$	6.7×10^{-26}	2.3×10^{-6}	110	5.6×10^{-26}	97
J1748−2446af	302.63	-4.8×10^{-15}	$5.8\times10^{-28}\ddagger$	7.3×10^{-26}	1.0×10^{-6}	120	6.0×10^{-26}	100
J1748−2446ag	224.82	-3.6×10^{-15}	$5.8\times10^{-28}\ddagger$	4.5×10^{-26}	1.2×10^{-6}	77	5.1×10^{-26}	88
J1748−2446ah	201.40	-3.2×10^{-15}	$5.8\times10^{-28}\ddagger$	4.5×10^{-26}	1.4×10^{-6}	76	4.0×10^{-26}	69
J1801−1417	275.85	3.1×10^{-16}	$4.8\times10^{-28}\dagger$	8.7×10^{-26}	4.9×10^{-7}	180	7.3×10^{-26}	150
J1803−30	140.82	-2.2×10^{-15}	$4.1\times10^{-28}\ddagger$	5.2×10^{-26}	4.8×10^{-6}	130	4.1×10^{-26}	100
J1823−3021A	183.82	-2.9×10^{-15}	$2.7\times10^{-28}\ddagger$	6.7×10^{-26}	5.6×10^{-6}	250	3.5×10^{-26}	130
J1824−2452A	327.41	-1.7×10^{-13}	$3.4\times10^{-27}\dagger$	8.6×10^{-26}	1.0×10^{-6}	25	5.5×10^{-26}	16
J1843−1113	541.81	-2.8×10^{-15}	$9.3\times10^{-28}\dagger$	1.6×10^{-25}	2.5×10^{-7}	170	1.1×10^{-25}	120
J1905+0400	264.24	-2.8×10^{-16}	$6.3\times10^{-28}\dagger$	6.2×10^{-26}	2.8×10^{-7}	99	5.0×10^{-26}	79
J1910−5959B	119.65	-1.9×10^{-15}	$7.1\times10^{-28}\ddagger$	3.6×10^{-26}	2.7×10^{-6}	50	2.5×10^{-26}	35
J1910−5959C	189.49	1.1×10^{-18}	$1.4\times10^{-29}\dagger$	4.3×10^{-26}	1.3×10^{-6}	3200	3.2×10^{-26}	2300
J1910−5959D	110.68	-1.8×10^{-15}	$7.1\times10^{-28}\ddagger$	3.1×10^{-26}	2.7×10^{-6}	43	2.1×10^{-26}	30

Continued

				SplInter & Nested Sampling			heterodyne & MCMC	
Pulsar	ν (Hz)	$\dot{\nu}$ (Hz/s)	$h_0^{\rm sd}$	$h_0^{95\%}$	$\varepsilon^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$	$h_0^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$
J1910−5959E	218.73	-3.5×10^{-15}	7.1×10^{-28}‡	4.2×10^{-26}	9.3×10^{-7}	59	3.6×10^{-26}	51
J1911+1347	216.17	-7.9×10^{-16}	9.6×10^{-28}	5.9×10^{-26}	4.8×10^{-7}	61	4.8×10^{-26}	50
J1913+1011	27.85	-2.6×10^{-12}	5.5×10^{-26}	4.3×10^{-25}	5.8×10^{-4}	7.7	1.6×10^{-25}	2.9
J1939+2134	641.93	-4.3×10^{-14}	1.3×10^{-27}†	1.3×10^{-25}	3.6×10^{-7}	95	1.3×10^{-25}	100
J1944+0907	192.86	-3.6×10^{-16}	8.6×10^{-28}†	5.5×10^{-26}	4.5×10^{-7}	64	5.5×10^{-26}	64
J2007+2722	40.82	-1.6×10^{-15}	7.4×10^{-28}	7.6×10^{-26}	7.4×10^{-5}	100	7.1×10^{-26}	96
J2010−1323	191.45	-1.8×10^{-16}	6.0×10^{-28}	6.9×10^{-26}	5.7×10^{-7}	110	6.3×10^{-26}	110
J2124−3358	202.79	-4.4×10^{-16}	4.0×10^{-27}†	4.8×10^{-26}	8.4×10^{-8}	12	3.9×10^{-26}	9.8
J2322+2057	207.97	-1.8×10^{-16}	9.6×10^{-28}†	4.1×10^{-26}	1.7×10^{-7}	43	5.4×10^{-26}	56

Table 5.2: Upper limits on the gravitational-wave amplitude of continuous waves from known isolated pulsars.
A † denotes that the pulsar's spin-down is corrected for proper motion effects.
A ‡ denotes that the pulsar's spin-down is calculated using a characteristic spin-down age of 10^9 years.

We see that the upper limits from the SplInter/nested sampling analysis are generally slightly higher than those from the heterodyne/MCMC analysis. This is to be expected from the drop in SNR (leading to an increase in upper limit) from the dead time consideration. Some pulsars have an improved upper limit in the new analysis, which can be due to the line removal scheme, such as pulsar J0024−7204C, for which the surrounding noise in an example SFT is plotted in figure 5.1.1, the peaks seen will be removed by the SplInter routine, but will adversely affect the upper limit produced by the heterodyned analysis.

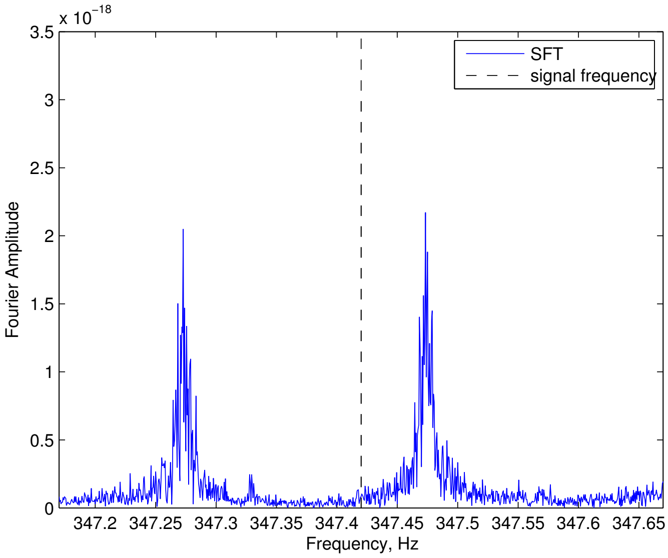

Figure 5.1.1: Noise levels around the gravitational-wave frequency for J0024−7204C, these peaks will be removed by the SplInter routine, but will adversely affect the upper limit produced by the heterodyned analysis.

For an overall comparison, we use the ratio between the upper limits of the two analyses, which would be unity if the analyses were exactly the same sensitivity. A histogram of values for this ratio,

$$\Omega = \frac{h_{0,\mathrm{S}}^{95\%}}{h_{0,\mathrm{h}}^{95\%}}, \tag{5.1.1}$$

is given in figure 5.1.2.

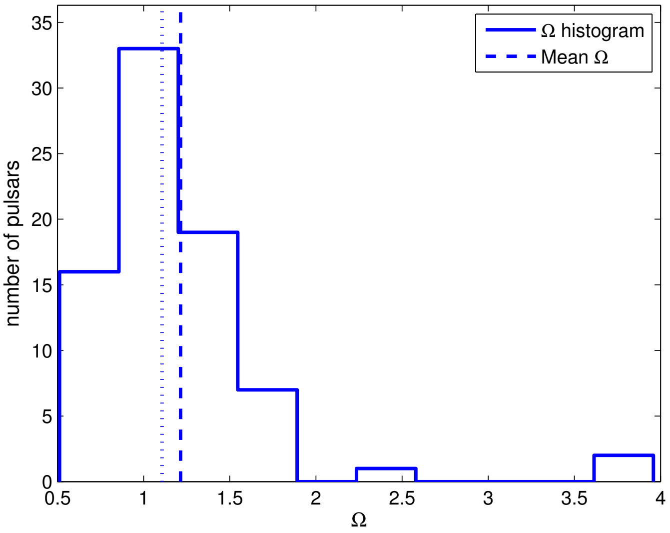

Figure 5.1.2: Histogram of Ω, the ratio of 95% upper limits from the SplInter/Nested sampling analysis and the heterodyne/MCMC analysis for isolated pulsars. Indicated by the dashed line is the mean value of Ω, and the dotted line indicates the mean value when ignoring the large outliers, which is 1.107.

The mean value of Ω (once outliers have been removed) is 1.107, and unity is within the uncertainty of the distribution. A slightly larger than unity Ω is expected due to the loss of the time for which data is dismissed as it is not in integer 1800 s chunks. The extra loss of SNR and hence higher $h_0^{95\%}$ values could be due to spectral leakage from the central peak in the expected sinc function as a result of the almost rectangular windowing

As noted in section 4.1.2, the recovery of signals from pulsars in binary systems by Spectral Interpolation may not be as accurate as for isolated pulsars. We analyse the pulsars in binary systems in the same way as we analyse the isolated pulsars, however we do not analyse pulsars which have been ruled out of analysis as a result of the discussion in section 4.1.2.

Table 5.3 shows the result of the targeted searches for continuous gravitational waves from binary pulsars, and figure 5.1.3 shows a histogram of Ω for the binary analysis. Of the sources in binary systems from table 4.4, we are unable to analyse J0605+3757 and J1630+3734 as the timing solutions are not applicable to the S6 dataset.

				SplInter & Nested Sampling			heterodyne & MCMC	
Pulsar	ν (Hz)	$\dot{\nu}$ (Hz/s)	$h_0^{\rm sd}$	$h_0^{95\%}$	$\varepsilon^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$	$h_0^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$
J0407+1607	38.91	-1.2×10^{-16}	3.5×10^{-28}	9.6×10^{-26}	6.1×10^{-5}	270	5.1×10^{-26}	145
J0437−4715	173.69	-4.1×10^{-16}	7.8×10^{-27}†	1.8×10^{-25}	2.2×10^{-7}	23	1.2×10^{-25}	15
J0614−3329	317.59	-1.7×10^{-15}	6.4×10^{-28}	7.0×10^{-26}	4.8×10^{-7}	110	8.5×10^{-26}	130
J0621+1002	34.66	-5.5×10^{-17}	5.4×10^{-28}†	1.1×10^{-25}	4.2×10^{-5}	210	9.6×10^{-26}	180
J0900−3144	90.01	-4.0×10^{-16}	2.1×10^{-27}	1.6×10^{-25}	3.8×10^{-6}	76	1.8×10^{-25}	86
J1017−7156	427.62	-3.8×10^{-16}	2.9×10^{-27}†	2.5×10^{-25}	8.6×10^{-8}	87	1.0×10^{-25}	34
J1022+1001	60.78	-1.6×10^{-16}	2.5×10^{-27}	7.0×10^{-26}	2.3×10^{-6}	28	4.8×10^{-26}	19
J1045−4509	133.79	-3.1×10^{-16}	5.3×10^{-27}†	3.1×10^{-26}	9.4×10^{-8}	5.8	3.0×10^{-26}	5.7
J1300+1240	160.81	-7.9×10^{-16}	3.0×10^{-27}†	4.5×10^{-26}	2.5×10^{-7}	15	4.9×10^{-26}	16
J1455−3330	125.20	-2.5×10^{-16}	1.5×10^{-27}†	3.6×10^{-26}	4.0×10^{-7}	24	3.6×10^{-26}	24
J1518+4904	24.43	-1.3×10^{-17}	8.5×10^{-28}†	4.5×10^{-25}	1.3×10^{-4}	540	4.5×10^{-25}	540
J1600−3053	277.94	-6.5×10^{-16}	5.1×10^{-28}†	7.7×10^{-26}	5.7×10^{-7}	150	6.7×10^{-26}	130
J1603−7202	67.38	-5.4×10^{-17}	4.4×10^{-28}†	3.8×10^{-26}	3.2×10^{-6}	86	2.3×10^{-26}	52
J1614−2230	317.38	3.9×10^{-16}	5.0×10^{-28}†	6.9×10^{-26}	2.9×10^{-7}	140	6.4×10^{-26}	130
J1623−2631	90.29	-5.1×10^{-15}	3.4×10^{-27}†	3.6×10^{-26}	1.9×10^{-6}	11	5.1×10^{-26}	15
J1640+2224	316.12	-1.6×10^{-16}	4.9×10^{-28}†	5.9×10^{-26}	1.7×10^{-7}	120	5.1×10^{-26}	100
J1643−1224	216.37	-8.5×10^{-16}	3.8×10^{-27}†	5.2×10^{-26}	1.1×10^{-7}	14	3.6×10^{-26}	7.8
J1701−3006A	190.78	-3.0×10^{-15}	4.6×10^{-28}‡	4.4×10^{-26}	2.0×10^{-6}	97	3.6×10^{-26}	78
J1709+2313	215.93	-6.9×10^{-17}	2.5×10^{-28}†	8.4×10^{-26}	7.8×10^{-7}	340	9.3×10^{-26}	370
J1713+0747	218.81	-3.9×10^{-16}	1.0×10^{-27}†	5.0×10^{-26}	2.6×10^{-7}	49	3.5×10^{-26}	35
J1719−1438	172.71	-1.5×10^{-16}	4.6×10^{-28}†	6.0×10^{-25}	7.8×10^{-6}	1300	1.6×10^{-25}	350
J1732−5049	188.23	-4.2×10^{-16}	6.7×10^{-28}†	5.8×10^{-26}	7.0×10^{-7}	86	4.6×10^{-26}	69
J1741+1351	266.87	-2.2×10^{-15}	2.5×10^{-27}	1.0×10^{-25}	3.1×10^{-7}	41	1.1×10^{-25}	44
J1745−0952	51.61	-7.6×10^{-17}	4.1×10^{-28}†	6.0×10^{-26}	1.3×10^{-5}	150	6.0×10^{-26}	150
J1748−2446E	455.00	-7.2×10^{-15}	5.8×10^{-28}‡	1.0×10^{-25}	6.6×10^{-7}	180	7.3×10^{-26}	130
J1748−2446Q	355.64	-5.6×10^{-15}	5.8×10^{-28}‡	1.0×10^{-25}	1.1×10^{-6}	180	9.4×10^{-26}	170
J1751−2857	255.44	-7.3×10^{-16}	9.5×10^{-28}	6.4×10^{-26}	3.3×10^{-7}	67	6.8×10^{-26}	72
J1801−3210	134.16	3.1×10^{-16}	2.4×10^{-28}†	4.2×10^{-26}	2.8×10^{-6}	180	3.5×10^{-26}	150
J1804−0735	43.29	-6.9×10^{-16}	4.1×10^{-28}‡	1.0×10^{-25}	1.0×10^{-4}	250	8.8×10^{-26}	220
J1804−2717	107.03	-4.7×10^{-16}	1.4×10^{-27}	3.4×10^{-26}	8.1×10^{-7}	23	2.2×10^{-26}	16
J1810−2005	30.47	-5.0×10^{-17}	2.6×10^{-28}†	2.4×10^{-25}	2.5×10^{-4}	950	1.6×10^{-25}	620
J1811−2405	375.86	-1.9×10^{-15}	1.1×10^{-27}	1.0×10^{-25}	2.9×10^{-7}	97	8.5×10^{-26}	77
J1841+0130	33.59	-9.2×10^{-15}	4.2×10^{-27}	1.4×10^{-25}	9.4×10^{-5}	34	1.3×10^{-25}	31
J1853+1303	244.39	-5.1×10^{-16}	7.3×10^{-28}†	6.2×10^{-26}	3.9×10^{-7}	85	8.5×10^{-26}	120
J1857+0943	186.49	-6.1×10^{-16}	1.6×10^{-27}†	5.6×10^{-26}	3.4×10^{-7}	35	5.7×10^{-26}	36
J1903+0327	465.14	-3.8×10^{-15}	3.6×10^{-28}†	1.2×10^{-25}	8.6×10^{-7}	340	1.6×10^{-25}	440
J1910+1256	200.66	-3.4×10^{-16}	5.4×10^{-28}†	5.8×10^{-26}	6.6×10^{-7}	110	7.7×10^{-26}	140
J1918−0642	130.79	-4.0×10^{-16}	1.0×10^{-27}†	3.4×10^{-26}	6.6×10^{-7}	34	4.0×10^{-26}	40
J1955+2908	163.05	-7.5×10^{-16}	3.2×10^{-28}†	4.9×10^{-26}	2.4×10^{-6}	150	5.4×10^{-26}	170
J2019+2425	254.16	-1.7×10^{-16}	7.2×10^{-28}†	5.1×10^{-26}	1.7×10^{-7}	71	5.6×10^{-26}	78

Continued

				SplInter & Nested Sampling			heterodyne & MCMC	
Pulsar	ν (Hz)	$\dot{\nu}$ (Hz/s)	$h_0^{\rm sd}$	$h_0^{95\%}$	$\varepsilon^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$	$h_0^{95\%}$	$h_0^{95\%}/h_0^{\rm sd}$
J2033+17	168.10	-2.3×10^{-16}	6.8×10^{-28}†	9.1×10^{-26}	1.0×10^{-6}	130	8.0×10^{-26}	120
J2129−5721	268.36	-1.5×10^{-15}	4.7×10^{-27}†	9.0×10^{-26}	1.2×10^{-7}	19	5.2×10^{-26}	11
J2145−0750	62.30	-1.0×10^{-16}	1.8×10^{-27}†	4.2×10^{-26}	1.4×10^{-6}	23	2.9×10^{-26}	16
J2229+2643	335.82	1.7×10^{-16}	4.1×10^{-28}†	6.8×10^{-26}	2.0×10^{-7}	170	6.5×10^{-26}	160
J2302+4442	192.59	-5.1×10^{-16}	1.8×10^{-27}	6.8×10^{-26}	3.2×10^{-7}	39	4.5×10^{-26}	25

Table 5.3: Upper limits on the gravitational-wave amplitude of continuous waves from known binary pulsars.

A † denotes that the pulsar's spin-down is corrected for proper motion effects.

A ‡ denotes that the pulsar's spin-down is calculated using a characteristic spin-down age of 10^9 years.

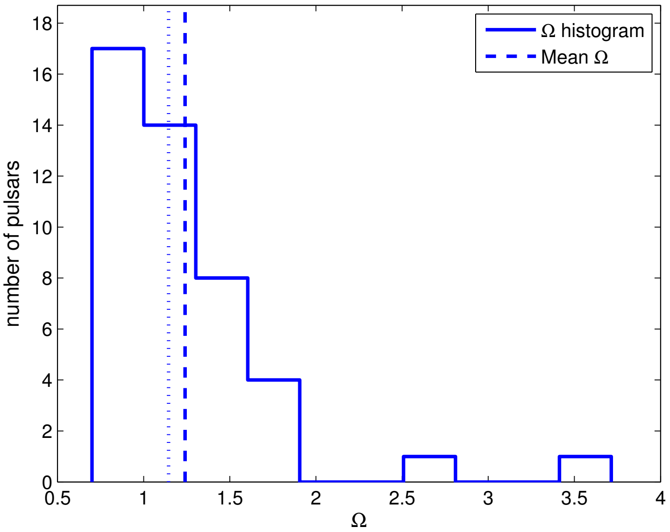

Figure 5.1.3: Histogram of Ω, the ratio of 95% upper limits from the SplInter/Nested sampling analysis and the heterodyne/MCMC analysis for binary pulsars. Indicated by the dashed line is the mean value of Ω, and the dotted line indicates the mean value when ignoring the large outliers, which is 1.15.

We again see that the 95% upper limits are all similar between the two searches, when considered with the expected differences discussed previously regarding the variability of $h_0^{95\%}$. The mean value of Ω in binary searches is 1.15, and unity is again within the uncertainty of the distribution, approximately as expected.

Chapter 6

Follow up of continuous-wave candidates from all-sky and directed searches

> *Space is big. You just won't believe how vastly, hugely, mind-bogglingly big it is. I mean, you may think it's a long way down the road to the chemist's, but that's just peanuts to space.*
>
> The Hitch-hikers Guide to the Galaxy, Douglas Adams

The speed increase of spectral interpolation compared to the heterodyne algorithm in both the $B_{k/K}$ calculation and the knock-on effect for the parameter estimation stages means that we are able to perform more computationally intensive procedures, such as rapid parameter estimation and evidence calculation on gravitational-wave candidates from all-sky and directed searches.

Using the Bayesian framework explained in section 2.2 and the SplInter algorithm in chapter 3, we can search the surrounding parameter space using an intelligent posterior sampling technique, rather than adaptive grid-based methods [91].

Section 6.1 introduces the all-sky and directed search methods used in the creation of these candidates, and explains the $\mathcal{F}$ Statistic used in frequentist searches for continuous gravitational waves. Section 6.2 then explains the method we use to follow up these candidates, and discusses how we can interpret the results.

Section 6.3 shows the results of a follow up search for a candidate from a directed search in the direction of the apparent stellar companion of Fomalhaut, which was found to have characteristics which could not rule out a neutron star ‘hiding’ behind the debris disk [75]. This is the first analysis using the Bayesian continuous-wave analysis to search over more than one frequency parameter. We also use this search as an example to show further look and detection protocols, i.e. how we take a candidate and go through detection criteria to decide whether this is a bona-fide gravitational-wave detection or not [65].

Section 6.4 shows the follow up of candidates from the Spotlight search aimed at two directions along the Orion spur [6]. In comparison to the Fomalhaut b search, we also include at this point a search over a relatively small patch of sky, according to the uncertainties for each candidate. This is the first search of its kind, using the continuous Bayesian framework to search in sky position.

6.1 Frequentist searches for signals with unknown origin: all-sky and directed searches

To follow up the candidates, we should understand the way in which they have been created so that we can know the extent of the parameter space which we must search and have an understanding of the meaning of the candidate information. An introduction to the searches which generated the candidates we use follows in sections 6.1.1 and 6.1.2, and section 6.1.1 includes a basic introduction to the $\mathcal{F}$ Statistic, which is a maximum likelihood estimator widely used in continuous gravitational-wave searches (e.g. in [1, 3, 11, 13, 15, 55, 69]).

6.1.1 Fully coherent directed searches for gravitational waves

The search method used to find the candidate used for section 6.3 is the fully coherent directed $\mathcal{F}$-statistic search. This search is in a particular direction which is known to a degree of accuracy such that it can be considered as a point source, but for a source for which the phase evolution is completely unknown. The details of the search are given in [4, 100], and an example of a search is given in [8].

The search templates are placed according to the three dimensional [$f^{(0)}$, $f^{(1)}$,$f^{(2)}$] parameter space metric, with a body centred cubic tiling[1], and then the $\mathcal{F}$ statistic is computed at each of these points. The templates with the largest $2\mathcal{F}$ are recorded, either by an $\mathcal{F}$-statistic cut-off ($2\mathcal{F} = 33.3768$ in [4]) or by selecting a proportion of the loudest candidates (0.01% in [8]).

These initial candidates are then filtered through a series of signal consistency checks, the first of which is an 'Fscan veto', which normalises SFTs and time-averages each frequency bin, the veto then removes any candidates within a number of bins whenever power deviates from the expected χ^2 distribution, this could be caused by non-stationary noise or spectral lines, this deviation is often defined by a threshold (± 7 standard deviations in [4], or 1.5 times the expected power in [8]). The next filter is to check consistency between the detectors, as the $2\mathcal{F}$ value should increase for the two detectors combined coherently compared to the individual detectors alone[2]. The remaining candidates are then assessed for statistical significance by comparing to a Gaussian noise distribution [8] and only those with less than 5% false alarm probability survive for human inspection. This human inspection involves comparison with known line artefacts in the data, and a check on the normalisation of the χ^2 distribution, as seen in figure 1 of [4].

The $\mathcal{F}$ Statistic

A popular tool for analysing the probability of a signal with certain Doppler parameters is the $\mathcal{F}$ statistic [55] used e.g. in [1, 15].

The $\mathcal{F}$ statistic is a maximum likelihood statistic given the Doppler parameters of the source, and is computed by comparing the data $s(t)$ to four basis waveforms per detector, $h_{1 \ldots 4}$. The basis waveforms are dependent upon the antenna patterns $a(t)$

---

[1] Though the extent of the $f^{(2)}$ parameter space will often be comparable to or less than the unit cell length, meaning it is effectively a two dimensional search

[2] We will use an analogous argument with regard to the SNR later when discussing the results of the follow ups.

and $b(t)$,[3] and the Doppler parameters λ, and are defined as

$$\begin{aligned} h_1(t;\lambda) &= a(t;\alpha;\delta)\cos\Phi(t;\lambda), & h_2(t;\lambda) &= b(t;\alpha;\delta)\cos\Phi(t;\lambda) \\ h_3(t;\lambda) &= a(t;\alpha;\delta)\sin\Phi(t;\lambda), & h_4(t;\lambda) &= b(t;\alpha;\delta)\sin\Phi(t;\lambda). \end{aligned} \tag{6.1.1}$$

The computation of these waveforms can be computationally limiting, and so the choice must be made between coherent or semi-coherent methods. The coherent method will require the same initial phase ϕ_0 throughout the search, but a semi-coherent search will allow the phase to change within specific intervals. This is advantageous for computational reasons as the longer phase templates are not required for the entire duration of the search. The $\mathcal{F}$-statistic method utilises an analytical maximisation over the amplitude parameters, **a**, and so these are not stated in results, though a 95% upper limit can be inferred through Monte Carlo simulations e.g. in [4].

The multi-detector $\mathcal{F}$ statistic is defined by [55][4]

$$\mathcal{F}(s;\lambda) := \frac{4}{|N(f)|T_{\mathrm{Obs}}}\frac{B|F_a|^2 + A|F_b|^2 - 2C\mathfrak{Re}[F_a F_b^*]}{D}, \tag{6.1.2}$$

where T_{Obs} is the observation time, $|N(f)|$ is the one sided spectral noise density, and F_a and F_b are defined by

$$F_a(f,\lambda) = \int_0^{T_{obs}} s(t)a(t;\lambda)e^{-i\Phi(t;\lambda)}\,\mathrm{d}t, \tag{6.1.3}$$

$$F_b(f,\lambda) = \int_0^{T_{obs}} s(t)b(t;\lambda)e^{-i\Phi(t;\lambda)}\,\mathrm{d}t, \tag{6.1.4}$$

with appropriate Doppler effects from detector and source motion applied. A, B, C

---

[3] These are $a(t)$ and $b(t)$ in section 1.3.2, and are defined fully in equations (12,13) of [55].

[4] A more elegant derivation is given in [85], but as the notation used in much of the literature is related to that of [55] we keep in line with that paper.

and D are defined by

$$\begin{aligned}
A &= \frac{2}{T_{\mathrm{Obs}}}\int_0^{T_{\mathrm{Obs}}} a(t)^2\,\mathrm{d}t,\\
B &= \frac{2}{T_{\mathrm{Obs}}}\int_0^{T_{\mathrm{Obs}}} b(t)^2\,\mathrm{d}t,\\
C &= \frac{2}{T_{\mathrm{Obs}}}\int_0^{T_{\mathrm{Obs}}} a(t)b(t)\,\mathrm{d}t,\\
D &= AB - C^2. \qquad (6.1.5)
\end{aligned}$$

The different computational implementations of this are discussed in [32, 70, 77]. The first effectively boils down to a summation estimate of the integrals of equations 6.1.3 and 6.1.4. For F_a, this summation is

$$F_a(f,\lambda) = \sum_{k=0}^{N_{SFTs}-1} a(t_k)e^{-i\Upsilon_k} \sum_{j=j^*-N_{terms}}^{j=j^*+N_{terms}} S_{kj}P_{kj}(f,\lambda), \qquad (6.1.6)$$

and F_b is similar, j denotes SFT frequency bin and k is the SFT, N_{SFTs} is the number of SFTs and N_{terms} is the number of SFT bins used, and j^* is the frequency bin nearest to the source frequency. Υ_k is a phase correction for each SFT, P_{kj} is the Dirichlet kernel, effectively the signal model to compare to the data, which is similar to the sinc interpolator we use in section 3.1.1, which both take into account the orbital motion effects.

The resampling implementation of the $\mathcal{F}$ statistic heterodynes SFTs to a desired frequency, inverts the Fourier transform to a time series which is then band-limited, downsampled and phase shifted for each datapoint according to the relative motion effects. This means that the $\mathcal{F}$ statistic can be efficiently calculated through the Fourier transform of this data.

The loudest $2\mathcal{F}$ values in a particular search will follow a χ^2 distribution with four degrees of freedom if searching in stationary Gaussian noise, however if there is a signal present, then a non-centrality parameter proportional to $\int_0^{T_{\mathrm{Obs}}} h(t)^2\,\mathrm{d}t$ is introduced.

6.1.2 The Spotlight search

The second method, which is used for the candidates in section 6.4 is the Spotlight search [6], which utilises the PowerFlux method [15] for a search pointed along the Orion spur. Although the search is aimed in a particular direction it is not a directed search, as the sky area searched is greater than the size of the unit cell of the template spacing. PowerFlux is a variant of the StackSlide method [28], in which the frequency bins of an SFT are slid according to barycentric corrections such that the prospective signal frequency is aligned and all of the power present is in one series of bins, the PowerFlux differs from a traditional StackSlide method in that the contributions from each SFT are weighted according to the sensitivity of the detector from the antenna pattern.

The version of the PowerFlux method implemented in the Spotlight search is a *loosely coherent* search, meaning that it does not use the power within each SFT[5], but includes the complex amplitude of the frequency bins, allowing the phase to vary within a pre-determined set [36]. The reason for this loosely coherent statistic is that it is designed to optimize processing for a set of templates rather than a single template as used in fully coherent searches. This optimization can be done for detection efficiency or for computational efficiency; however the gains for detection efficiency are much less than those for computational efficiency [37].

6.2 Bayesian follow up parameter estimation

In order to perform a parameter estimation search in Doppler parameter space, we alter the targeted search to include a narrow frequency band and a small patch of sky. Our signal must still be near to the initial Doppler parameters used in the spectral interpolation stage, as we require the frequency to be close enough to the used frequency so as not to degrade the SNR (see section 4.2.2), we also want a fairly small search space for computational expense reasons. In general, the candidates will come with a small enough uncertainty that both of these considerations are already satisfied.

The searches for unknown sources provide us with candidate locations, and we use

---

[5] This allusion to coherence means that it is no longer really a flux of power.

a nested sampling algorithm to search the amplitude parameters $[h_0, \phi_0, \cos\iota, \psi]$ as before, but also in the Doppler parameters of the source $[f^{(0)}, f^{(1)}, f^{(2)}, \alpha, \delta]$[6].

The search region for the Doppler parameters will preferably be a multivariate Gaussian with associated errors from the candidate generating algorithm if provided (as in section 6.4), though we can also use a uniform prior with limits set by repeated tests to find the optimal search area (as in section 6.3).

The most efficient way to search the parameter space would be to use an ellipsoid in parameter space defined by a covariance or correlation matrix according to the candidate position and the grid resolution, though these were not available for the searches presented here.

When we search in frequency and/or frequency derivatives, the high correlation between these parameters and the initial phase parameter ϕ_0 mean that we will not get much, if any, useful information about ϕ_0. We still retain ϕ_0 in the illustration of results in section 6.4, as a strong signal should have a peak in the posterior probability of ϕ_0 parameters.

The results of this will be a series of nest points taken during the nested sampling algorithm (see section 2.2.1), which are then converted into a posterior by accepting the nest point as a posterior point with a probability according to the prior volume multiplied by the likelihood. This means that the points with the lowest likelihood and the lowest prior probabilities will be mostly discarded. These posterior points are then placed into histogram bins, with the number of points in each bin of the histogram indicating the probability density that there is a signal with the given parameter within that range. These histograms will then be represented either in a one dimensional plot of the parameter value versus probability density or as a two dimensional contour plot, in which lines of constant probability are plotted, showing correlation between parameters.

---

[6] In this chapter we only consider the gravitational wave frequency f, not the rotational frequency ν, this is as the relationship between f and ν is ostensibly unknown.

6.3 Fomalhaut b directed search follow up

> *The eye was rimmed with fire, but was itself glazed, yellow as a cat's, watchful and intent, and the black slit of its pupil opened to a pit, a window into nothing.*
>
> *Lord of the Rings, J.R.R. Tolkien*

In a 2015 paper, Neuhäuser et. al, [75] postulated that an apparent companion object in the debris disk of Fomalhaut[7], known as *Fomalhaut b* or *Fom b*, could possibly be a background neutron star hiding behind the debris disk of the star and with very similar proper motion.

This would be an extremely interesting potential source of gravitational waves due to its proximity to the Earth, around 11 pc [75]. Considering the closest pulsar targeted in chapter 5 was 160 pc (J0437-4715), and the $1/r$ dependence of the amplitude of a signal (equation 2.1.3), a signal from this pulsar would be 14.5 times stronger than a signal from J0437-4715 and 25.5 times stronger than one from the closest isolated pulsar J0030+0451 (280 pc) given similar source and amplitude parameters. This proximity means that it is worth searching for, even if the likelihood is that it is not a neutron star.

A coherent directed search for gravitational waves from Fom b [52] found a single outlier which could not initially be attributed to lines in the data, the parameters of which are given in table 6.1. T_{coh} denotes the coherence time of the search, the time from the start to the end of the search, including any time for which the detector is not collecting data.

R.A., α (rad)	Dec, δ (rad)	$f^{(0)}_{\mathrm{cand}}$ (Hz)	$f^{(1)}_{\mathrm{cand}}$ (Hz/s)	$f^{(2)}_{\mathrm{cand}}$ (Hz/s^2)
6.0111077204	-0.51696653	738.9957991249271	-7.57352×10^{-11}	-1.803×10^{-19}

$2\mathcal{F}$	h_0 (est.)	Start (GPS time)	Finish (GPS time)	T_{coh}, days
71.1	1.31×10^{-24}	965 529 827	967 981 375	28.37

Table 6.1: Candidate parameter estimates and search information for the outlier from the directed search for Fomalhaut b [52].

---

[7] Dubbed 'the Great Eye of Sauron' [90] due the similarity of the images of the debris disk to the eye in the Lord of the Rings films released in the same year.

We perform parameter estimation using nested sampling on the candidate for the span of time given in table 6.1, so that we can confirm that we agree with the directed search that there could be a signal present, as well as to see the SNR and Bayes factor (evidence ratio, equation 2.2.4) for the signal hypothesis against the noise hypothesis. We use uniform priors as given in table 6.2 to give us all of the contained probability ranges for the Doppler parameters, these were found by repeated tests, expanding the prior ranges until all of the probability was contained, again we use uniform priors on the amplitude parameters as as discussed in section 2.2 to provide unbiased parameter estimation. The choice of upper limit on h_0 was so that we could conservatively take into account the results of [3] figure 9, setting the upper limit to be over two orders of magnitude above the all-sky upper limit at that frequency ($\sim 3 \times 10^{-24}$).

The candidate information contained $f^{(2)}$ values, so we decide to search over $f^{(2)}$, even though it is unlikely we would see strong effects from this parameter given the range from table 2.1 in this search time. We do not search over the right ascension and declination, as these are accurately known to sub-milli-arc-second precision, and so the search would not be able to distinguish a more accurate sky position within this range given this coherence time. An approximate sky resolution of the search from the estimates given in all-sky searches (such as [3]) would be

$$\Delta\theta = \frac{2c}{v f^{(0)} T_{\mathrm{coh}}}, \tag{6.3.1}$$

where v is the velocity of the detector, for which we use $30\,000\,\mathrm{ms}^{-1}$. For this frequency and coherence time, we would get a sky patch of size $\mathrm{d}\theta = 1.1 \times 10^{-5}\,\mathrm{sr}$, or 0.04 square degrees, much larger than any uncertainties on the position of the object from EM observation.

The results of this parameter estimation are given in figure 6.3.1. In figure 6.3.1 and in later plots (all plots to figure 6.4.5), we denote the frequency by the difference between the frequency and the candidate frequency, $f^{(0)} - f^{(0)}_{\mathrm{cand}}$, this is for aesthetic reasons, as we would otherwise not be able to read the axis labels due to the numerical precision required.

Figure 6.3.1 shows that we are able to agree with the directed $\mathcal{F}$-statistic search

parameter	lower limit	upper limit
h_0	0	10^{-21}
ψ	0	$\frac{\pi}{2}$
ϕ_0	0	2π
$\cos\iota$	-1	1
$f^{(0)}$ (Hz)	738.9957970	738.9958010
$f^{(1)}$ (Hz s^{-1})	-8×10^{-11}	-7×10^{-11}
$f^{(2)}$ (Hz s^{-2})	-6×10^{-18}	2×10^{-18}

Table 6.2: Prior values used in the search for the Fom b outlier. Our prior choices were made to be unbiased in amplitude parameters and to contain all of the probability in the Doppler parameter space.

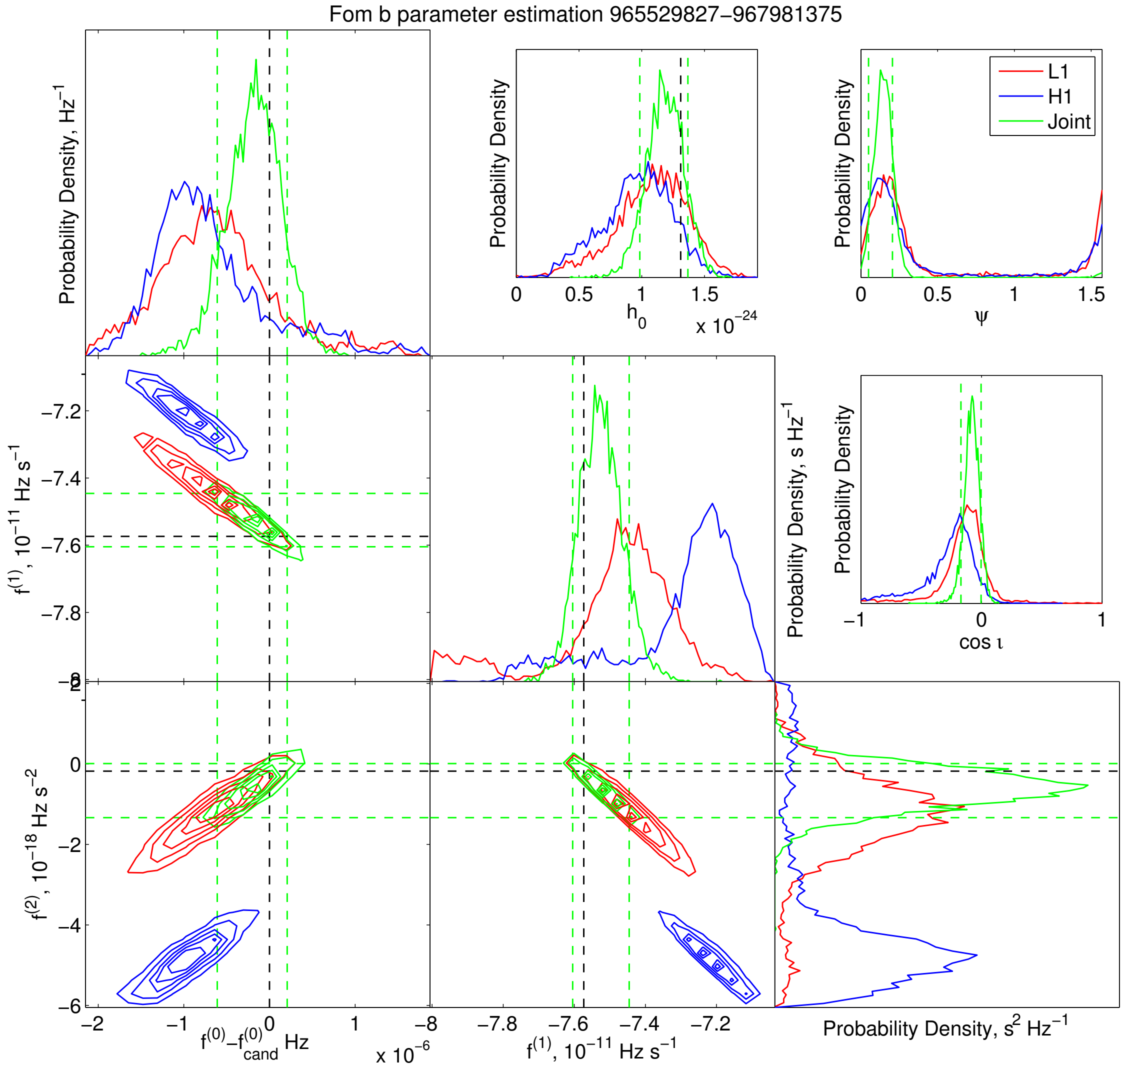

Figure 6.3.1: Posterior probability density distributions of the parameters for the Fom b candidate from H1 (blue), L1 (red) and joint (green) analyses. The black dashed line indicates the candidate parameters and estimated h_0 for the candidate from the fully coherent $\mathcal{F}$-statistic search. The green dashed lines indicate the 68% credible interval regions from the joint analysis.

that we see what looks like a signal in the data at this approximate frequency and spin-down, coming from both detectors. The parameter estimation shows that for h_0, $f^{(0)}$, $f^{(1)}$ and $f^{(2)}$ the candidate parameters are within the 68% credible intervals of the joint detector posteriors. The joint analysis uses the data from both detectors as the B_k input, meaning that the posteriors come from the joint likelihoods of all the analysis. There is a slight discrepancy between the $f^{(2)}$ parameters, though we should not be too disheartened by this, as it is difficult to resolve a second order spin-down at this coherence length. The contour plots show the approximate Gaussian distributions with correlations as expected from the signal model, and the joint posterior for h_0 has very small probability density at $h_0 = 0$. We also see from figure 6.3.2 that ϕ_0 is strongly peaked and consistent between detectors, again qualitatively improving the evidence that this is a signal.

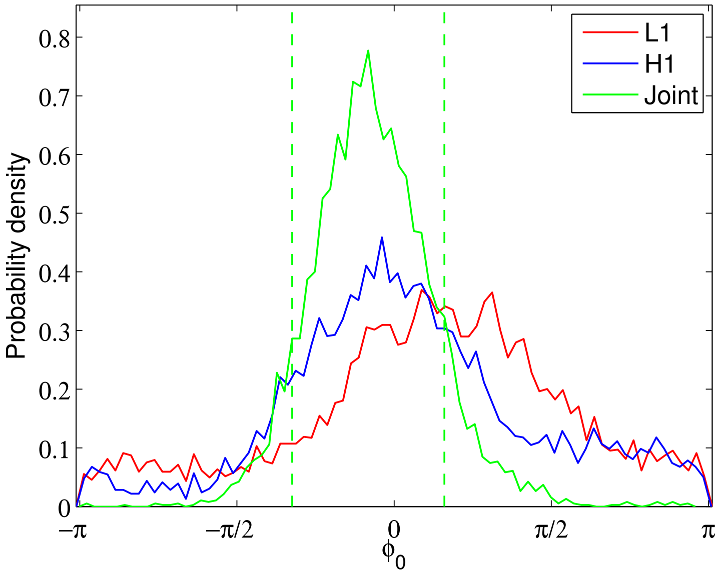

Figure 6.3.2: Posterior probability density distribution of the initial phase ϕ_0 parameters for the Fom b candidate from H1 (blue), L1 (red) and joint (green) analyses. The strong peak and the agreement between detectors for this parameter imply a true signal. The green dashed lines indicate the 68% credible interval regions from the joint analysis. As the probability peak is around $\phi_0 = 0$, we plot the probability density including ϕ_0 in the negative regions, rather than above $\phi_0 = \pi$, this is entirely equivalent.

The nested sampling algorithm gives us a $\log_{10}$ Bayes factor of 5.572 from the joint analysis, which means that the evidence for the signal model is higher than the noise-only model by a factor of $10^{5.572}$ (approximately $373, 232$). In the individual detectors, the log Bayes factors are -0.402 and -0.724 for H1 and L1 respectively, so individually we favour the noise model. This improved Bayes factor in the joint analysis implies a real signal, as the joint Bayes factor is much higher than for either of the individual detectors, or the combined incoherent Bayes factor $(-0.402 + -0.724 = -1.126)$.

The coherent SNR of the joint detector analysis was 8.49, of which 6.03 was from the H1 detector and 5.98 from L1, again indicating the presence of a signal as the joint SNR is greater than the individual values. The SNR from individual analyses was 6.59 for H1 and 6.18 for L1, as expected, these are higher than the individual interferometer SNRs stated in the joint analysis, this is as the most likely point in parameter space is different in the individual analyses and in the joint analysis.

6.3.1 Signal consistency check: increasing coherence time

To check for real signals we extend the coherence time of the search, for which the SNR should scale as $\propto \sqrt{T_{\text{coh}}}$. Therefore we perform the analysis again with the coherence length at twice and four times its initial value, which we now refer to as $T_{\text{coh,init}}$. This is a search over 56.75 days of data from GPS time 964 304 053 to 969 207 149 and 113.5 days from GPS time 961 852 505 to 971 621 841 respectively. The parameter estimation results of these analyses are presented in figures 6.3.3 and 6.3.4.

We see in figure 6.3.3 that the islands of probability are smaller than previously seen in figure 6.3.1, this is as the longer coherence time leads to a better resolution in parameter space. The posterior distributions are not consistent between detectors, which is not consistent with a true astrophysical signal.

We see in figure 6.3.4 that the islands of probability are again smaller than previously seen in figures 6.3.1 and 6.3.3, due to the longer coherence time. The posterior distributions are again not consistent between detectors, which is not consistent with an astrophysical signal.

Tables 6.3 and 6.4 show the $\log_{10}$ Bayes factors and SNRs for the individual and joint analyses as we increase the coherence times. We see that in the joint analysis, the Bayes factor falls from 5.572 to 2.652 for $T_{\text{coh}} = 2T_{\text{coh,init}}$ but then falls to -1456.008 for $T_{\text{coh}} = 4T_{\text{coh,init}}$, this means that the signal hypothesis is almost completely ruled out for this coherence time, this implies either a transient signal[8], or no signal at all.

In the joint analysis with $T_{\text{coh}} = 4T_{\text{coh,init}}$ we see a relatively high SNR signal

---

[8] Future work has been planned to investigate the possibility of transient continuous-wave signals.

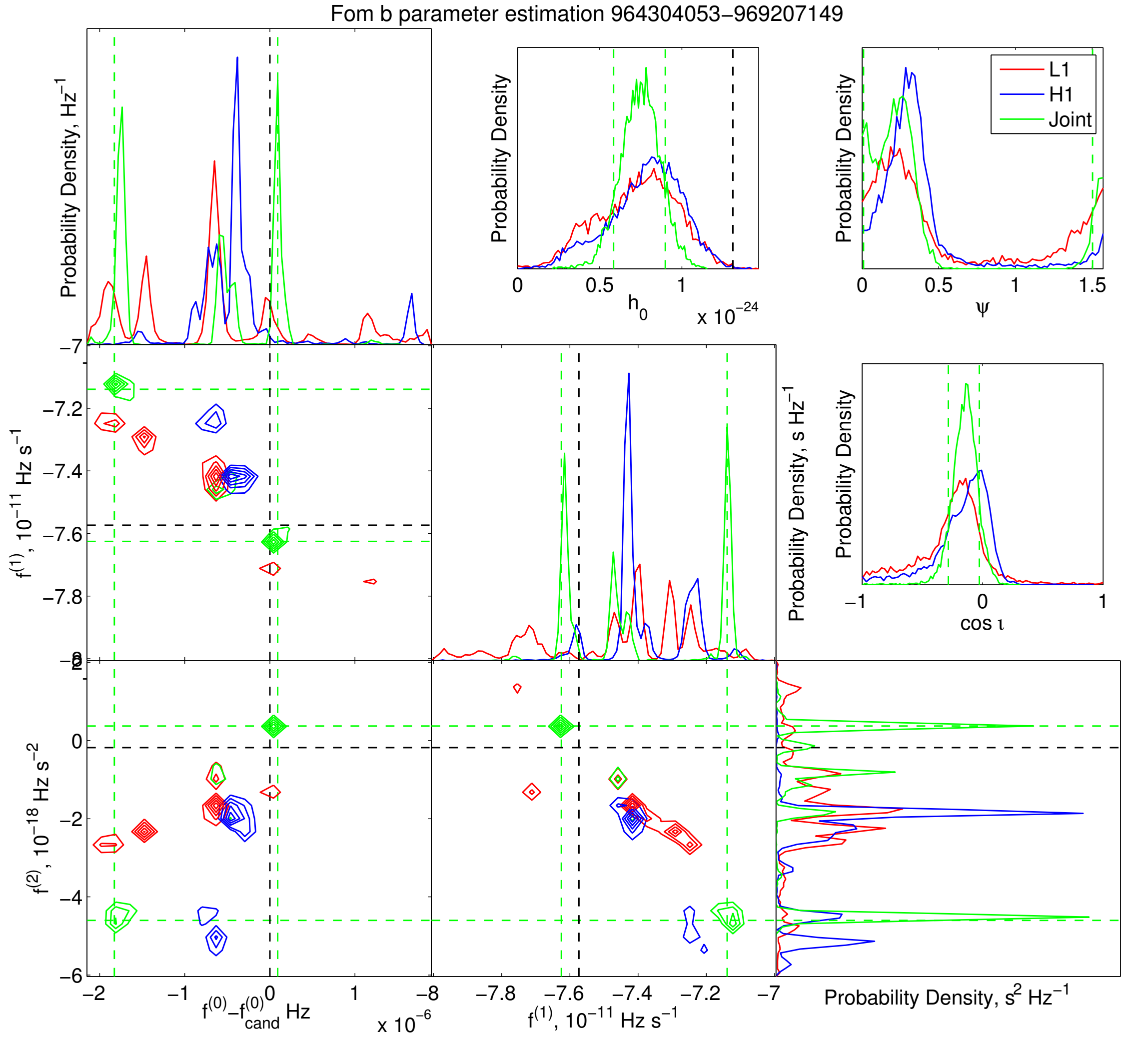

Figure 6.3.3: Posterior probability density distributions of parameters of the Fom b candidate with twice the initial coherence time (GPS times 964 304 053 to 969 207 149) from H1 (blue), L1 (red) and joint (green) analyses. The black dashed line indicates the candidate parameters and estimated h_0 for the candidate from the fully coherent $\mathcal{F}$-statistic search. The green dashed lines indicate the 68% credible interval regions from the joint analysis.

Individual analysis

Detector	H1		L1	
Coherence time	ρ	$\log_{10}(B)$	ρ	$\log_{10}(B)$
965 529 827 to 967 981 375	6.5866	−0.402	6.1837	−0.724
964 304 053 to 969 207 149	7.4006	−0.887	6.4665	−1.274
961 852 505 to 971 658 697	7.4806	−1.263	7.3100	−1458.746

Table 6.3: SNR and Bayes factor values from individual detector analyses for the initial Fom b followup, and analyses with twice and four times the initial coherence time.

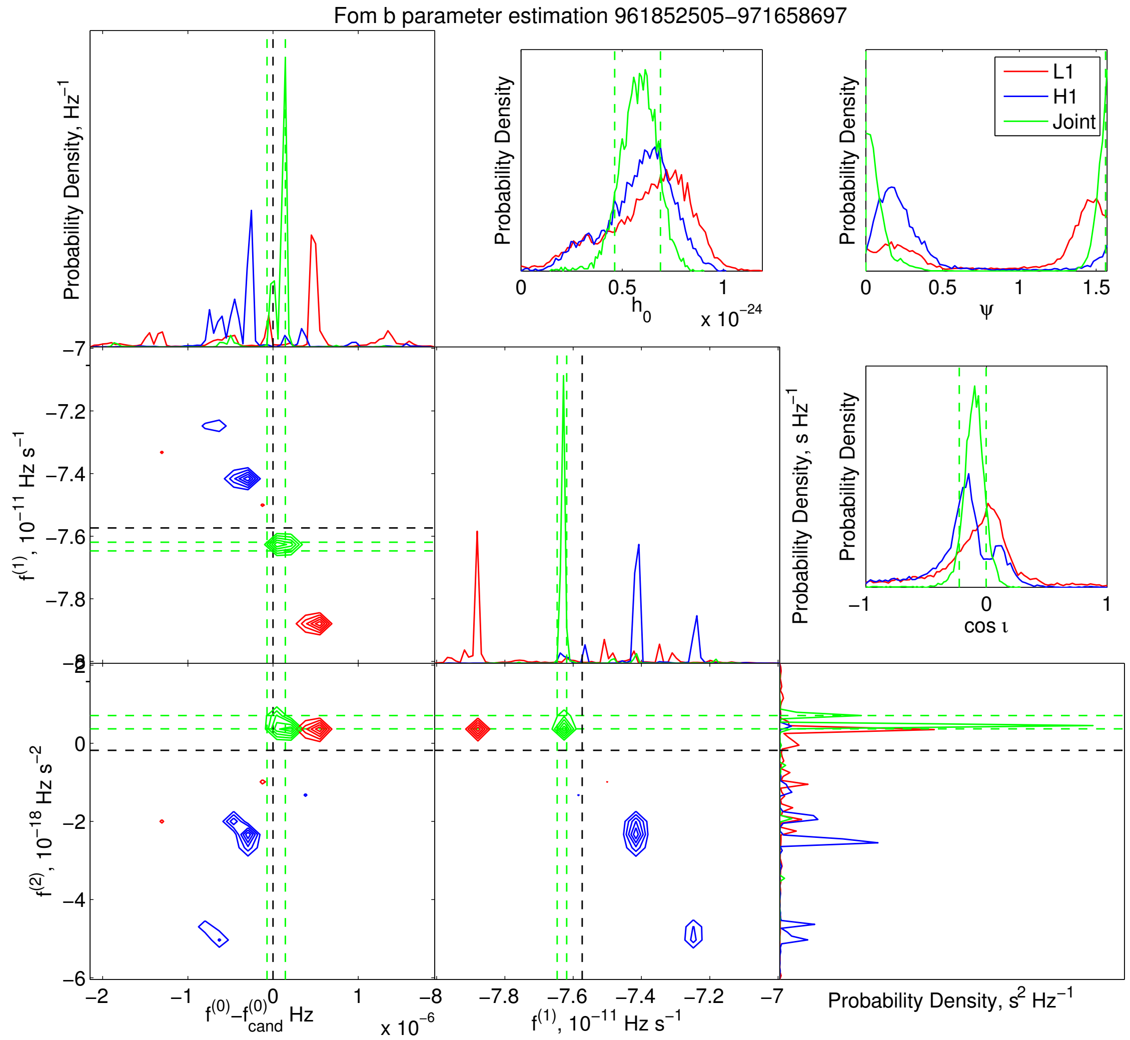

Figure 6.3.4: Posterior probability density distributions of parameters of the Fom b candidate with four times the initial coherence time (GPS times 961 852 505 to 971 658 697) from H1 (blue), L1 (red) and joint (green) analyses. The black dashed line indicates the candidate parameters and estimated h_0 for the candidate from the fully coherent $\mathcal{F}$-statistic search. The green dashed lines indicate the 68% credible interval regions from the joint analysis.

Joint analysis

Detector	H1	L1	Joint	
Coherence time	ρ	ρ	ρ	$\log_{10}(B)$
965 529 827 to 967 981 375	6.0308	5.9790	8.4923	5.572
964 304 053 to 969 207 149	6.2763	5.6093	8.4177	2.652
961 852 505 to 971 658 697	6.0360	5.6182	8.2460	−1456.008

Table 6.4: SNR and Bayes factor values from the joint detector analysis for the initial Fom b followup, and analyses with twice and four times the initial coherence time.

(8.2460) with an extremely disfavoured signal hypothesis ($\log_{10}(B) = -1456.008$), this seems somewhat counter-intuitive, and to investigate further, we perform the analysis with exactly the same set-up, but on white noise. This noise was created with the B_k drawn from a normal distribution characterised by the σ_k in the spectral interpolation output file from the analysis of Fom b, so that we track the same changes in the detector sensitivity over this band.

A similar test would be to shuffle the datapoints, randomly switching the timestamps of the data. This would lead to each of the datapoints having the same characterising noise as in the initial data, but any coherence would be removed, this version of the check is performed in [79]. The white noise analysis results are presented in table 6.5.

Joint analysis

Detector	H1	L1	Joint	
	ρ	ρ	ρ	$\log_{10}(B)$
Real analysis	6.0360	5.6182	8.2460	-1456.008
White noise	4.3728	3.9413	5.8868	-1278.960

Table 6.5: Comparison of SNR and Bayes factor in detector data and white noise for an analysis with four times the initial coherence time.

We see that we can find relatively high SNR values even in the region of the extremely disfavoured signal hypothesis, this is due to the SNR being calculated from the nested sampling point with the highest likelihood value

$$\rho = \sqrt{\sum_{i=1}^{N} \frac{y_i(\vec{\theta}_{\mathrm{max}})^2}{\sigma_i^2}}, \tag{6.3.2}$$

where $\vec{\theta}_{\mathrm{max}}$ is the parameter space point $\vec{\theta} = [\mathbf{a}, \lambda]$ with the maximum likelihood, whereas the Bayes factor is calculated from the evidences for the signal model, which marginalises over all the parameter space, and compares this to the noise hypothesis as in equation 2.2.4

$$B = \frac{p(\{B_k\} \mid 0 < h_0 \leq h_{\mathrm{max}})}{p(\{B_k\} \mid h_0 = 0)} = \frac{Z}{Z_{\mathrm{noise}}}. \tag{6.3.3}$$

This is due to the difference between a maximum likelihood and a marginalised likelihood estimate.

Using the longest stretch of data, we set a 95% upper limit on h_0, $h_0^{95\%}$ on gravitational waves from this direction, at this frequency and spin-down. For this analysis, we find that $h_0^{95\%} = 7.67 \times 10^{-25}$. This upper limit is greater than those in chapter 5, which is as expected considering that we use only around 24% of that coherence time and only use two detectors (H1 and L1) instead of three (H1,L1,V1). The spin-down limit for this source with the frequency and spin-down as in the candidate parameters is 2.35×10^{-23}, so if the source is truly a neutron star, we will have constrained h_0 to be less than 3.3% of the spin-down limit.

6.3.2 Finding the source of the outlier: decreasing coherence time

As the searches for longer coherence time seem to have implied that the signal does not last for longer than the initial search time, we consider the possibility that the signal could be a transient, monochromatic gravitational-wave signal, or a transient noise feature in the data. We investigate these possibilities by splitting the initial time limits in half, each with a coherence time of $T_{\mathrm{coh,init}}/2$.

Figures 6.3.5 and 6.3.6 show the results of parameter estimation in these two halves of the data.

We see that in the first half of the data, there is no obvious signal present from the H1 detector, and the posteriors look fairly uniform, with the contours spread over the entire prior range, representing the small variations in the noise. From L1 however we see a relatively strong, but not obviously signal-like posterior, which is particularly visible in the $f^{(0)}$ vs. $f^{(1)}$ and h_0 plots, and this is reflected in the joint posterior.

The second half of the data also shows no obvious signal, and the spread and disagreement between the detectors leads to qualitative disfavouring of the signal hypothesis. Despite this, this stretch of data had a SNR of 6.877 and a Bayes factor of 1.546 from the joint analysis, which shows that the signal evidence in the initial coherence time is mainly coming from this half of the data.

Tables 6.6 and 6.7 show the SNR in the first and second halves of the initial coherence time for individual and joint analyses. The SNR in the joint analysis of the first half of the data for both detectors is 5.594, of which 4.156 is from H1 and 3.744 is

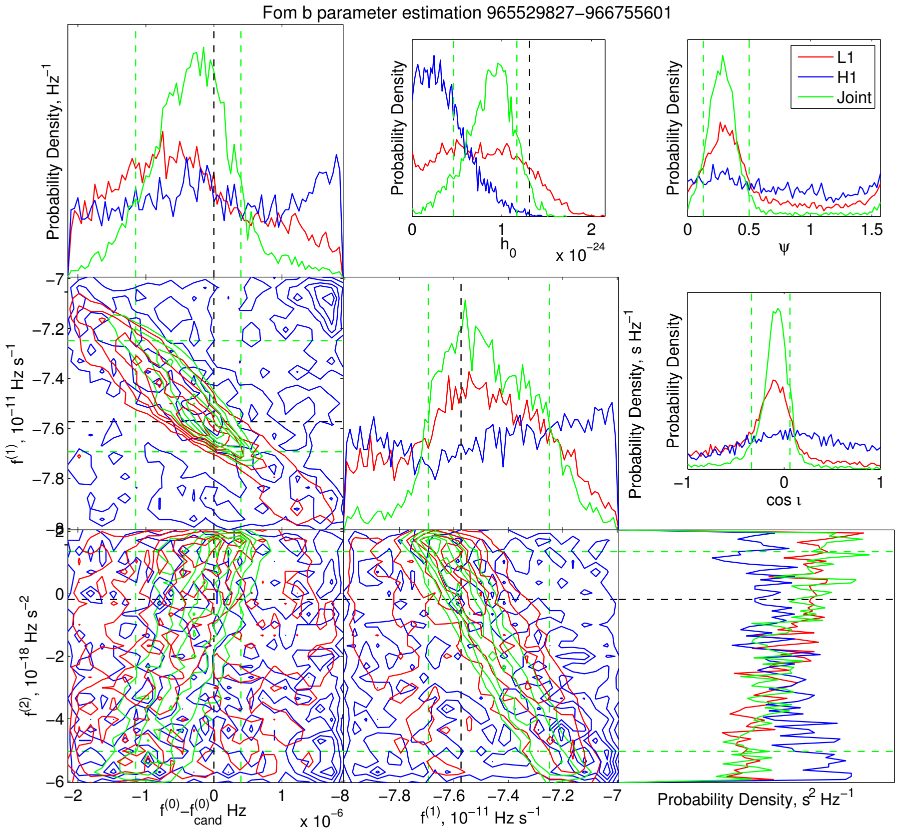

Figure 6.3.5: Posterior probability density distributions of amplitude parameters of the Fom b candidate for the first half of the initial coherence time from H1 (blue), L1 (red) and joint (green) analyses. Black dashed lines indicates the candidate parameters and estimated h_0 from the fully coherent $\mathcal{F}$-statistic search. The green dashed lines indicate the 68% credible interval regions from the joint analysis.

Individual analysis

Detector	H1		L1	
Coherence time (GPS time)	ρ	$\log_{10}(B)$	ρ	$\log_{10}(B)$
965 529 827 to 967 981 375	6.587	−0.402	6.184	−0.724
965 529 827 to 966 755 601 (1st half)	3.606	−2.990	4.521	−2.332
966 755 601 to 967 981 375 (2nd half)	6.290	−1.100	5.725	−1.229

Table 6.6: SNR and Bayes factor values from individual detector analyses for the Fom b followup in the first and second half of the initial coherence time.

Joint analysis

Detector	H1	L1	Joint	
Coherence Time (GPS time)	ρ	ρ	ρ	$\log_{10}(B)$
965,529,827 to 967,981,375	6.031	5.979	8.492	5.572
965,529,827 to 966,755,601 (1st half)	4.156	3.744	5.594	−1.346
966,755,601 to 967,981,375 (2nd half)	4.747	4.972	6.877	1.546

Table 6.7: SNR and Bayes factor values from joint detector analyses for the Fom b followup in the first and second half of the initial coherence time.

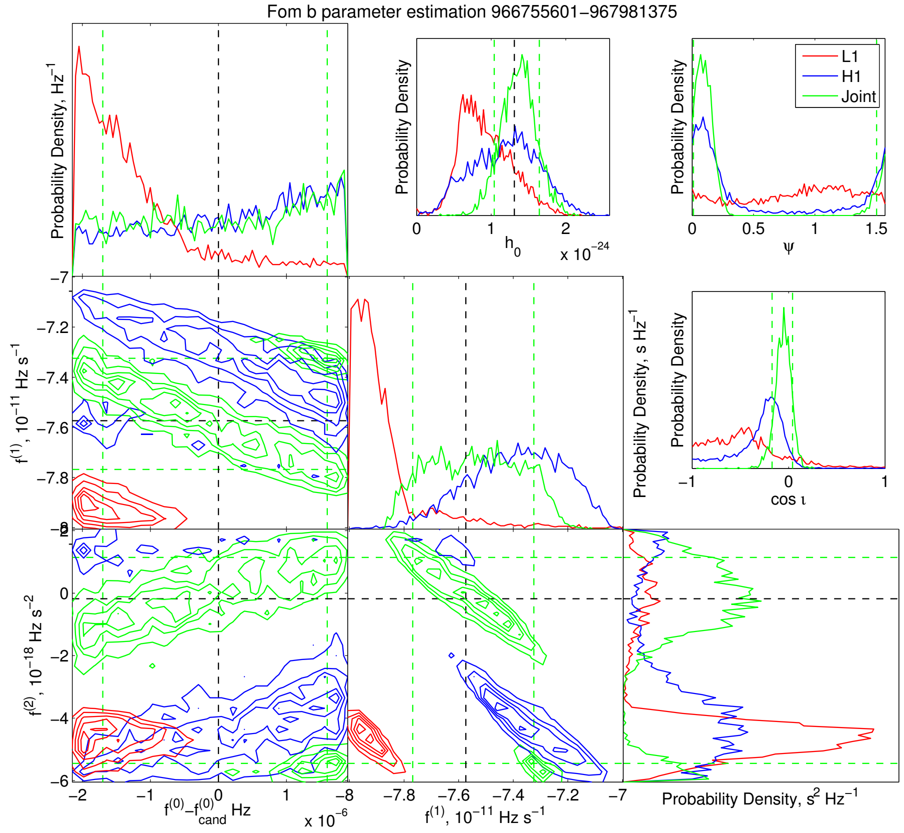

Figure 6.3.6: Posterior probability density distributions of amplitude parameters of the Fom b candidate for the second half of the initial coherence time from H1 (blue), L1 (red) and joint (green) analyses. Black dashed lines indicates the candidate parameters and estimated h_0 from the fully coherent $\mathcal{F}$-statistic search. The green dashed lines indicate the 68% credible interval regions from the joint analysis.

from L1. For the second half this is 6.877 for both detectors, 4.747 from H1 and 4.972 from L1, implying a signal favoured in the second half of the analysis. The $\log_{10}$ Bayes factors in the joint analysis of −1.346 and 1.546 again imply that the signal evidence is mainly coming from the second half of the coherence time.

Detector	Joint			
Initial time	5.572(8.492)			
Halves	−1.346(5.594)		1.546(6.877)	
Quarters	−2.620(3.796)	−1.994(4.615)	**2.610(6.761)**	−2.059(4.581)

Table 6.8: $\log_{10}$ Bayes factors and SNRs (in brackets) from halves and quarters of the initial coherence time for joint H1 and L1 analysis.

Detector	H1			
Initial time	−0.402(6.587)			
Halves	−2.990(3.606)		−1.100(6.290)	
Quarters	−2.844(3.141)	−2.644(3.989)	**−1.828(4.741)**	−6.144(3.927)

Detector	L1			
Initial time	−0.724(6.184)			
Halves	−2.332(4.521)		−1.229(5.725)	
Quarters	−2.903(2.882)	−2.456(4.201)	**−0.301(5.941)**	−77.296(3.886)

Table 6.9: $\log_{10}$ Bayes factors and SNRs (in brackets) from halves and quarters of the initial coherence time for individual H1 and L1 analyses.

We then cut the coherence length down again, investigating each quarter of the initial coherence time, performing a 'binary chop' to find the length of time most affected by this transient signal or noise feature.

The breakdowns of SNR and Bayes factors into halves and quarters in tables 6.8 and 6.9 show that the third quarter accounts for much of the signal evidence in both detectors, highlighted in bold in the tables. If we are to find a physical reason behind this outlier, whether a transient signal or a noise source, this is the best place to start looking.

Some further investigation into this approximate frequency with this span of data found noise lines in an auxiliary channel monitoring the output mode cleaner in L1 and a short duration artefact in H1 data from the pre-stabilised laser.

From this evidence, we conclude that this outlier is probably not an astrophysical signal. It is possible that the postulated neutron star glitched twice in this period of data (as in [50]), which could lead to a coherent template matching between the glitches and not matching outside of this period. However due to the evidence from the auxiliary channels, it is much more likely that a coincidence between the two noise sources is the cause of the signal evidence in this stretch of data.

6.4 Orion spur Spotlight search candidate followup

The Orion spur is a spoke-like region of the galaxy in which our solar system is found. As neutron stars are more likely to be found in populated concentrations of stars such as globular clusters [7] or the galactic centre [1], a speculative search has been carried out along this spur [6].

This search analysed S6 data, aiming in two directions; these are referred to as the ‘A’ and ‘B’ directions, specified in table 6.10. We reiterate here that although this search is pointed along the spur it is not a directed search, as the search areas (6.87° and 7.45° diameters respectively) are much too large to be searched with a single coherent template in sky position.

The Spotlight search utilised the Powerflux algorithm, [15], a variant of the Stack-Slide method, in which SFT frequency bins are shifted according to the Doppler shift at that time and summed at the (now monochromatic) frequency. The frequency range for this search was $f^{(0)} = [50, 1500]\,\mathrm{Hz}$ and the spin-down range was $f^{(1)} = [-5 \times 10^{-9}, 0]\,\mathrm{Hz\,s^{-1}}$, the time limits of this search run from GPS 951534120 (2010 Mar 02 03:01:45 UTC) to GPS 971619922 (2010 Oct 20 14:25:07 UTC), the second half of S6. The result after four stages of post-processing from the search is a set of seventy outliers, of which 37 come from the ‘A’ direction, and 33 from the ‘B’ direction.

Search region	Right ascension		Declination		Search radius
	Radians	Hour angle	Radians	Degrees	Degrees
A	5.283600	20h 10m 54.715s	0.585700	$33^\circ\ 33'\ 29.297''$	3.438
B	2.248610	08h 35m 20.607s	−0.788476	$-46^\circ\ 49'\ 25.151''$	3.724

Table 6.10: Spotlight search outlier search direction in radians and hour angle/degrees and the size of the patch on the sky covered by this search.

Of these seventy outliers, fourteen are not readily dismissed from line artefacts within the data; the parameters of these are given in table 6.11. The Spotlight search computed SNRs from seven equal-length segments of S6 data, which were then combined coherently with neighbouring segments to find the highest combined SNR value, given in the table with an indication of the inclusive range of subsegments used to calculate this SNR. Some of the outliers have parameters outside the initial range, e.g. A13 has $f^{(1)} = 9 \times 10^{-11}$, this is due to the initial post-processing procedure, which searches over spin-down, so it is possible for an outlier which is initially found with small spin-down to become a candidate with a small spin-up.

Outlier	$f^{(0)}_{\mathrm{cand}}$, Hz	$f^{(1)}_{\mathrm{cand}}$, Hz s^{-1}	α (rad)	δ (rad)	ρ	subsegments
A13	1138.509931	9×10^{-11}	5.225329	0.606463	9.72	[1,6]
A14	1404.892257	-1.205×10^{-09}	5.299464	0.642611	10.03	[0,6]
A24	1321.567031	-1.82×10^{-09}	5.318136	0.55852	8.83	[0,6]
A27	1474.94224	-2.05×10^{-09}	5.293492	0.563278	8.67	[0,5]
A28	990.761302	-2.705×10^{-09}	5.229674	0.580059	8.92	[0,6]
A29	1429.678924	-2.01×10^{-09}	5.301244	0.57326	8.61	[1,6]
A30	1325.509688	-4.325×10^{-09}	5.241072	0.598875	8.65	[0,6]
A33	1456.266111	1.95×10^{-10}	5.276759	0.586927	8.49	[0,5]
B15	613.261319	-3.95×10^{-09}	2.187829	-0.735556	9.54	[0,6]
B17	933.338229	1×10^{-10}	2.226266	-0.851431	9.30	[0,5]
B20	1249.438351	-1.55×10^{-09}	2.248791	-0.819053	8.64	[1,6]
B21	880.401753	-2.865×10^{-09}	2.284457	-0.828541	8.28	[1,6]
B23	1333.279063	-1.65×10^{-09}	2.23864	-0.835651	8.56	[0,6]
B30	1458.536476	-3.8×10^{-09}	2.298322	-0.754296	8.12	[0,6]

Table 6.11: Spotlight search outlier candidate parameters, with frequency, spin-down, sky position and the stretch of data for which they had the highest SNR [6].

The followup analysis was run using the parameters in table 6.11 to define a multivariate Gaussian with uncertainties from table 6.12 for priors on the Doppler parameters, and uniform priors were used on amplitude parameters as discussed in section 2.2. We note here that the uncertainty in sky position is dependent upon frequency, and so the sky patch is not consistently the same size, this is as the template mismatch in sky position increases for higher frequency, and so the template spacing is smaller.

Parameter	uncertainty	
$f^{(0)}$	5×10^{-5}	Hz
$f^{(1)}$	3×10^{-11}	Hz s^{-1}
α	$8/f^{(0)}$	rad Hz^{-1}
δ	$8/f^{(0)}$	rad Hz^{-1}

Table 6.12: Spotlight search outlier candidate uncertainties. The frequency and spin-down uncertainties are constant for all candidates, whereas the sky position uncertainties are inversely proportional to frequency. $f^{(0)}$ and $f^{(1)}$ from [6], α and δ from [37].

This search was performed using the nested sampling algorithm for individual and joint detector analyses of H1 and L1 data, for the stretch of data indicated by the subsegments in table 6.11. The analysis was run on the `atlas` cluster, and took an average of 2.6 days for the individual analyses and 5.9 days for the joint analyses.

Figures 6.4.1 to 6.4.4 show the parameter estimation from the nested sampling for the candidates A13, A14, A27 and B23[9]. Shown are the results from the H1, L1 and joint detector analyses, the candidate parameters are indicated by the solid black line, and the priors are indicated for the Doppler parameters as black dashed lines. We also show contour plots for $f^{(0)}$ vs. $f^{(1)}$ and α vs. δ, with contours given for the 1 and 2 σ probability levels. The contour levels are found by using a best fit of the amplitude of a normal distribution with parameters as in the priors to the posterior pdfs. As well as this we show the 1 and 2 σ contours for the prior normal distribution as the dashed black lines. Shown by green dashed lines are the 1σ credible intervals for the parameters from the joint analysis. These credible intervals and the priors are shown in the figures to illustrate how well the data fits a normal probability distribution. We also include contours which are at one and a half, and five times maximum of the best fit normal pdf and five times the best fit, which are shown by the thicker lines on the Doppler parameter contour plots, and indicate the parameter space points with excesses of probability compared to the parameter space around them, such as the large spikes in the joint posterior in figure 6.4.1.

We also perform the nested sampling on white noise with the same length of data and parameters and priors the same as candidate A13, the white noise is created in the same way as that discussed in section 6.3.1, utilising the σ_k estimates from the SplInter algorithm to randomly select a B_k from that distribution. We do this so that we can compare the results to the case in which we know that there is no signal present, the parameter estimation for this analysis is shown in figure 6.4.5.

We see that the posterior probabilities (except h_0) are dominated by the priors, with the amplitude parameters varying over the full range and the Doppler parameters being very close the the multivariate Gaussian used as the prior; this is visibly evidenced by the fact that the 1 and 2 σ data contours are very close to the prior contours, implying that there is no signal present in the data. We note that the candidates from a semi-coherent (or loosely coherent) search such as PowerFlux are more likely to give a

---

[9] The rest of the candidates in table 6.11 are presented in appendix D, figures D.0.1 to D.0.10.

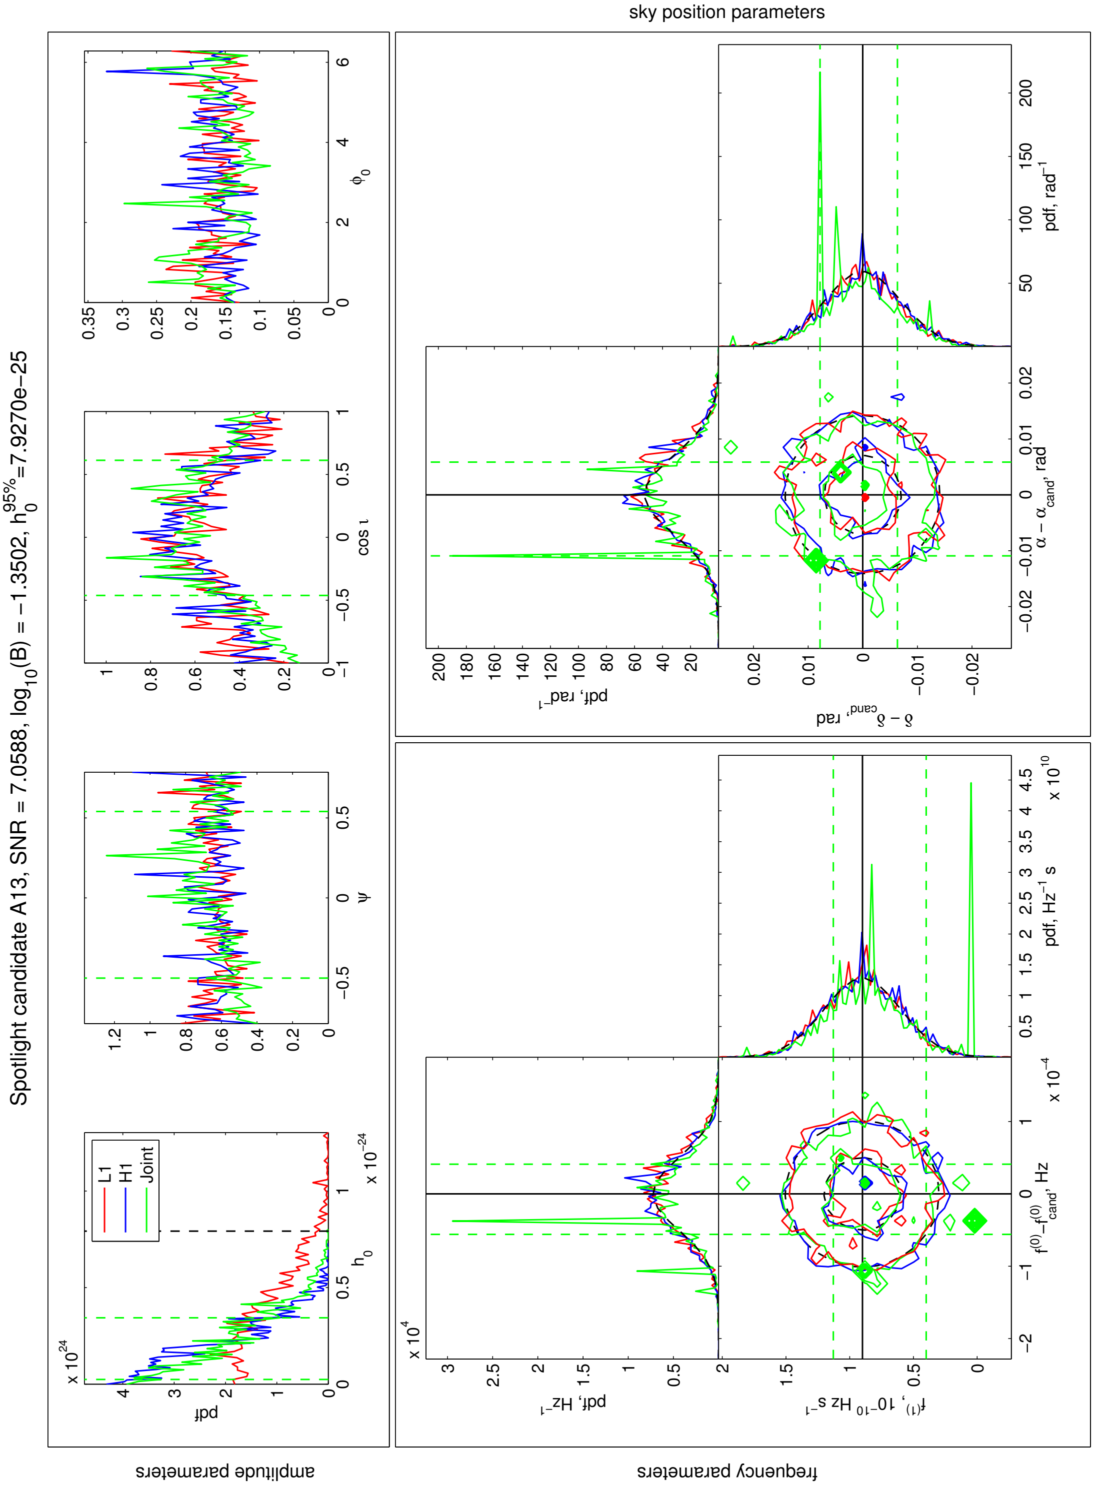

Figure 6.4.1: Posterior probability density distributions and contour plots for parameters of Spotlight candidate A13. $f^{(0)}_{\mathrm{cand}} = 1138.509931\,\mathrm{Hz}$.

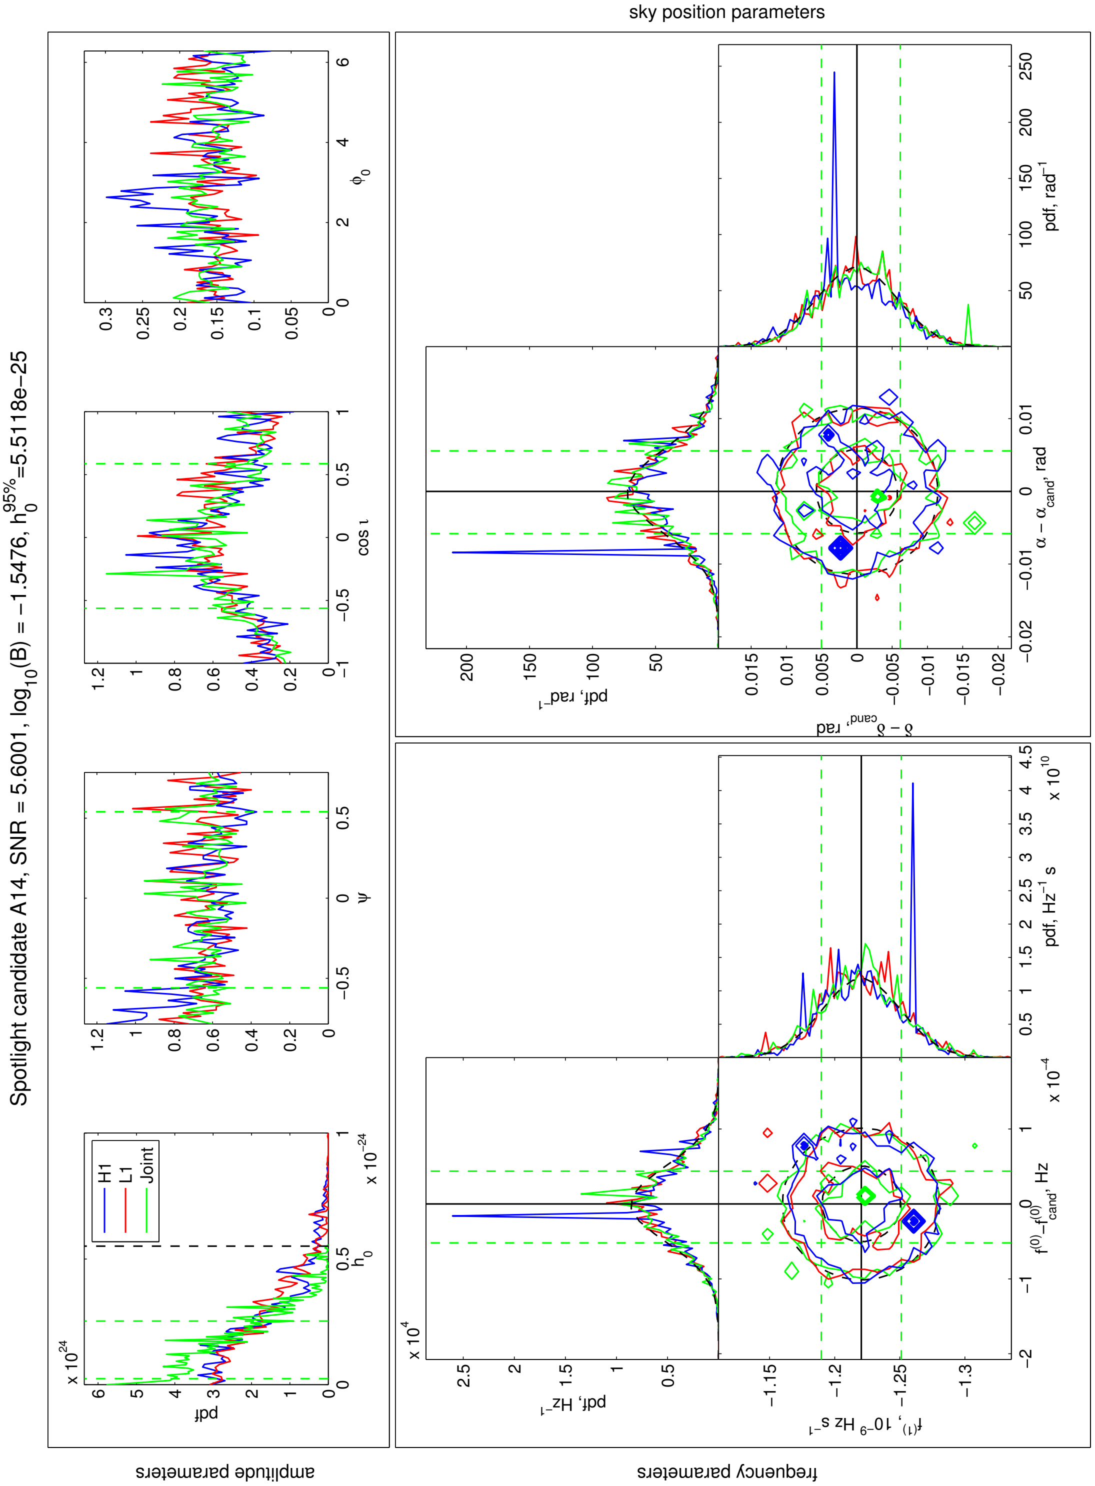

Figure 6.4.2: Posterior probability density distributions and contour plots for parameters of Spotlight candidate A14. $f^{(0)}_{\rm cand} = 1404.892257\,$Hz.

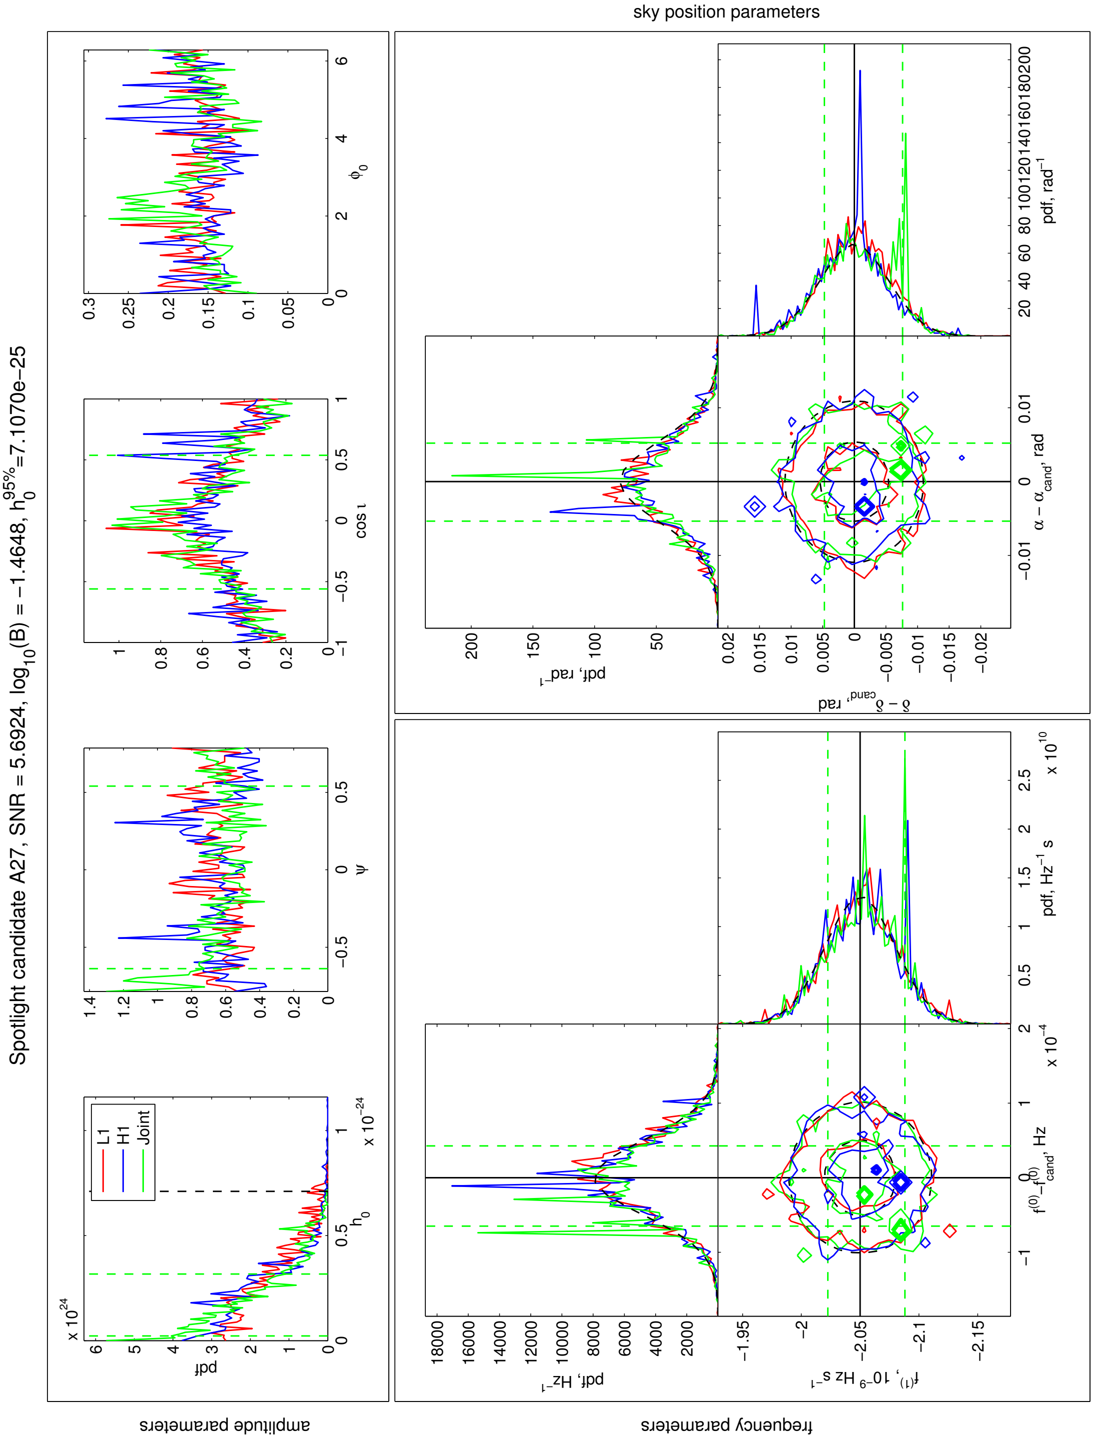

Figure 6.4.3: Posterior probability density distributions and contour plots for parameters of Spotlight candidate A27. $f^{(0)}_{\rm cand} = 1474.94224\,$Hz.

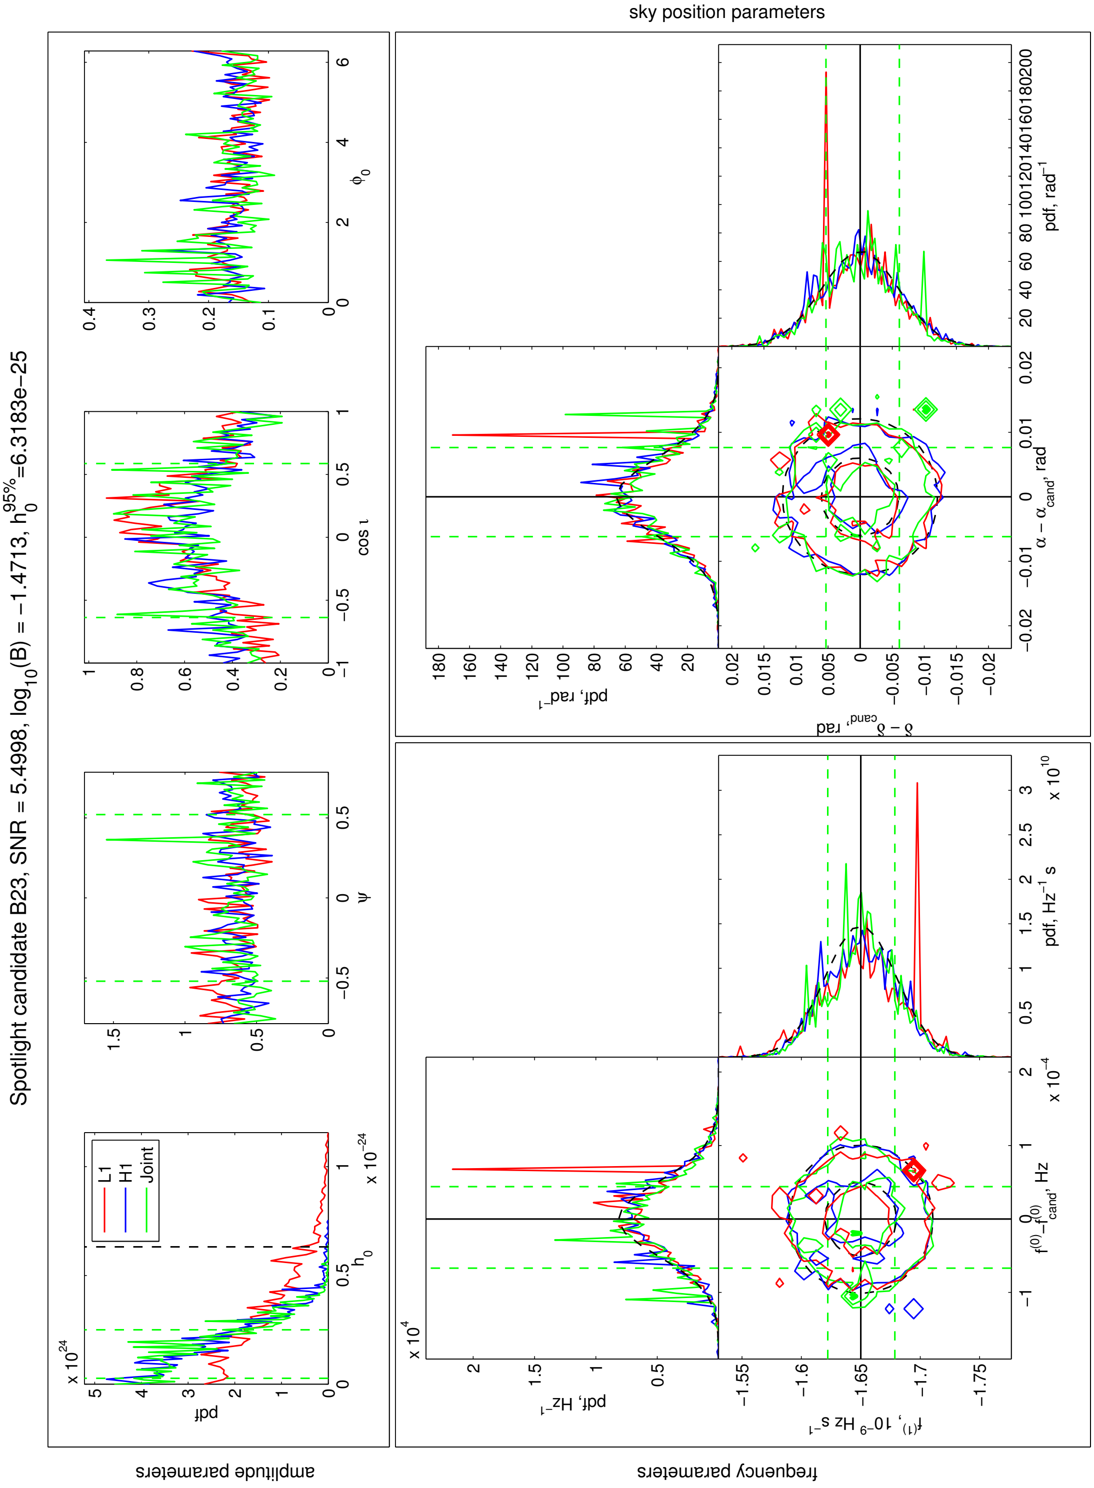

Figure 6.4.4: Posterior probability density distributions and contour plots for parameters of Spotlight candidate B23. $f^{(0)}_{\mathrm{cand}} = 1333.279063\,\mathrm{Hz}$.

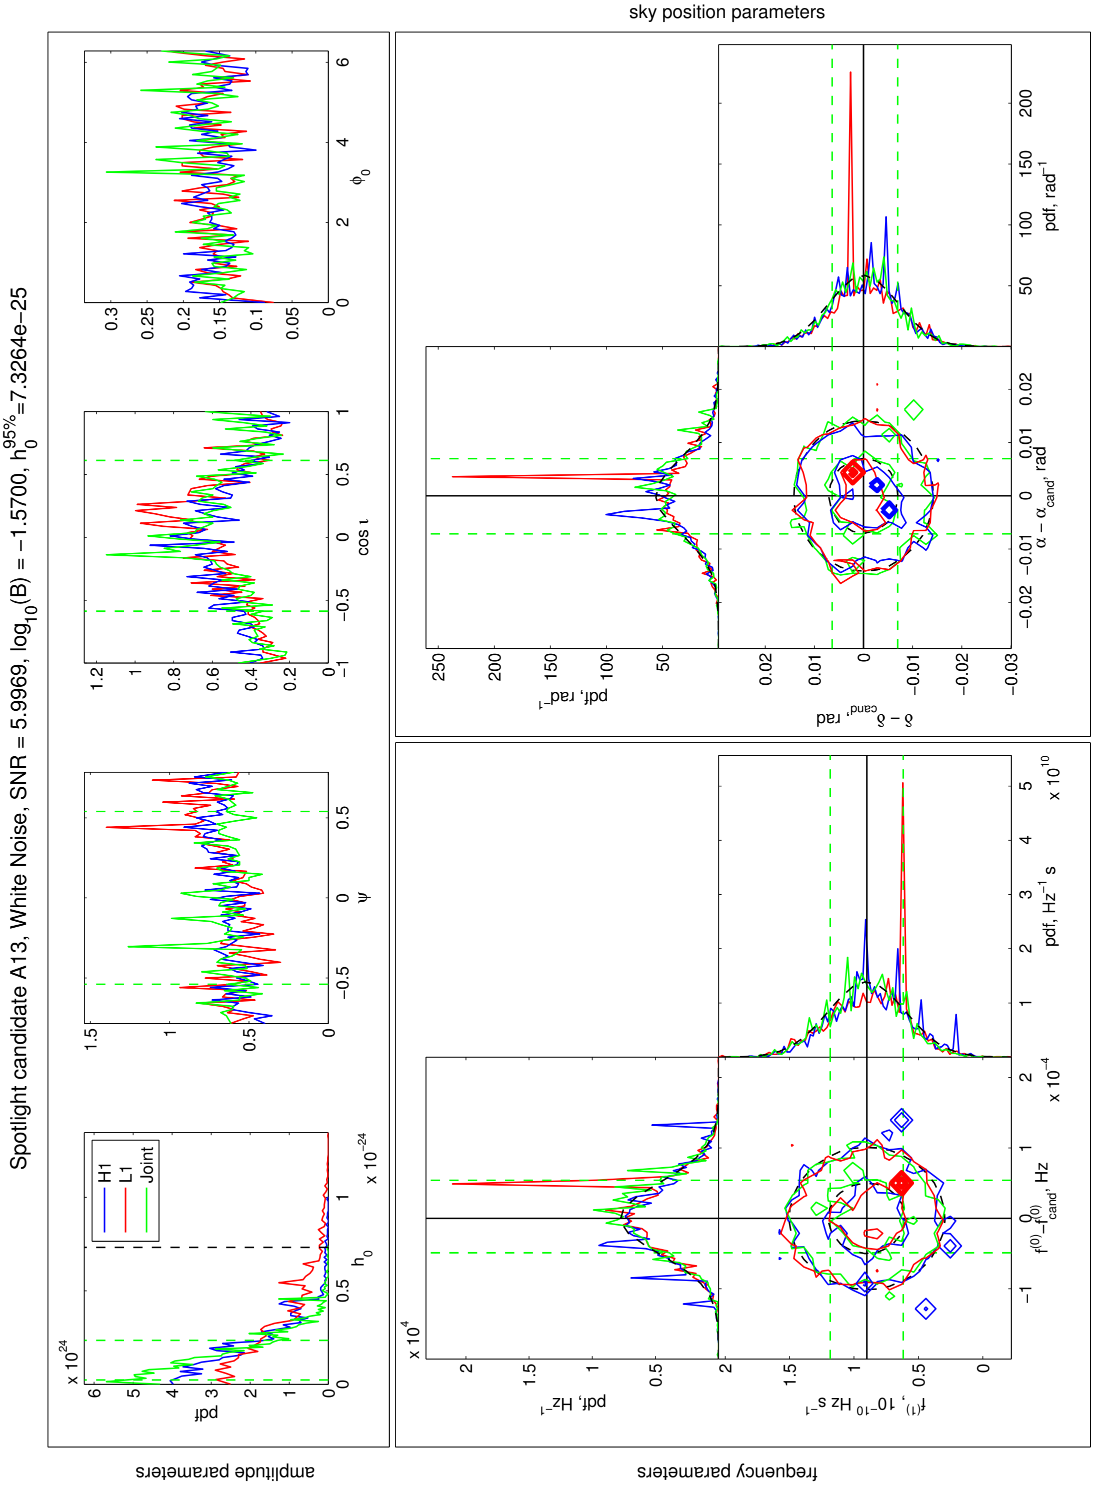

Figure 6.4.5: Posterior probability density distributions and contour plots for white noise analysed in the same way as Spotlight candidate A13. $f^{(0)}_{\rm cand} = 1138.509931\,$Hz.

candidate which looks like noise in the fully coherent data, this is as there may be a significant jump in phase between successive coherent data stretches, and so the loosely coherent template is favoured in a region where the fully coherent template is not. The Bayes factors and SNRs for each of the candidates in the individual and joint analyses are shown in table 6.13, with 95% upper limits on h_0 for that region of parameter space, these upper limits are indicated by black dashed lines on the h_0 posterior plots.

Outlier	H1		L1		Joint		
	ρ	$\log_{10}(B)$	ρ	$\log_{10}(B)$	ρ	$\log_{10}(B)$	$h_0^{95\%}$
A13	5.67272	-1.4270	4.89238	-1.2179	7.05884	-1.3502	7.9270×10^{-25}
A14	6.92681	-1.4287	5.38657	-1.3415	5.60011	-1.5476	5.5118×10^{-25}
A24	6.71552	-1.5730	5.69546	-1.4229	4.58003	-1.5984	4.7701×10^{-25}
A27	7.23744	-1.3730	3.77208	-1.2476	5.69235	-1.4648	7.1070×10^{-25}
A28	6.06536	-1.5156	4.29310	-1.5196	5.90139	-1.5872	5.1807×10^{-25}
A29	6.03214	-1.2712	6.46573	-1.3282	6.52650	-1.4117	8.3487×10^{-25}
A30	5.68745	-1.4047	6.58232	-1.4469	6.53262	-1.5861	5.7512×10^{-25}
A33	5.75062	-1.4347	4.92137	-1.3392	5.75062	-1.4347	6.4434×10^{-25}
B15	5.77548	-1.7557	6.13126	-1.5156	6.10400	-1.7656	3.3277×10^{-25}
B17	5.42201	-1.5793	6.24954	-1.5034	6.33777	-1.6786	3.5913×10^{-25}
B20	7.17421	-1.5582	5.15920	-1.4315	6.63466	-1.5601	6.2784×10^{-25}
B21	6.88002	-1.5702	6.34269	-1.4856	5.56983	-1.7117	3.5987×10^{-25}
B23	5.65441	-1.5115	6.79574	-1.2544	5.49981	-1.4713	6.3183×10^{-25}
B30	5.79615	-1.5088	6.52938	-1.3698	——		7.5222×10^{-25}
White Noise	7.41321	-1.4110	6.42666	-1.3280	5.99686	-1.5700	7.3264×10^{-25}

Table 6.13: Spotlight search outlier SNRs, Bayes factors and upper limits. The joint analysis for candidate B30 was cut off, and as such we present the individual analyses only. The $h_0^{95\%}$ is the value from incoherently combined individual analyses.

The joint analysis for candidate B30 was unfortunately cut off, without time for reanalysis for representation in this thesis, however an incoherent upper limit was set by multiplying the posterior pdfs on h_0 together and finding 95% of the integral under the resultant posterior density function. When doing this, we must be careful not to apply the priors twice, and we note that because the prior on h_0 is uniform, no 'prior removal' routine is required which would affect the 95% upper limit.

From table 6.13 we again see that the upper limits on the follow up candidates are much larger than those in chapter 5, due to using only half of the coherence time and two of the three detectors.

To be confident that we are returning the priors due to noise and not due to searching a parameter range which is too small for the search we consider the sky position accuracy of the search. If we compare the uncertainty in table 6.12 of $8/f^{(0)}$ to the uncertainty from equation 6.3.1, we can see the effective number of sky position templates that will fit into the prior search range,

$$N_{\text{templates}} \approx \frac{\Delta\theta}{8/f^{(0)}} = \frac{c}{4vT_{\text{coh}}}. \tag{6.4.1}$$

For a coherence length of the full duration of the search, this is around 8 000 templates, so we can be confident that the agreement between the posteriors and the priors is down to a lack of signal. A more basic, phenomenological argument would be that as we have single-point spikes in the data, the resolution must be smaller than the grid size of the plot.

We also consider the frequency and spin-down accuracy of the search; the uncertainties of 5 $\times 10^{-5}$ Hz in frequency and $3\times10^{-11}\,\text{Hz}\,\text{s}^{-1}$ in $f^{(1)}$, using an allowed mismatch of 0.3 gives an approximate frequency resolution of (e.g. from [3])

$$\Delta f^{(0)} = \frac{\sqrt{12m}}{\pi T_{\text{coh}}} \approx 3 \times 10^{-8}\,\text{Hz and} \tag{6.4.2}$$

$$\Delta f^{(1)} = \frac{\sqrt{720m}}{\pi T_{\text{coh}}^2} \approx 10^{-14}\,\text{Hz}\,\text{s}^{-1}, \tag{6.4.3}$$

which are both many orders of magnitude less than the width of the priors. To use the Bayes factor as a definitive detection or non-detection statistic, we would require many repeated studies of noise given the exact same experimental conditions, which is not feasible for the amount of repetitions needed to recreate the conditions properly. As figure 6.4.5 gives us a direct comparison for candidate A13, and an approximate comparison for the other candidates, we can however use this as a qualitative comparison. From the negative $\log_{10}$ Bayes factors, we can conclude that although we cannot entirely rule out the signal hypothesis, the noise hypothesis is favoured by factors of 16.5 ($\log_{10}(B) = -1.2179$) or more even in the individual detector analysis, and so a signal from any of these candidates is unlikely.

Chapter 7

Conclusion and future work

Bilbo: *Have you thought of an ending?*

Frodo: *Yes, several, and all are dark and unpleasant.*

Lord of the Rings, J.R.R Tolkien

As the anticipated observation of gravitational waves becomes more tangible, the prospect for astrophysics is intriguing, seeing the Universe in ways not possible with electromagnetic observation. One of the most exciting types of object to observe using gravitational waves is a neutron star. For example we can target an analysis for quasi-sinusoidal gravitational-wave emission from a known pulsar. The observation of these continuous waves will hopefully provide a fundamental probe of nuclear matter, as well as providing a fascinating new window for astronomy.

Until we observe gravitational waves though, discussing the potential astrophysics is like a song for the deaf, and therefore many algorithms have been developed for digging deep into the detector noise for traces of these signals. This work has focussed on the search for continuous waves from known pulsars or prospective signals using Bayesian inference.

No gravitational waves have been found in the work of this thesis using the continuous-wave searches presented, and we are not able to confirm or disprove their existence.

7.1 Spectral interpolation of fast Fourier transforms

In this thesis, particularly chapters 3 and 4, a new algorithm intended as a replacement for one of the vital cogs of one of the flagship searches for continuous gravitational waves from neutron stars, has been introduced, tested and applied to targeted searches. The testing of the Spectral Interpolation, *SplInter*, algorithm has shown that it is a viable replacement for the current heterodyne routine in most cases, particularly showing that the SplInter B_k output produces almost identical posterior probability densities for the amplitude parameters of hardware injected signals. The tests have also shown that we are able to utilise a Fresnel integral-based interpolator to search for signals with high instantaneous spin-down from the observatory's relative motion, $\dot{f}_k$, e.g. from isolated sources with low declination. A test was performed to see the limit to which we can utilise SplInter for signals from neutron stars in binary systems. Of the 97 sources in binary systems analysed, 47 were found to be in systems which could be analysed using the SplInter algorithm, and from these results, an empirical limit of $\nu^{(0)}a_1 < 10P_{\rm b}^3$ was set for binary systems in which sources can be analysed. This limit is based on the SFT, the Fourier transformed data, being taken using a time window of 30 mins, and so the limit would be much larger if a B_k separation of 60 s, as used in the heterodyne algorithm, could be used.

This algorithm has been used to analyse data from LIGO S5 in the first search for signals from pulsars at both twice the rotation frequency and the rotation frequency [79], as would be expected from a triaxial neutron star with non-aligned crust and core angular momenta, or a biaxial precessing star [57]. This work showed the advantages of the SplInter line removal routine, improving upon the S5 upper limits [17] for pulsar J1748−2446ac by a factor of 1.7 (7.19×10^{-26} to 4.2×10^{-26}) due to a wandering spectral line feature. This line was not present in S6 data, so this improvement is not seen in chapter 5.

This algorithm is now due to be used in the search for known pulsars with data from the advanced generation detectors. One option for future searches would be to develop an 'online' search, updating the parameter estimation regularly as data comes from the detector. For example a weekly update would use the posterior results from the previous week as the prior on the analysis for the new data, and by utilising a prior based on posteriors which involve all parameters it would remain a fully coherent search. This may be limited by the provision of SFTs, as the creation of these data files is currently not a regular or online occurrence.

The next stage of work to develop this algorithm would be to investigate the use of switching between the frequency and time domains so that we can reconstruct signals with higher $\dot{f}_k$ or $\ddot{f}_k$ in a more efficient way, such as sources in binaries with shorter periods, a process which could be similar to the resampling implementation of the $\mathcal{F}$-statistic. The concept for this work is to perform an inverse FFT on the SFT data, split the data into shorter segments, and then perform another FFT to return to the frequency domain. Because of the efficiency of the FFT and inverse FFT algorithms, this stage should not be particularly computationally intensive. At this point we could carry out the SplInter algorithm with a shorter Δt, and hence the allowed $\ddot{f}_k$ would increase according to $\ddot{f}_{k,\max} = 0.1/\Delta t^3$ as we saw in section 4.1.3. For optimal computational efficiency, it may even be possible to choose the new Δt based on a calculated $\dot{f}_k$ and $\ddot{f}_k$ of the signal, leading to analysis of sources in binaries with shorter periods and the Fresnel interpolation may no longer be required.

7.2 Targeted searches for continuous-wave signals from known pulsars

A targeted analysis using Spectral Interpolation output with nested sampling was performed in chapter 5, partly to show how the new procedure compares to the one which it is intended to replace, and partly as an example of an astrophysical search for representation in this thesis. We found that the 95% upper limits set on h_0, $h_0^{95\%}$ were comparable to those set by the heterodyne and MCMC routine in [2], and that S6 data was within a factor of 10 of the spin-down limit in some of the cases presented. A total of 97 sources were analysed, of which 45 were in binary systems; this is less than the number indicated in section 4.1.2 as two pulsars, J0605+3757 and J1630+3734 did not have timing solutions during the S6 dataset. The spin-down limits for known pulsars were found on average to be larger than those from the heterodyne and MCMC analysis by a factor of around 1.1 in both the isolated and binary cases; this is due to the use of 30 min SFTs causing a significant proportion of data in science segments to be lost, as discussed in section 3.4.

The next stage of this work could be to analyse the seven sources considered for special attention in [2]. Such a careful analysis, taking account of multiple glitches, is not required for the comparison between methods. If the work described above for adjustments to the SplInter algorithm is implemented, it would then be possible to include all sources in binaries in this targeted search.

The switch on of the advanced LIGO detectors in late 2015, and advanced Virgo in 2016-17 [64], should lead to a much higher probability of seeing gravitational waves, including continuous waves. The reduction in the noise floor of the detectors by an order of magnitude will mean more sensitive searches, and the widening of the search band will include pulsars with frequencies down to $f^{(0)} \simeq 10\,\mathrm{Hz}$, meaning that it will then be possible to analyse the population of young, non-recycled pulsars (we see this in figure 2.1.1). These young pulsars generally have high spin-down values, which could imply large gravitational-wave luminosities from kinetic energy considerations. As well as the wider frequency band and better sensitivity, more pulsars are being observed, which will provide more targets, increasing the likelihood of some of these targets having detectable gravitational radiation. For example with results from LOFAR [96], SKA [89] and AstroSAT [78], the expected number of known pulsars will increase by at least an order of magnitude [62] in the coming years. This increase in the number of targets increases the importance of the computational efficiency of targeted continuous-wave analysis pipelines, and the ability of the SplInter algorithm to analyse multiple sources at once is vital to this effort.

7.3 Rapid Bayesian follow up of continuous-wave candidates from all-sky and directed searches

Presented in chapter 6 was a follow up procedure for candidates from directed searches, utilising the SplInter algorithm for a rapid analysis of a candidate signal from a potential neutron star, Fomalhaut b, and candidates from a beamed search along the Orion spur. These searches found no evidence of continuous gravitational-wave signals from these candidates, and so an upper limit on h_0 on emission from these prospective sources was set. The search for signals from Fomalhaut b included a search over a frequency range of $4 \times 10^{-6}\,\mathrm{Hz}$, over a spin-down range of $10^{-11}\,\mathrm{Hz\,s^{-1}}$ and a frequency second derivative of $8 \times 10^{-18}\,\mathrm{Hz\,s^{-2}}$. This search was the first search using the Bayesian pipeline searching over more than one frequency parameter. The search for the Spotlight outliers searched a frequency range set by a normal distribution with an uncertainty of $5 \times 10^{-5}\,\mathrm{Hz}$, over a spin-down set by a normal distribution with an uncertainty of $3 \times 10^{-11}\,\mathrm{Hz\,s^{-1}}$ and a patch of sky defined by a normal distribution with an uncertainty of $8/f^{(0)}\,\mathrm{rad\,Hz^{-1}}$. The Spotlight follow up search was the first using the Bayesian pipeline to search within a patch of sky. This search was intended as an outline of a prospective follow up procedure, and more work would be required to optimise and automate the analysis.

The next stage of this work would be to optimise and automate the analysis, and to apply this search algorithm to candidates from other all-sky and directed searches. Again, if the adjustments to the SplInter algorithm as suggested above are implemented, then searches for prospective signals from sources in binary systems from the TwoSpect analysis are possible. Initially though it would be possible to make comparison to existing methods such as [91], and apply the follow up to upcoming results from the Einstein@Home bucket search.

The follow up analysis of the possible neutron star companion Fomalhaut b indicated an interesting line of investigation for future searches, including the possibility of searching for a transient monochromatic signal with all of the hallmarks of a traditional continuous wave, but with finite duration. This work is planned for the near future.

Appendix A

Regarding the derivation of the gravitational-wave equation in linearized gravity

Here we show a few of the points involved in the derivation of the gravitational-wave equation in linearized gravity which were skipped over in section 1.1.

Obtaining the Riemann curvature tensor in linearised gravity

The Riemann curvature tensor is defined by the Christoffel symbols, $\Gamma^{\theta}_{\mu\nu}$

$$R^{a}_{bcd} = \partial_c \Gamma^{a}_{bd} - \partial_d \Gamma^{a}_{bc}, \tag{A.0.1}$$

which are themselves defined by the metric

$$\Gamma^{a}_{bc} = \frac{1}{2} g^{ad} \left[\partial_b g_{dc} + \partial_c g_{db} - \partial_d g_{bc} \right]. \tag{A.0.2}$$

If use a metric which we consider as a perturbation $h_{\mu\nu}$ to flat spacetime $\eta_{\mu\nu}$, such that $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, this is then

$$\Gamma^{a}_{bc} = \frac{1}{2}(\eta^{ad} + h^{ad}) \left[\partial_b(\eta_{dc} + h_{dc}) + \partial_c(\eta_{db} + h_{db}) - \partial_d(\eta_{bc} + h_{bc}) \right]. \tag{A.0.3}$$

In linearised gravity, $h_{\mu\nu}$ and its derivatives are small, so we remove terms of higher than linear order in $h_{\mu\nu}$ or its derivatives. Also as $\eta_{\mu\nu}$ is constant, we can remove all terms which are a derivative of $\eta_{\mu\nu}$

$$\Gamma^a_{bc} = \frac{1}{2}\eta^{ad}\left[\partial_b h_{dc} + \partial_c h_{db} - \partial_d h_{bc}\right]. \tag{A.0.4}$$

We then use these Γ^a_{bc} in equation A.0.1

$$\begin{aligned} R^a_{bcd} =& \frac{1}{2}\big\{\partial_c \eta^{ae}\left[\partial_b h_{de} + \partial_d h_{be} - \partial_e h_{bd}\right] - \partial_d \eta^{ae}\left[\partial_b h_{ce} + \partial_c h_{be} - \partial_e h_{bc}\right]\big\} \\ =& \frac{1}{2}\eta^{ae}\big\{\partial_c\partial_b h_{de} + \partial_c\partial_d h_{be} - \partial_c\partial_e h_{bd} - \partial_d\partial_b h_{ce} - \partial_d\partial_c h_{be} + \partial_d\partial_e h_{bc}\big\} \\ =& \frac{1}{2}\eta^{ae}\big\{\partial_c\partial_b h_{de} - \partial_c\partial_e h_{bd} - \partial_d\partial_b h_{ce} + \partial_d\partial_e h_{bc}\big\}, \end{aligned} \tag{A.0.5}$$

where the $+\partial_c\partial_d h_{be}$ and $-\partial_d\partial_c h_{be}$ terms cancel because the partial derivatives commute. This final line is the result we stated in equation 1.1.4.

Gauge transformations

In the gauge transformation $h_{ab} \to h'_{ab} = h_{ab} - \partial_a\xi_b - \partial_b\xi_a$, we state that the Einstein equation remains valid as all the ξ terms cancel by symmetry, we see this in the Riemann curvature tensor. Equation 1.1.4 gives us

$$\begin{aligned} 2R^a_{bcd} = \eta^{ae}\big\{&\partial_c\partial_b h'_{de} - \partial_c\partial_b\partial_d\xi_e - \partial_c\partial_b\partial_e\xi_d - \partial_c\partial_e h'_{bd} + \partial_c\partial_e\partial_b\xi_d + \partial_c\partial_e\partial_d\xi_b \\ &- \partial_d\partial_b h'_{ce} + \partial_d\partial_b\partial_c\xi_e + \partial_d\partial_b\partial_e\xi_c + \partial_d\partial_e h'_{bc} - \partial_d\partial_e\partial_b\xi_c - \partial_d\partial_e\partial_c\xi_b\big\}, \end{aligned} \tag{A.0.6}$$

for which many of the terms on the brackets cancel – the second with the eighth (the ξ_e terms) the third with the fifth (the ξ_d terms), the sixth with the final term (the ξ_b terms) and the ninth and the eleventh (the ξ_c terms) leading to

$$R^a_{bcd}\frac{1}{2}\eta^{ae}\big\{\partial_c\partial_b h'_{de} - \partial_c\partial_e h'_{bd} - \partial_d\partial_b h'_{ce} + \partial_d\partial_e h'_{bc}\big\}, \tag{A.0.7}$$

which is the same as equation 1.1.4, but in the h' gauge, meaning that the physical results of being in the h' gauge are the same as in the original frame of reference.

The Ricci tensor in the harmonic gauge

We now consider obtaining equation 1.1.6 for the Ricci tensor in the harmonic gauge. Given $R_{bc} = R^a_{bca}$, from equation 1.1.4 (or A.0.5), we can see that

$$\begin{aligned}
2R^a_{bca} =& \eta^{ae}\left\{\partial_c\partial_b h_{ae} - \partial_c\partial_e h_{ba} - \partial_a\partial_b h_{ce} + \partial_a\partial_e h_{bc}\right\} \\
2R_{bc} =& \partial_c\partial_b h^a_a - \partial_c\partial_a h^a_b - \partial_a\partial_b h^a_c + \partial_a\partial^a h_{bc} \\
=& \partial_c\partial_b h - \partial_c\partial_a h^a_b - \partial_a\partial_b h^a_c + \square h_{bc},
\end{aligned} \tag{A.0.8}$$

where $h = h^a_a$ and $\square = \partial_a\partial^a$ as in section 1.1.

We then use the trace-reverse of the perturbation, $\bar{h}^a_b = h^a_b - \frac{1}{2}\eta^a_b h$, and analogously $h^a_b = \bar{h}^a_b - \frac{1}{2}\eta^a_b\bar{h}$, and the Ricci tensor becomes

$$\begin{aligned}
2R_{bc} =& \square h_{bc} + \partial_c\partial_b\bar{h} - \partial_c\partial_a\bar{h}^a_b - \partial_a\partial_b\bar{h}^a_c - \frac{1}{2}\left(\partial_c\partial_b\eta^a_a\bar{h} - \partial_c\partial_a\eta^a_b\bar{h} - \partial_a\partial_b\eta^a_c\bar{h}\right) \\
=& \square h_{bc} + \partial_c\partial_b\bar{h} - \partial_c\partial_a\bar{h}^a_b - \partial_a\partial_b\bar{h}^a_c - \frac{1}{2}\left(4\partial_c\partial_b\bar{h} - \partial_c\partial_b\bar{h} - \partial_c\partial_b\bar{h}\right) \\
=& \square h_{bc} - \partial_c\partial_a\bar{h}^a_b - \partial_a\partial_b\bar{h}^a_c + \partial_c\partial_b\bar{h} - \frac{1}{2}\left(2\partial_c\partial_b\bar{h}\right) \\
=& \square h_{bc} - \partial_c\partial_a\bar{h}^a_b - \partial_a\partial_b\bar{h}^a_c.
\end{aligned} \tag{A.0.9}$$

In the harmonic gauge, $\partial_a\bar{h}^a_b = 0$, so the final two terms are zero, leaving

$$R_{bc} = \frac{1}{2}\square h_{bc}, \tag{A.0.10}$$

which is equation 1.1.6. When we use this in the Einstein equation 1.1.1, we get

$$\begin{aligned}
G_{ab} =& R_{ab} - \frac{1}{2}g_{ab}R \\
=& \frac{1}{2}\square h_{ab} - \frac{1}{2}g_{ab}\frac{1}{2}\square h \\
=& \frac{1}{2}\square\left(h_{ab} - \eta_{ab}\frac{1}{2}h\right) \\
=& \frac{1}{2}\square\bar{h}_{ab},
\end{aligned} \tag{A.0.11}$$

where we have replaced g_{ab} with η_{ab} on the third line by remembering that in linearised gravity, anything higher than linear order in h is zero. We can therefore use equation 1.1.1 to relate this to the stress-energy-momentum tensor T_{ab}

$$\Box \bar{h}_{ab} = -\frac{16\pi G}{c^4} T_{ab}, \tag{A.0.12}$$

which is equation 1.1.7, and which is $\Box \bar{h}_{ab} = 0$ in free space and $\Box h_{ab} = 0$ in the transverse traceless gauge.

Appendix B

Regarding the use of discrete Fourier transforms as a sample of the continuous Fourier transform

This appendix concerns the use of continuous Fourier transforms to approximate the outcome of a discrete Fourier transform, and shows that the sinc approximation would be the same whether it is calculated using discrete or continuous methods. We follow a similar method to that used in [66].

The discrete Fourier transform is an algorithm which translates a series of data into its inverse domain, i.e. a time series into the frequency domain or a displacement series into wave number space. The discrete Fourier transform is defined by

$$F_j := \sum_{n=0}^{N-1} f_n e^{-2i\pi jn/N}, \tag{B.0.1}$$

where F_j is the Fourier amplitude of the j^{th} frequency bin of the FFT, N is the total number of datapoints (such that $N = r\Delta t$) and f_n is the discretised function being transformed.

Using the sinc approximation, and only using the equivalent of the first part of equation 3.1.13[1] the function being transformed is

$$f_n = y_k \exp\left[i\phi_k + 2i\pi f_k t_n\right], \tag{B.0.2}$$

where t_n is the n^{th} time point, $t_n = n/r$, r is the sample rate, and N is the total number of datapoints, $N = r\Delta t$. This means that F_j is

$$F_j = y_k \exp\left[i\phi_k\right] \sum_{n=0}^{N-1} \exp\left[2i\pi \frac{n}{r} \left(f_k - f_j\right)\right], \tag{B.0.3}$$

where $f_j = \frac{j}{\Delta t}$ is the frequency of the j^{th} bin.

Fortunately, the summation here has a closed form, and can be rewritten as

$$F_j = y_k \exp\left[i\phi_k\right] \frac{1 - \exp\left[2i\pi \frac{1}{r} \left(f_k - f_j\right) N\right]}{1 - \exp\left[2i\pi \frac{1}{r} \left(f_k - f_j\right)\right]}. \tag{B.0.4}$$

By factoring out appropriate half angles in the numerator and denominator, one can modify this to a simpler form, here (so that the equations fit on the page) we use a symbolic substitution of $q = \pi \frac{1}{r} \left(f_k - f_j\right)$

$$F_j = y_k \exp\left[i\phi_k\right] \frac{\exp\left[iqN\right] \left\{\exp\left[-iqN\right] - \exp\left[iqN\right]\right\}}{\exp\left[iq\right] \left\{\exp\left[-iq\right] - \exp\left[iq\right]\right\}}. \tag{B.0.5}$$

Using Euler's formula, the terms in the braces can be rewritten as sinusoids,

$$F_j = y_k \exp\left[i\phi_k\right] \frac{\exp\left[i\pi \left(f_k - f_j\right) \Delta t\right] \sin\left[\pi \left(f_k - f_j\right) \Delta t\right]}{\exp\left[i\pi \frac{1}{r} \left(f_k - f_j\right)\right] \sin\left[\pi \frac{1}{r} \left(f_k - f_j\right)\right]}. \tag{B.0.6}$$

If we make the assumption that r is large, we can use the small angle approximation for the sinusoid in the denominator, and assume that the exponent in the denominator is zero

$$F_j = y_k \exp\left[i\phi_k + i\pi \left(f_k - f_j\right) \Delta t\right] \frac{\sin\left[\pi \left(f_k - f_j\right) \Delta t\right]}{\exp[0] \frac{1}{r} \pi \left(f_k - f_j\right)}. \tag{B.0.7}$$

Using the sinc function as defined before, $\operatorname{sinc}(x) := \frac{\sin(x)}{x}$, we rewrite as

$$F_j = y_k r \Delta t \exp\left[i\phi_k + i\pi \left(f_k - f_j\right) \Delta t\right] \operatorname{sinc}\left[\pi \left(f_k - f_j\right) \Delta t\right]. \tag{B.0.8}$$

---

[1] This is the same assumption we made in the continuous case, but doing so at an earlier stage of the calculation.

In section 3.1.1, we calculated $H(f)$ in the sinc approximation as (this is equation 3.1.16)

$$H_k(f) \approx y_k \Delta t \exp\left[i\phi_k + i\pi(f_k - f)\Delta t\right] \operatorname{sinc}\left[\pi(f_k - f)\Delta t\right], \tag{B.0.9}$$

and comparing the two, we see a spurious factor of r remaining in the discrete calculation which is not in the continuous form. This is expected, and within the LIGO collaboration, the `lalapps_MakeSFTs` code removes this factor.

Appendix C

Regarding the derivation of the B_k estimator

Some people in the review of this work for the collaboration have raised doubts over the legitimacy of taking the derivative of something with respect to its conjugate, and that this would be zero, as was done in section 3.2 to find an estimator for B_k from the least squares fit. In fact this is something that can be shown fairly easily. By using the Wirtinger derivatives with the complex variable $z = x + iy$

$$\frac{\mathrm{d}}{\mathrm{d}z} = \frac{1}{2}\left(\frac{\mathrm{d}}{\mathrm{d}x} - i\frac{\mathrm{d}}{\mathrm{d}y}\right) \text{ and } \frac{\mathrm{d}}{\mathrm{d}z^*} = \frac{1}{2}\left(\frac{\mathrm{d}}{\mathrm{d}x} + i\frac{\mathrm{d}}{\mathrm{d}y}\right), \tag{C.0.1}$$

we see that

$$\frac{\mathrm{d}z}{\mathrm{d}z^*} = \frac{1}{2}\left(\frac{\mathrm{d}}{\mathrm{d}x} + i\frac{\mathrm{d}}{\mathrm{d}y}\right)(x + iy) \tag{C.0.2}$$

$$= \frac{1}{2}\left(\frac{\mathrm{d}x}{\mathrm{d}x} + i\frac{\mathrm{d}x}{\mathrm{d}y} + i\frac{\mathrm{d}y}{\mathrm{d}x} + i^2\frac{\mathrm{d}y}{\mathrm{d}y}\right) \tag{C.0.3}$$

$$= \frac{1}{2}(1 + 0 + 0 - 1) = 0, \tag{C.0.4}$$

and equivalently

$$\frac{\mathrm{d}z^*}{\mathrm{d}z} = \frac{1}{2}\left(\frac{\mathrm{d}}{\mathrm{d}x} - i\frac{\mathrm{d}}{\mathrm{d}y}\right)(x - iy) \tag{C.0.5}$$

$$= \frac{1}{2}\left(\frac{\mathrm{d}x}{\mathrm{d}x} - i\frac{\mathrm{d}x}{\mathrm{d}y} - i\frac{\mathrm{d}y}{\mathrm{d}x} + i^2\frac{\mathrm{d}y}{\mathrm{d}y}\right) \tag{C.0.6}$$

$$= \frac{1}{2}(1 + 0 + 0 - 1) = 0. \tag{C.0.7}$$

In case the reader still has the same doubts, the derivation is repeated here separately for the real and imaginary parts of B_k and then recombined.

The least squares fit means that we must find the minimum value of $\sum_j |S_{kj} - B_k\mu_{kj}|^2$ given different B_k. To do this, we differentiate here with respect to each of the real and imaginary parts of B_k and set these to zero. Firstly, we expand the terms of the sum

$$\sum_j |S_{kj} - B_k\mu_{kj}|^2 = \sum_j S_{kj}S_{kj}^* - B_k^*\mu_{kj}^*S_{kj} - B_k\mu_{kj}S_{kj}^* + B_k^*B_k\mu_{kj}^*\mu_{kj}, \tag{C.0.8}$$

which, with the real and imaginary parts of B_k explicitly separated becomes

$$\sum_j |S_{kj} - B_k\mu_{kj}|^2 = \sum_j S_{kj}S_{kj}^* - \mathfrak{Re}(B_k)\left[\mu_{kj}^*S_{kj} + \mu_{kj}S_{kj}^*\right] + i\mathfrak{Im}(B_k)\left[\mu_{kj}^*S_{kj} - \mu_{kj}S_{kj}^*\right] + \left[\mathfrak{Re}(B_k)^2 + \mathfrak{Im}(B_k)^2\right]\mu_{kj}^*\mu_{kj}. \tag{C.0.9}$$

We then take the derivative with respect to the real part of B_k, $\mathfrak{Re}(B_k)$,

$$\frac{\mathrm{d}}{\mathrm{d}\mathfrak{Re}(B_k)}\sum_j |S_{kj} - B_k\mu_{kj}|^2 = \sum_j \{-\left[\mu_{kj}^*S_{kj} + \mu_{kj}S_{kj}^*\right] + 2\mathfrak{Re}(B_k)\mu_{kj}^*\mu_{kj}\} = 0, \tag{C.0.10}$$

which can be solved for $\mathfrak{Re}(B_k)$ as

$$\mathfrak{Re}(B_k) = \frac{\sum_j \left[\mu_{kj}^*S_{kj} + \mu_{kj}S_{kj}^*\right]}{2\sum_j \mu_{kj}^*\mu_{kj}}. \tag{C.0.11}$$

Next we take the derivative with respect to the imaginary part of B_k, $\mathfrak{Im}(B_k)$,

$$\frac{\mathrm{d}}{\mathrm{d}\mathfrak{Im}(B_k)} \sum_j |S_{kj} - B_k\mu_{kj}|^2 = \sum_j \{ i \left[\mu^*_{kj} S_{kj} - \mu_{kj} S^*_{kj} \right] + 2\mathfrak{Im}(B_k)\mu^*_{kj}\mu_{kj} \} = 0, \quad \text{(C.0.12)}$$

which can be solved for $\mathfrak{Im}(B_k)$ as

$$\mathfrak{Im}(B_k) = \frac{\sum_j \left[\mu^*_{kj} S_{kj} - \mu_{kj} S^*_{kj} \right]}{2i \sum_j \mu^*_{kj}\mu_{kj}}. \quad \text{(C.0.13)}$$

This placing of the imaginary unit in the denominator helps us to compare equations C.0.11 and C.0.13 to z, an arbitrary complex number, for which we can easily see

$$\mathfrak{Re}(z) = \frac{z + z^*}{2} \text{ and } \mathfrak{Im}(z) = \frac{z - z^*}{2i}, \quad \text{(C.0.14)}$$

so we can equivalently see from equations C.0.11 and C.0.13

$$B_k = \frac{\sum\limits_j [S_{kj}\mu^*_{kj}]}{\sum\limits_j [\mu^*_{kj}\mu_{kj}]}. \quad \text{(C.0.15)}$$

which is equation 3.2.9.

In this separated real and imaginary method, it is more intuitive than it was in the previous method to see that this stationary point will always be a minimum. If we take the next (second) derivative of $\sum_j |S_{kj} - B_k\mu_{kj}|^2 = 2\mu^*_{kj}\mu_{kj}$ with respect to $\mathfrak{Re}(B_k)$ and $\mathfrak{Im}(B_k)$, we find

$$\frac{\mathrm{d}^2}{\mathrm{d}\mathfrak{Re}(B_k)^2} \sum_j |S_{kj} - B_k\mu_{kj}|^2 = 2\mu^*_{kj}\mu_{kj} \quad \text{(C.0.16)}$$

$$\frac{\mathrm{d}^2}{\mathrm{d}\mathfrak{Im}(B_k)^2} \sum_j |S_{kj} - B_k\mu_{kj}|^2 = 2\mu^*_{kj}\mu_{kj}. \quad \text{(C.0.17)}$$

As $\mu^*_{kj}\mu_{kj}$ will always be positive ($\mu^*_{kj}\mu_{kj} = |\mu_{kj}|^2$), these stationary points will always be a minimum.

Appendix D

Additional results from the Spotlight outlier search

The Spotlight outlier plots in section 6.4 were all very similar, as such we presented a representative example at that time. For completeness we present the rest of the parameter estimation plots here in figures D.0.1 to D.0.10.

We plot the posterior probability densities for individual detector analyses from H1 and L1 and a joint analysis from H1 and L1. This is shown in the form of one dimensional pdfs and two dimensional contour plots for frequency against frequency first derivative $f^{(0)}$ vs. $f^{(1)}$ and right ascension against declination α vs. δ. We plot $f^{(0)} - f^{(0)}_{\mathrm{cand}}$ for aesthetic reasons, as the numerical precision required would not be representable on the figure.

The joint analysis for candidate B30 was unfortunately cut off, without time for re-analysis, as such figure D.0.10 contains an *incoherent* joint posterior on the parameters. This difference has been visually highlighted by plotting the incoherent joint posterior in magenta, rather than green as used in the other figures. The incoherent posterior is found by multiplying the two posterior pdfs together, effectively using one analysis as the prior on the other. We must be careful about applying the priors twice though, as applying a Gaussian prior twice will effectively reduce the width of the Gaussian. In consideration of the priors, we divide each point on the combined posterior by the prior value at that point.

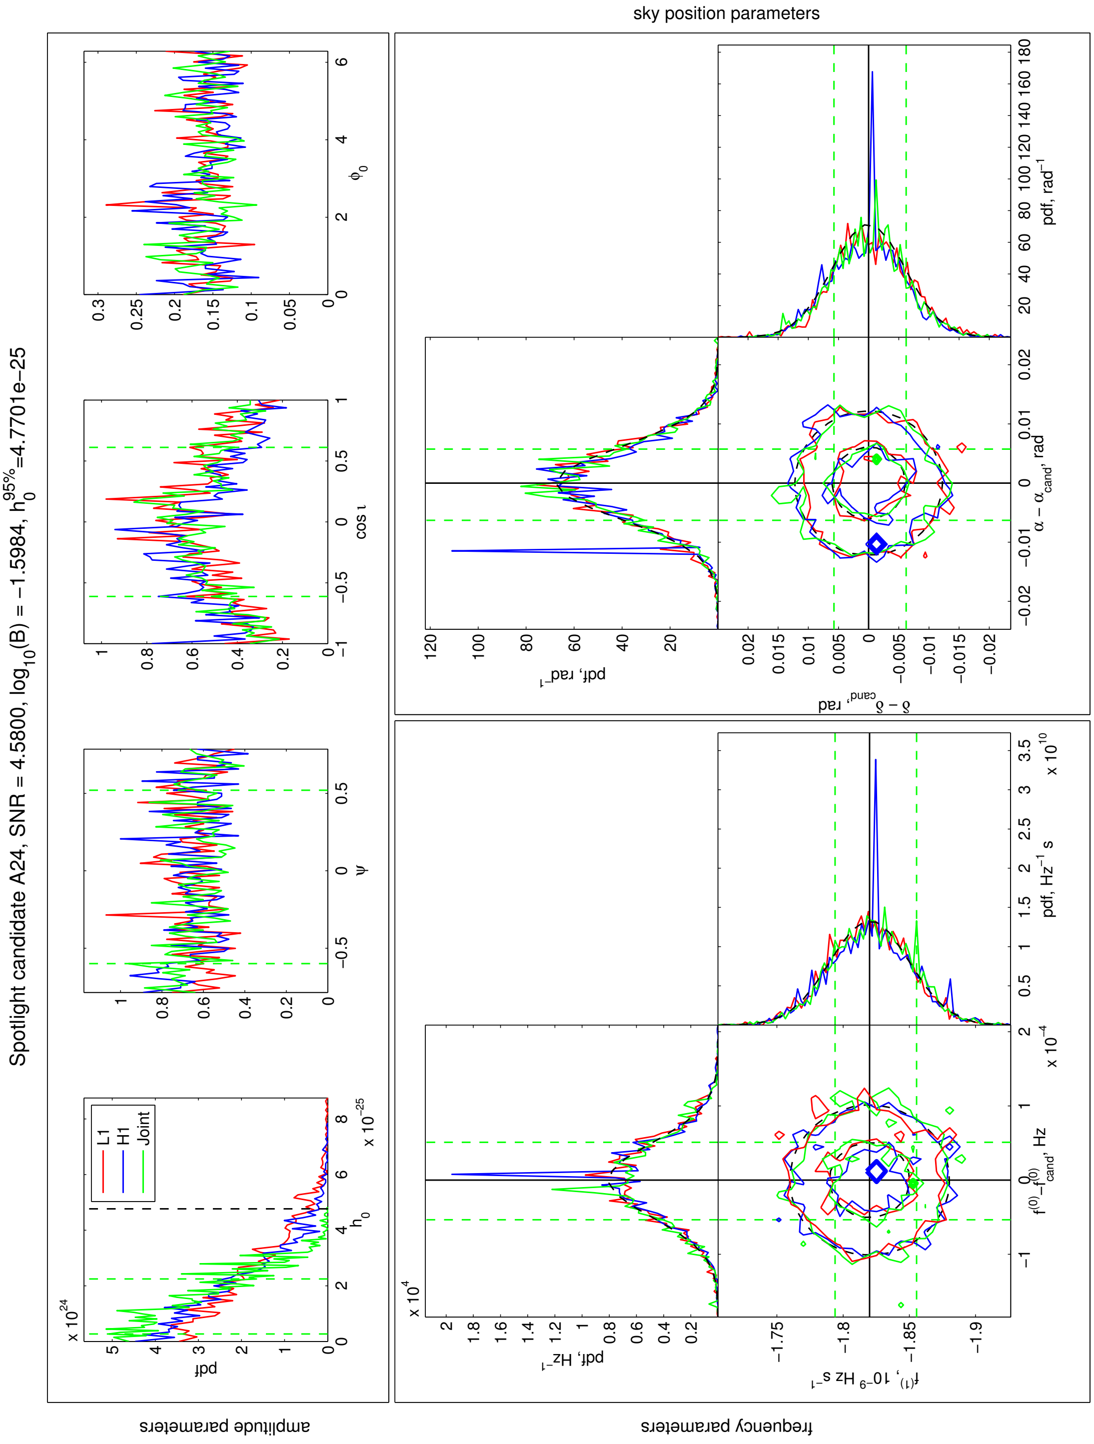

Figure D.0.1: Posterior probability density distribution and contour plots for parameters of Spotlight candidate A24, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\rm cand} = 1321.567031\,\mathrm{Hz}$.

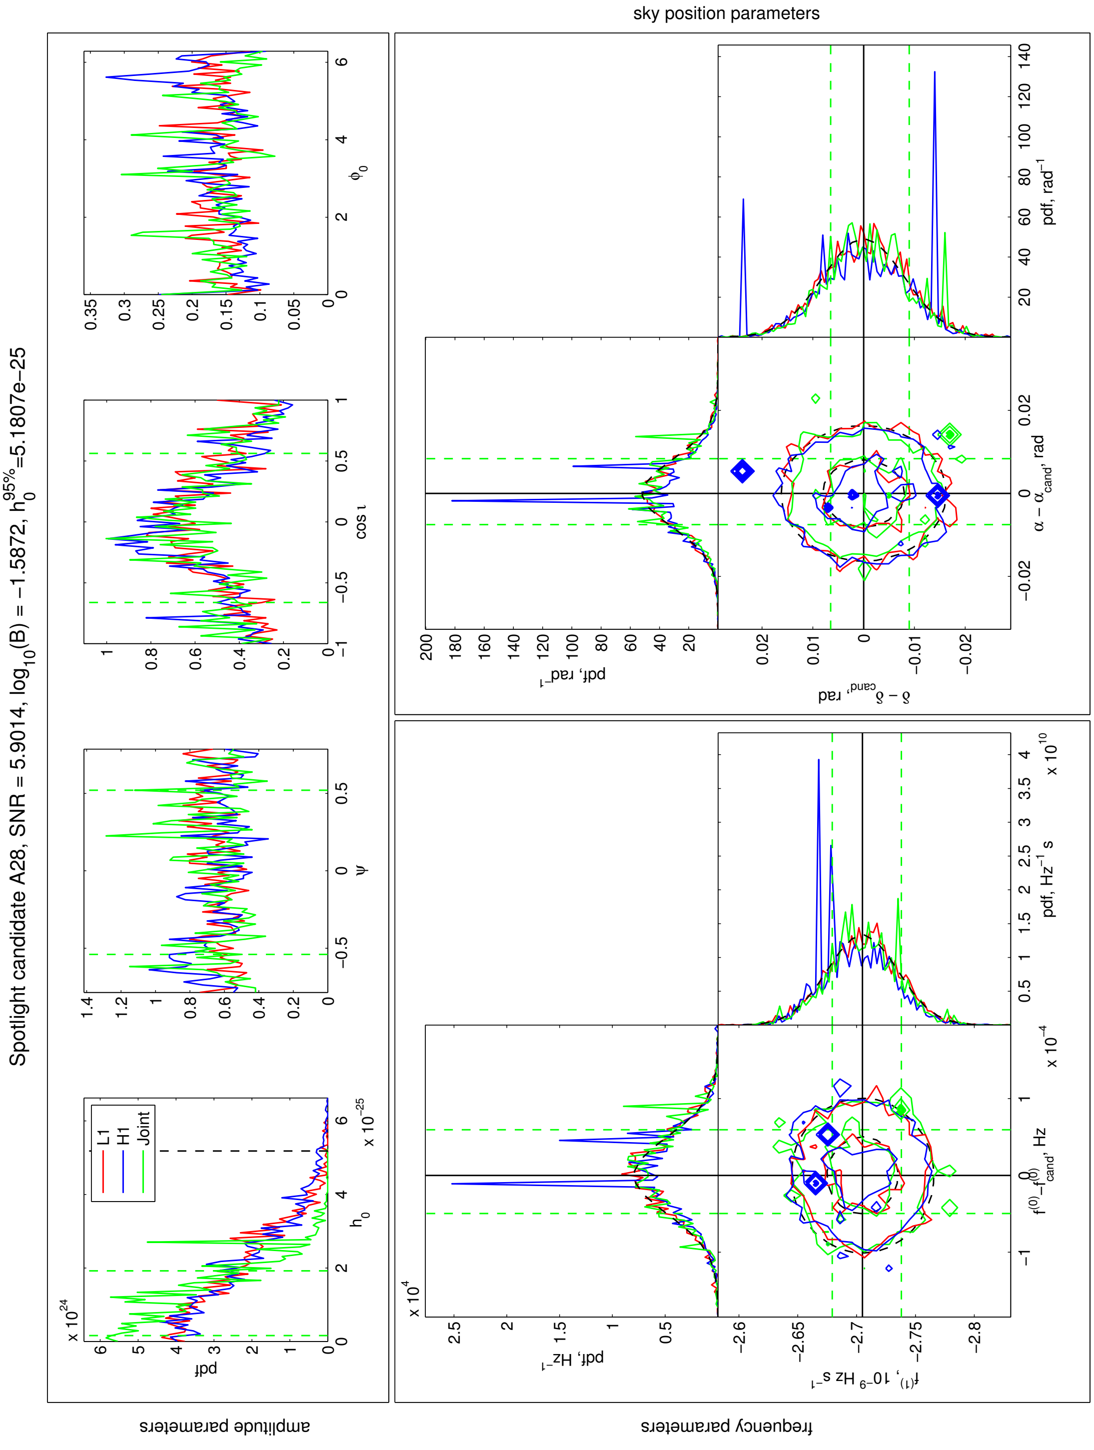

Figure D.0.2: Posterior probability density distribution and contour plots for parameters of Spotlight candidate A28, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\mathrm{cand}} = 990.761302\,\mathrm{Hz}$.

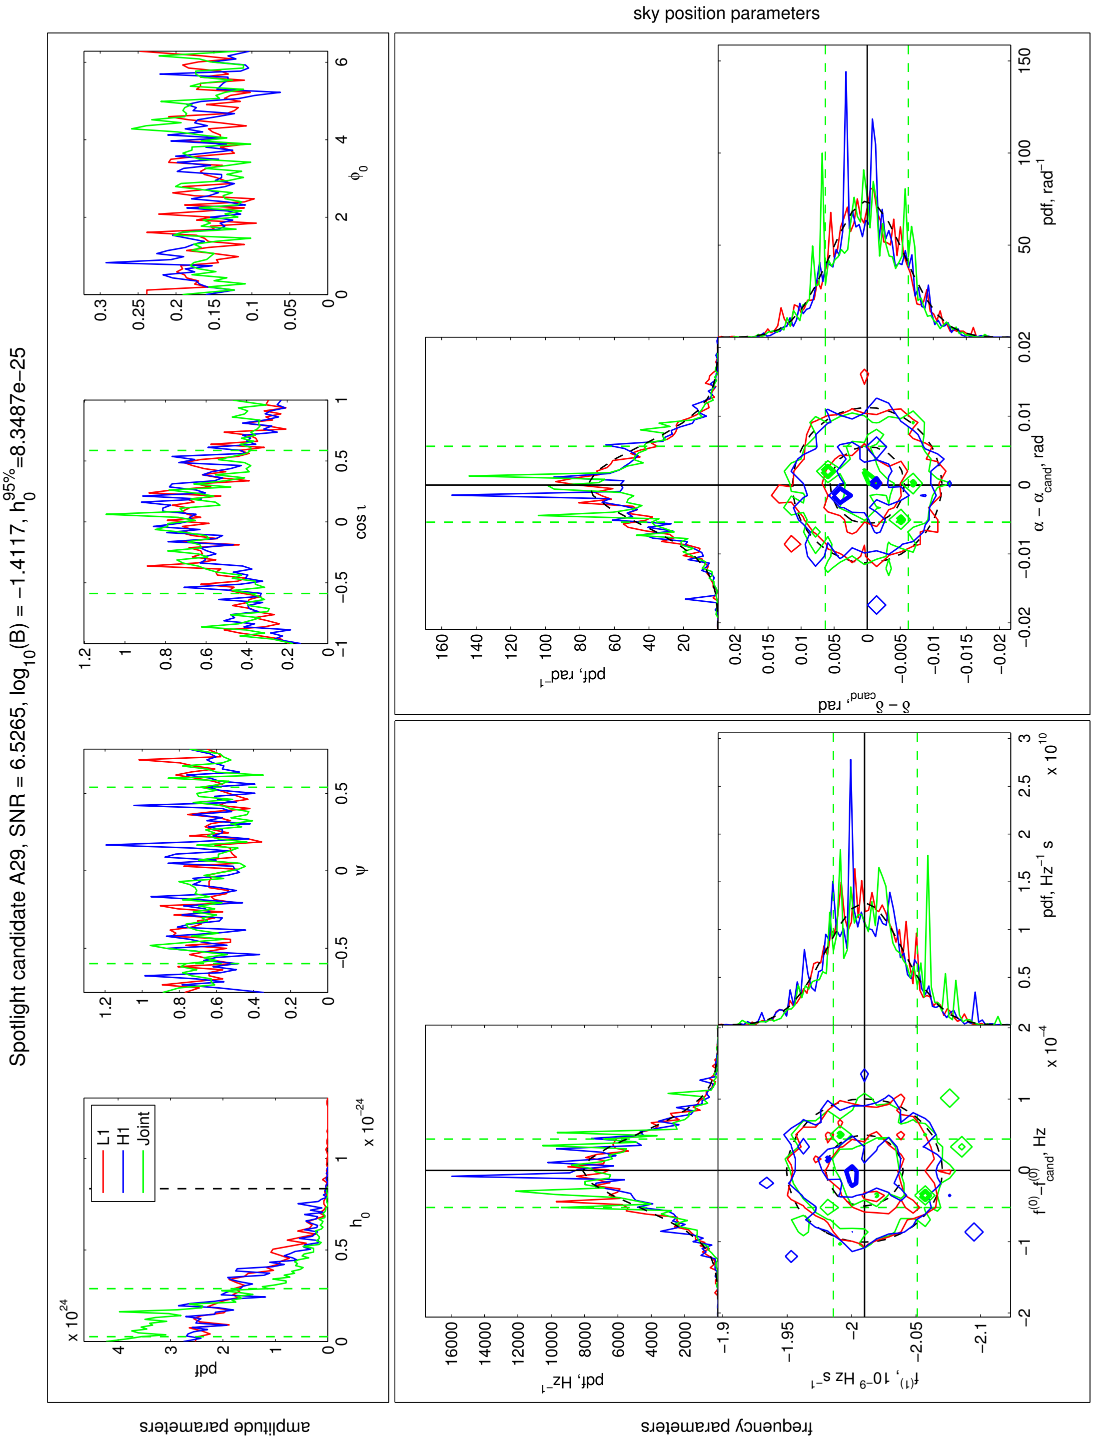

Figure D.0.3: Posterior probability density distribution and contour plots for parameters of Spotlight candidate A29, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\mathrm{cand}} = 1429.678924\,\mathrm{Hz}$.

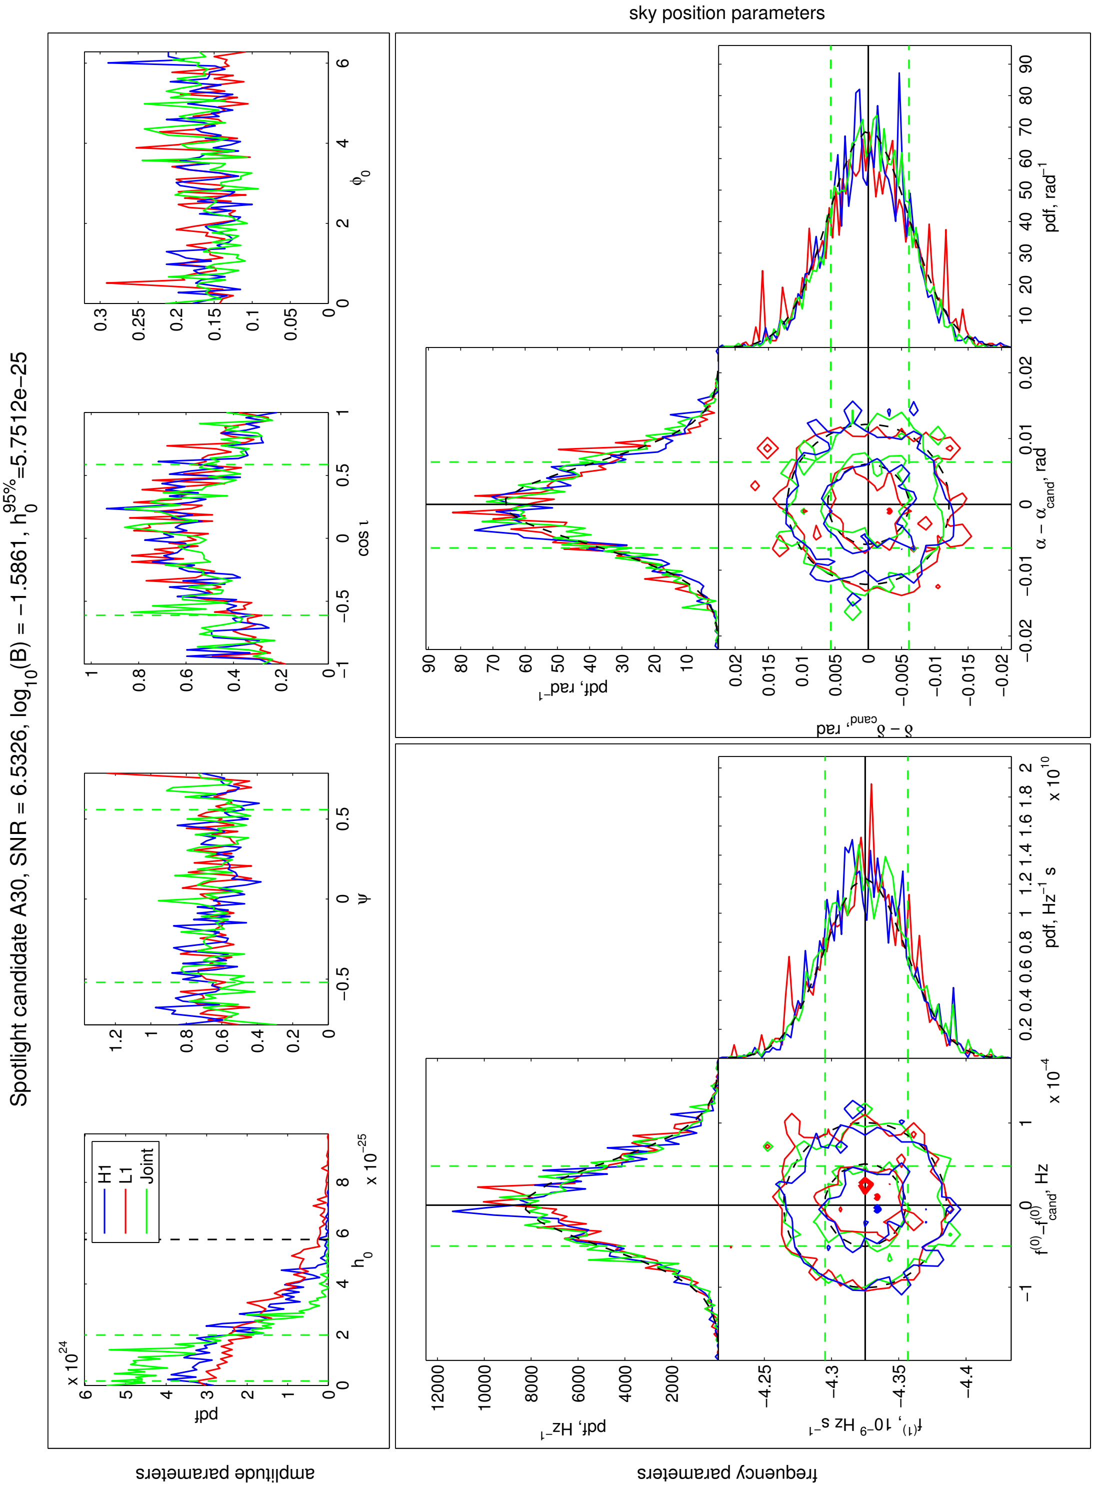

Figure D.0.4: Posterior probability density distribution and contour plots for parameters of Spotlight candidate A30, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\rm cand} = 1325.509688\,{\rm Hz}$.

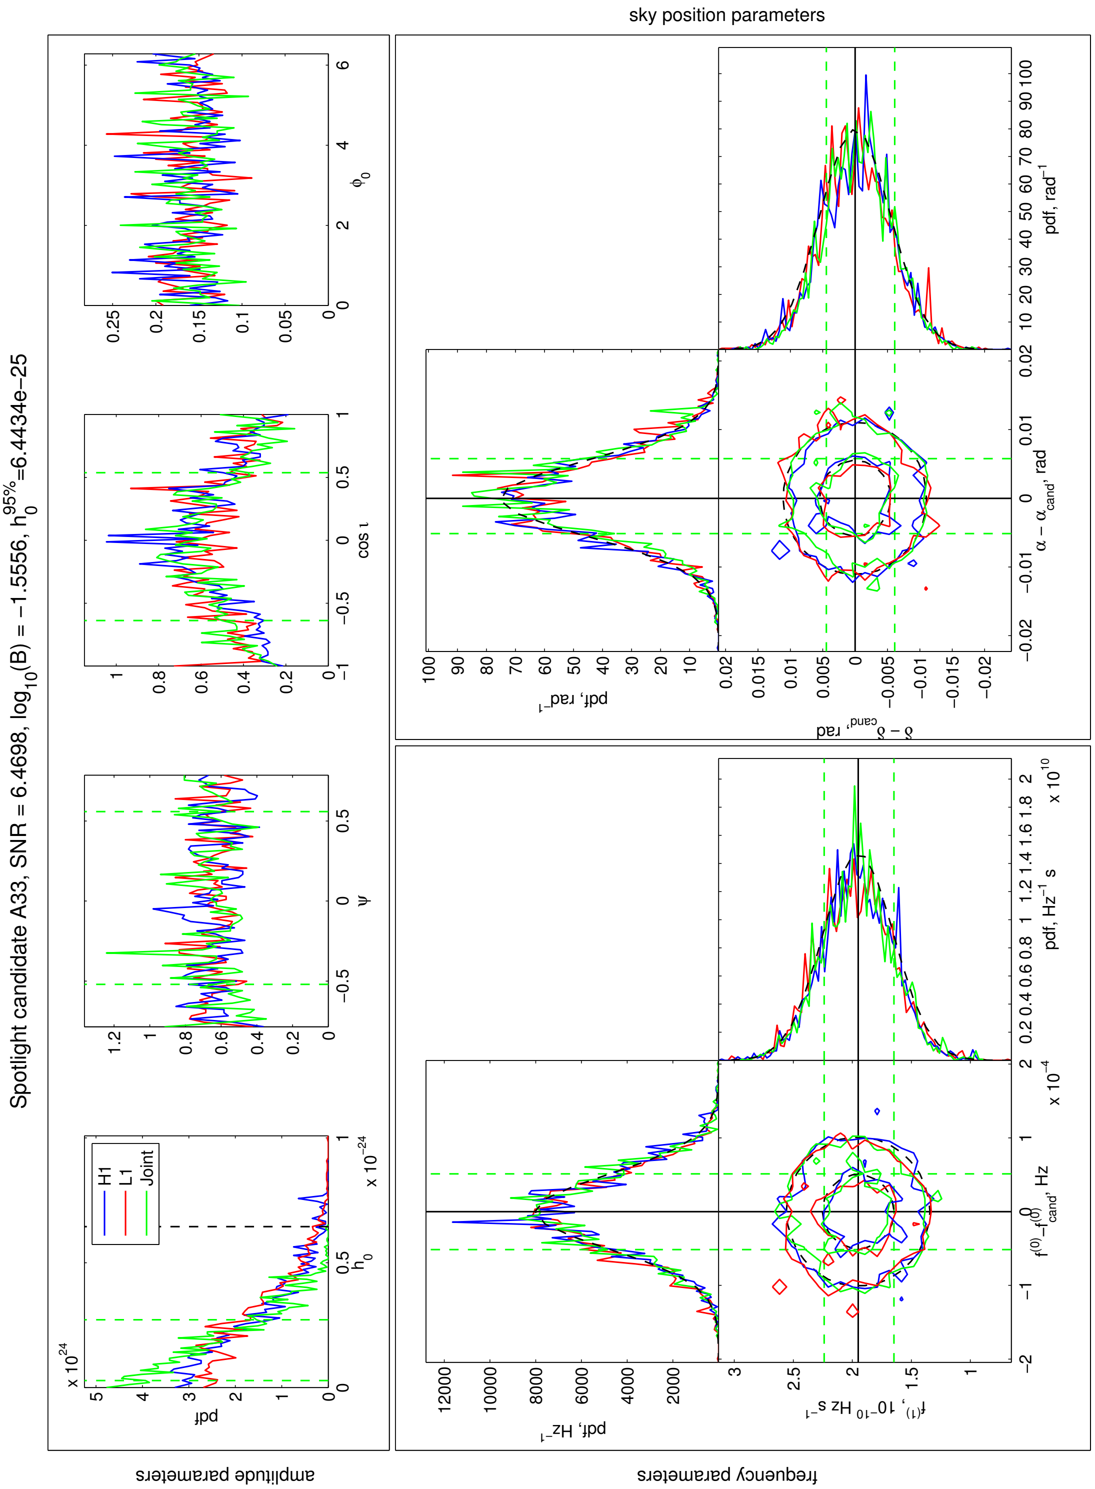

Figure D.0.5: Posterior probability density distribution and contour plots for parameters of Spotlight candidate A33, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\mathrm{cand}} = 1456.266111\,\mathrm{Hz}$.

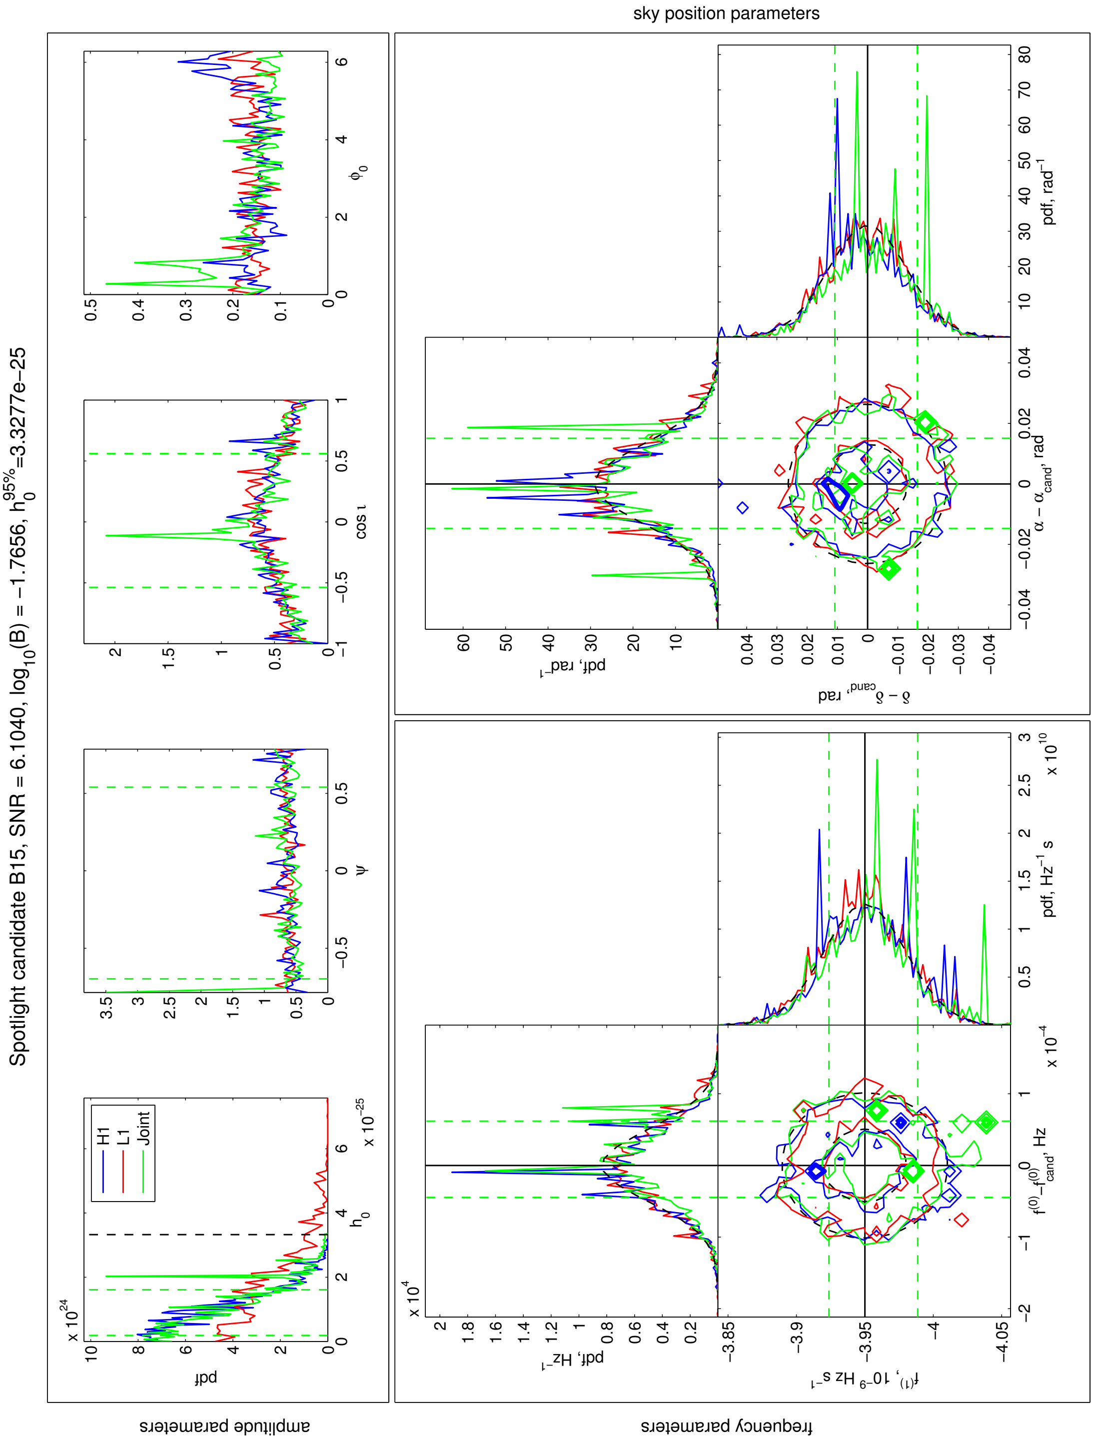

Figure D.0.6: Posterior probability density distribution and contour plots for parameters of Spotlight candidate B15, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\mathrm{cand}} = 613.261319\,\mathrm{Hz}$.

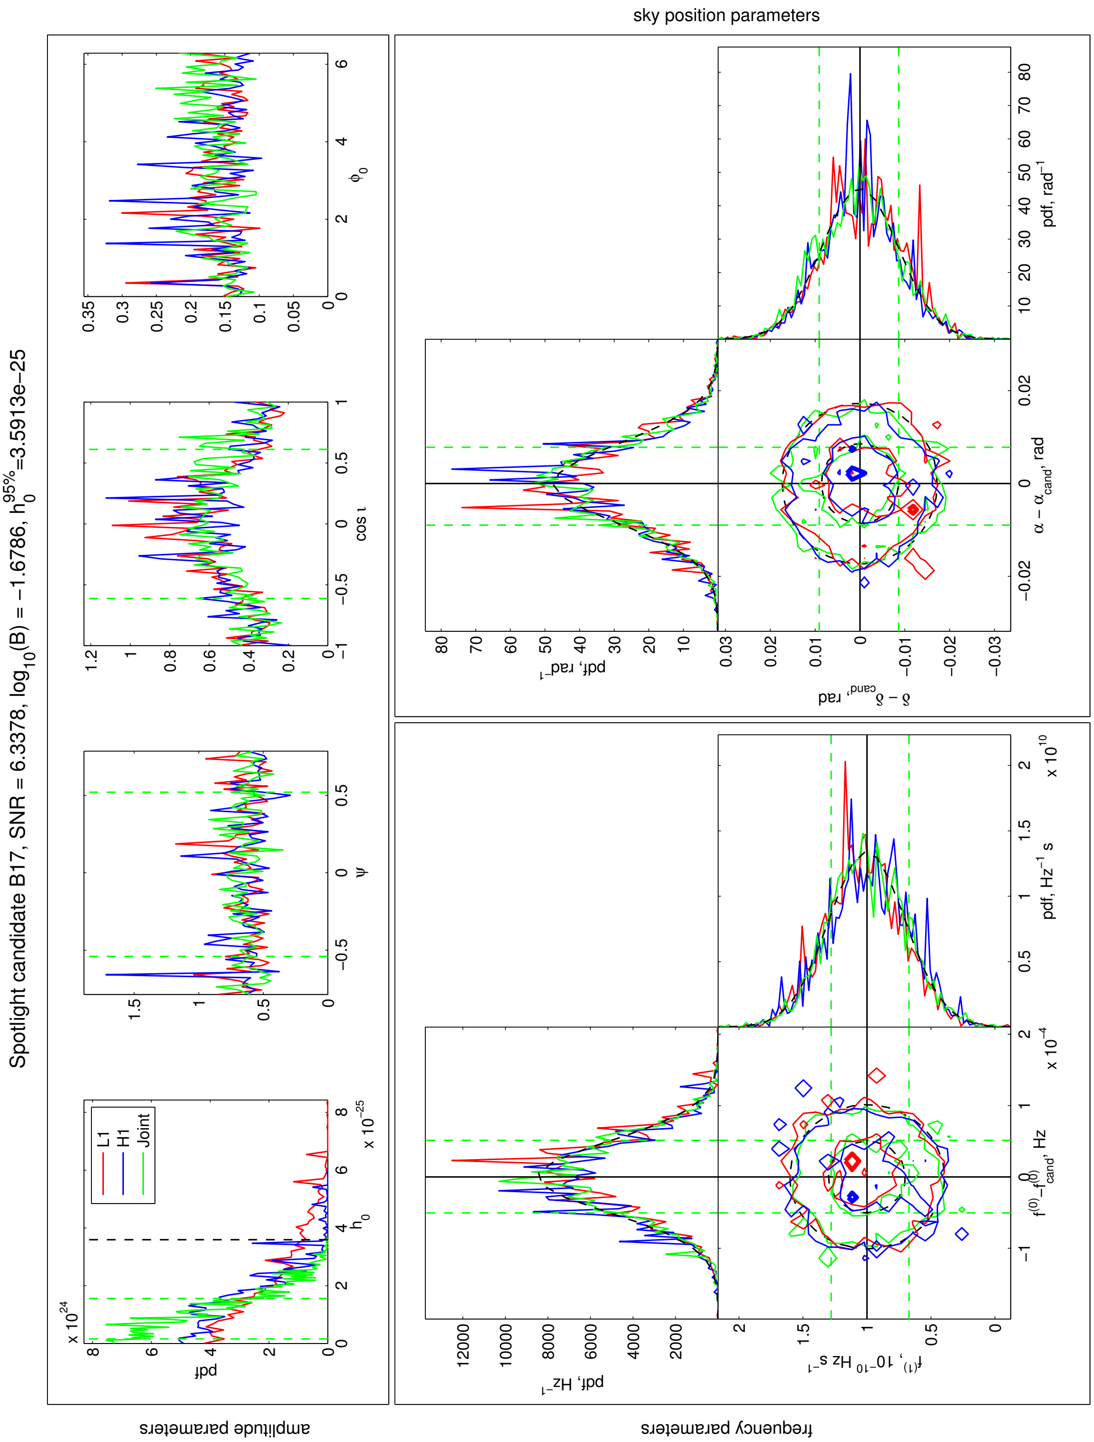

Figure D.0.7: Posterior probability density distribution and contour plots for parameters of Spotlight candidate B17, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\mathrm{cand}} = 933.338229\,\mathrm{Hz}$.

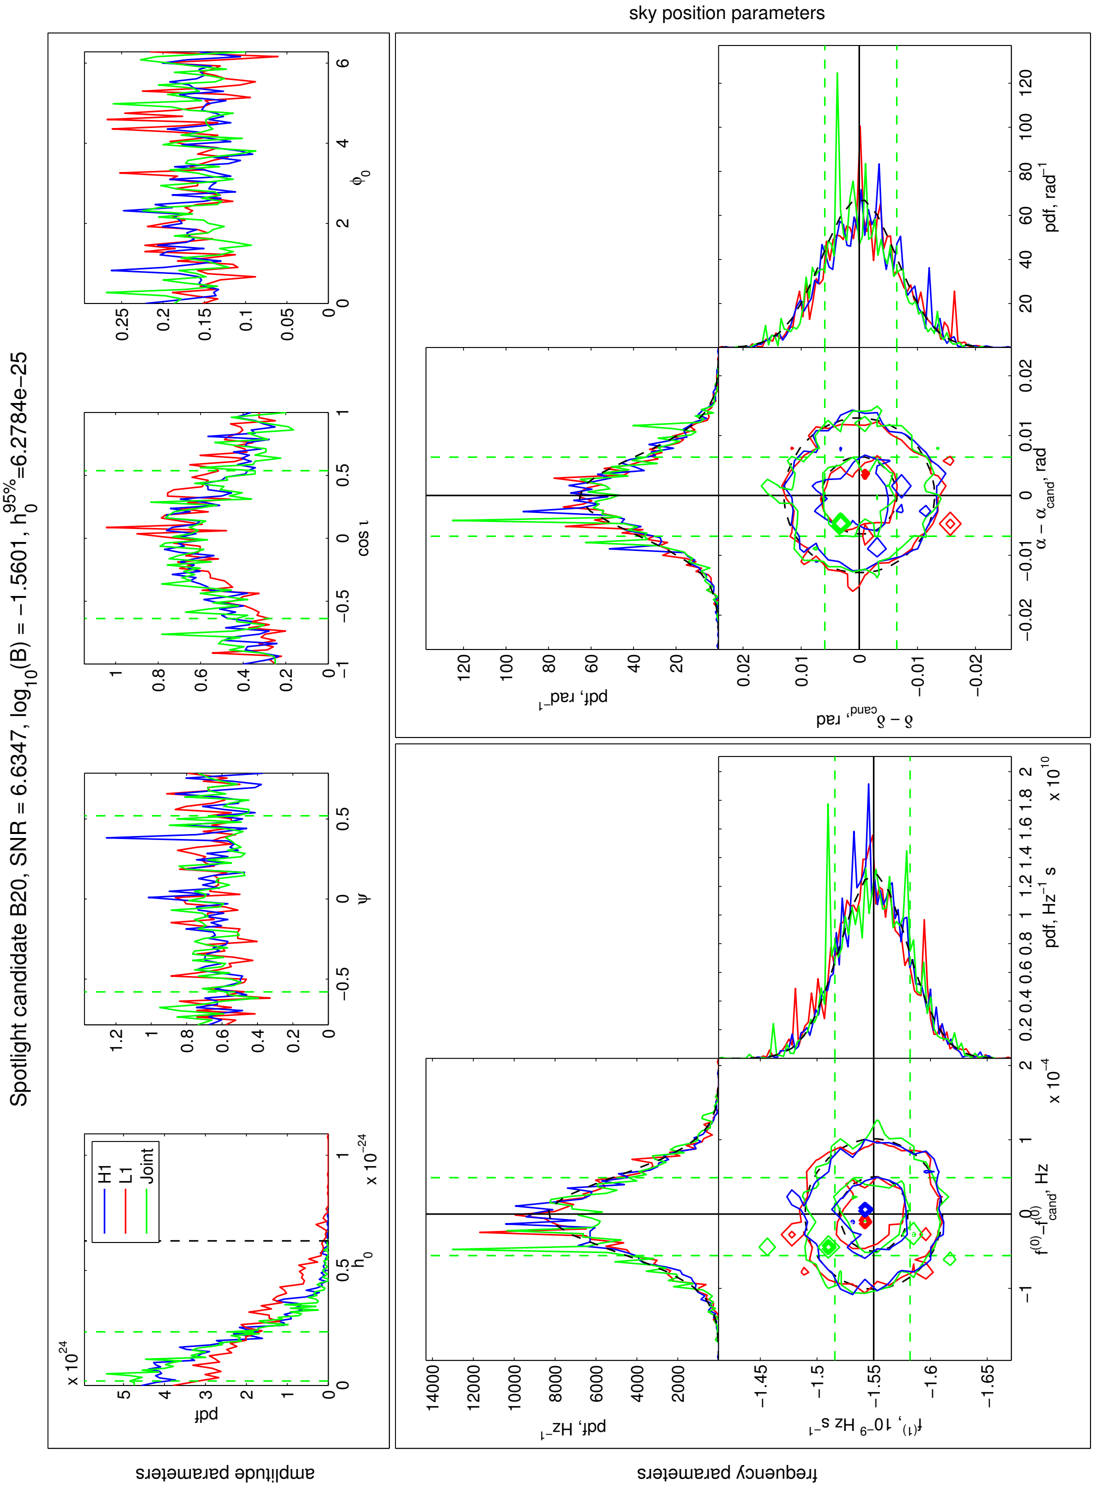

Figure D.0.8: Posterior probability density distribution and contour plots for parameters of Spotlight candidate B20, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\text{cand}} = 1249.438351\,\text{Hz}$.

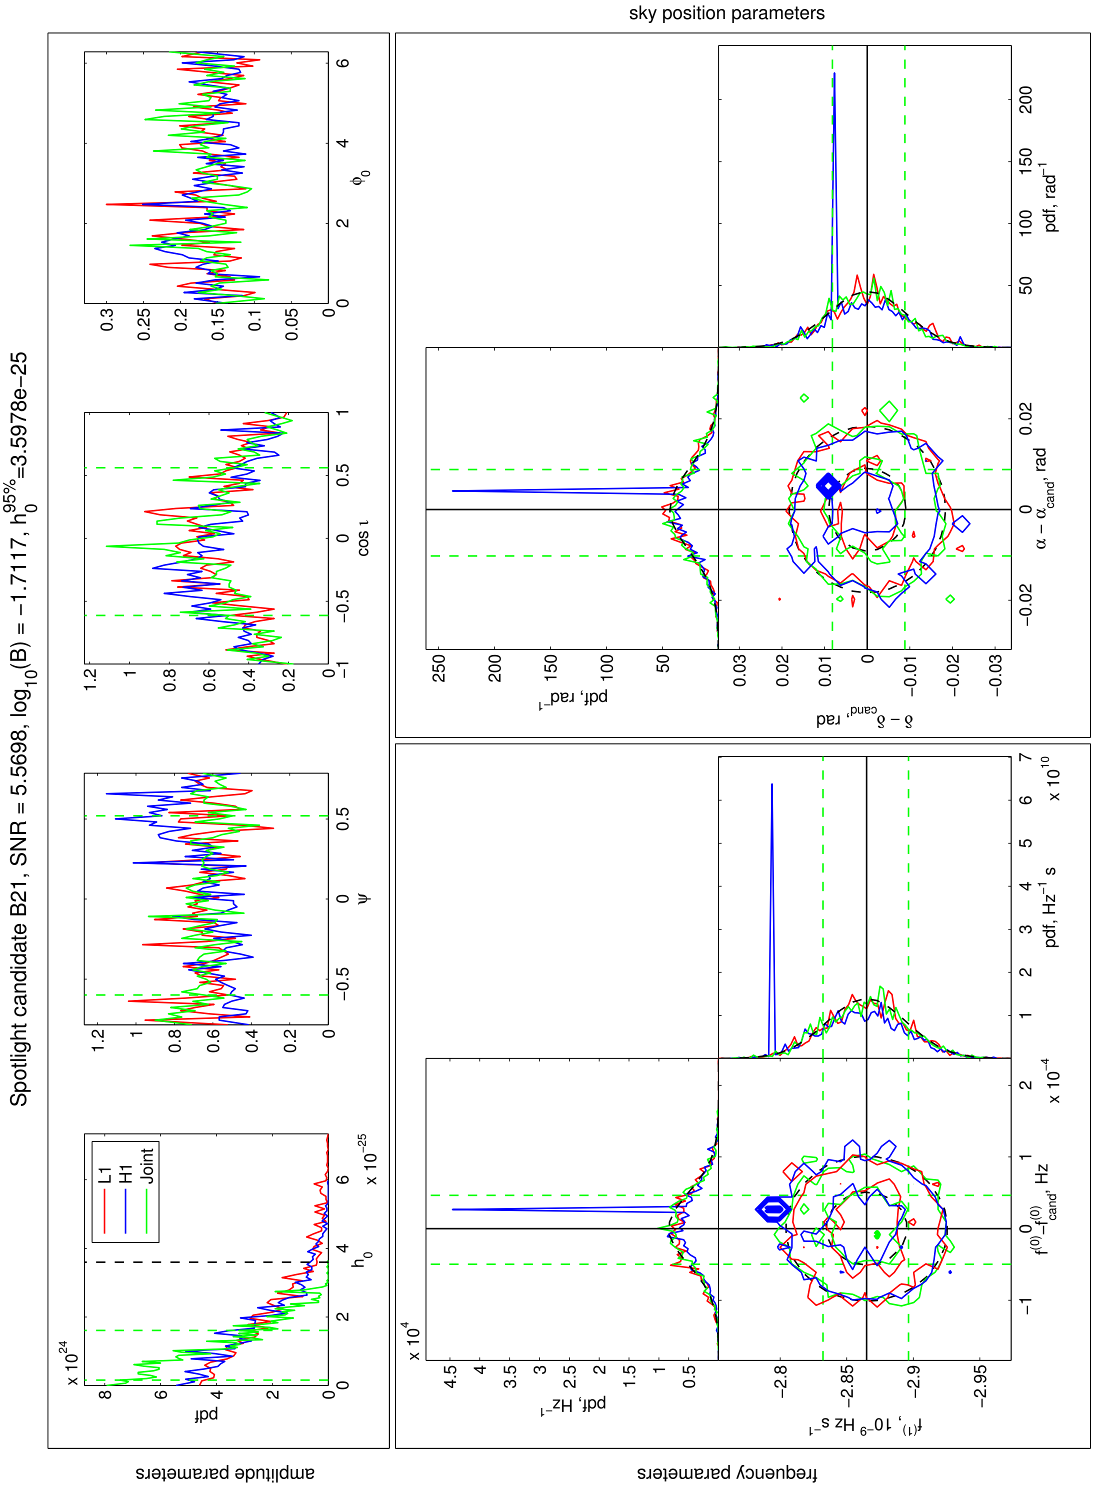

Figure D.0.9: Posterior probability density distribution and contour plots for parameters of Spotlight candidate B21, from individual analyses for H1 (blue) and L1 (red), and a joint analysis (green). $f^{(0)}_{\text{cand}} = 880.401753\,\text{Hz}$.

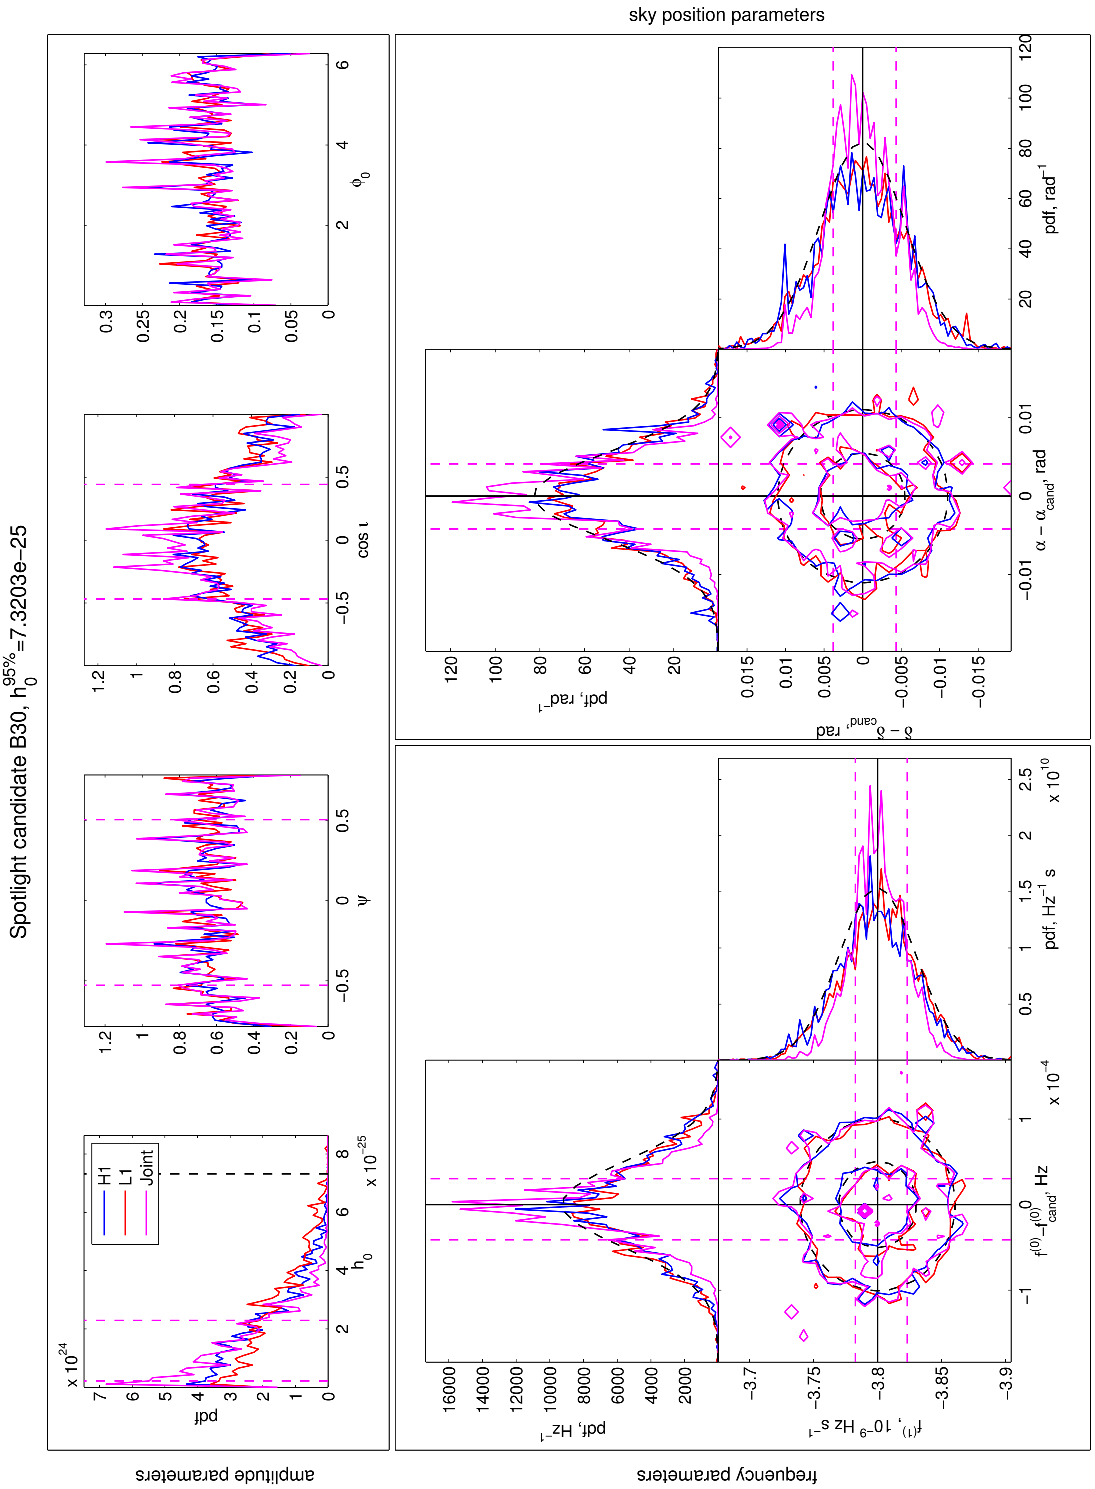

Figure D.0.10: Posterior probability density distribution and contour plots for parameters of Spotlight candidate B30, from individual analyses for H1 (blue) and L1 (red), and an incoherent joint posterior (magenta). $f^{(0)}_{\mathrm{cand}} = 1458.536476\,\mathrm{Hz}$.

Abbreviations and Symbols

Presented here is a list of definitions and abbreviations used throughout the text. All definitions and abbreviations are indicated at the point of their first use in the main body of the text, but are included here for reference.

Abbreviations

Abbreviation	Meaning
ANTF	Australia Telescope National Facility, the Commonwealth Scientific and Industrial Research Organisation's radio astronomy observatories, most often used to discuss the pulsar catalogue, [26, 67].
CBC	Compact Binary Coalescence. The inspiral and merging of two compact objects such as neutron stars or black holes.
DFT	Discrete Fourier Transform. An algorithm for transforming a discretely sampled series of data from the time domain to the frequency domain.
EM	Electromagnetic, often referring to optical, X-ray and γ-ray observations of astrophysical objects.
FFT	Fast Fourier Transform. An algorithm for fast computation of a discrete Fourier transform, also used to describe its output.
GR	General Relativity, Einstein's theory of.
GW	Gravitational Wave, a solution to the wave equation of General Relativity, and its effect on the spacetime metric.
H1	LIGO detector at Hanford, WA.

Continued

Abbreviation	Meaning
IFO	Interferometer.
L1	LIGO detector at Livingston, LA.
pdf	probability density function.
MCMC	Markov-Chain Monte Carlo. A method for sampling a probability distribution, used mostly in this thesis to describe the use of MCMC to explore posterior probability distribution parameter space.
SFT	Short Fourier Transform. A specific form of FFT, used in the LIGO and Virgo Scientific collaborations for storing frequency-domain data.
S5	Science run five. A time period (4th Nov 05, 4pm UTC to 30th Sep 07, 0:00am UTC) for which the LIGO detectors were collecting scientifically useful data from the detectors.
S6	Science run six. A time period (7th Jul 09, 9pm UTC to 20th Oct 10, 3:04pm UTC) for which the LIGO detectors were collecting scientifically useful data from the detectors.
SplInter	Spectral Interpolation, a method of downsampling gravitational wave data for a prospective continuous-wave source, removing rotation and relative motion effects.
SSB	Solar System Barycentre, an inertial frame of reference at the centre of mass of the solar system. Generally very close to, but not neccessarily at, the centre of the Sun.
V1	Virgo detector, at Cascina, Tuscany.
VSR2	Virgo Science Run two. A time period (7th Jul 09, 9pm UTC to 22nd Oct 09, 7:36am UTC) for which the Virgo detector was collecting data at or close to design sensitivity.

Symbols

This table presents a list of commonly used symbols throughout the text. Other symbols are used, but their scope is limited to use in and around the point of first definition, so are not included here. Some symbols are utilised for more than one meaning in different chapters, though their different meanings are defined at that point and usually separated by different chapters in the text.

Symbol	Meaning
A_{ab}	The GW amplitude tensor, equation 1.1.8.
$A(t)$	The complex amplitude of a gravitational wave at the detector, equation 2.3.2.
B_k	The output of the SplInter routine. Noisy estimate of complex amplitude of GW signal at the detector, y_k, $B_k = y_k + n_k$, equation 3.2.9.
B_K	The output of the heterodyne routine. Noisy estimate of complex amplitude of GW signal at the detector, y_K, $B_K = y_K + n_K$, equations 2.3.5 and 2.3.6.
B_κ	The output of the heterodyne routine with half hour data separation. Found by averaging B_K during the window associated with each B_k, equation 4.1.1.
$\mathcal{C}[w]$	Fresnel cosine integral, equation 3.1.18.
$F_{+/\times}$	Antenna functions for plus and cross polarisations respectively, equations 1.3.2 and 1.3.3.
$\mathcal{F}$	The $\mathcal{F}$ statistic, the maximum likelihood probability for a given set of Doppler parameters for a continuous-wave signal, see section 6.1.1.
$H(f)$	Fourier transform of $h(t)$, as in equation 3.1.1.
$H_k(f)$	Fourier transform of $h(t)$ during the time window $t_k - \frac{\Delta t}{2} < t < t_k + \frac{\Delta t}{2}$, equation 3.1.4.
$\mathfrak{Im}[z]$	Imaginary part of z, e.g. in 3.3.1.
$N(f)$	Fourier transform of the noise, $n(t)$, $N_k(f)$ is frequency-domain noise in the k^{th} SFT, introduced in section 3.1.
R^a_{bcd}	Riemann Curvature Tensor, equation 1.1.4.
R_{ab}	Ricci Tensor and similarly the Ricci Scalar $R = R^a_a$, defined in section 1.1.1.
$\mathfrak{Re}[z]$	Real part of z, e.g. in 3.3.1

Continued

Symbol	Meaning
$S(f)$	Fourier transformed data from the detector, $S(f) = \mathrm{FFT}[s(t)]$, defined in equation 3.1.1.
$S_k(f)$	Fourier transformed data, $S(f)$ in the k^{th} SFT.
$\mathcal{S}[w]$	Fresnel sine integral, equation 3.1.19.
T_{coh}	coherent observation time, used in Fomalhaut b search, defined in table 6.1.
Z	evidence, the probability of the data given the information I, $\mathrm{p}(\{B_k\}, I)$, defined in section 2.2.
$\mathbf{a}$	set of parameters $[h_0, \psi, \phi_0, \iota]$, define in section 2.2.
$f^{(n)}$	GW frequency, n^{th} derivative, defined in section 2.1.
f_k	instantaneous GW frequency at time t_k, equation 3.1.8.
$\dot{f}_k$	instantaneous GW frequency derivative at time t_k, calculated using equation 3.1.9.
$\ddot{f}_k$	instantaneous GW frequency second derivative at time t_k, see section 4.1.3.
$g_{\mu\nu}$	spacetime metric, equation 1.1.2.
$h_{\mu\nu}$	a small perturbation to the metric, equation 1.1.2.
$\bar{h}_{\mu\nu}$	trace reverse of the metric perturbation, equation 1.1.5.
$h(t)$	GW strain in detector, equation 2.3.1.
h_0	GW strain amplitude, equation 2.1.3.
$h_0^{95\%}$	95% upper limit on h_0, i.e. 95% of the probability for values of h_0 is below this point, equation 2.2.2.
k	subscript to denote a specific B_k , used for summation in the likelihood calculation. This subscript also used to indicate values or variables specifically related to this B_k value, such as t_k being the centre of the time window, $\mu_k(f)$ and $S_k(f)$ for model and data related to the k^{th} SFT. First used in equation 3.1.3.

Continued

Symbol	Meaning
K	subscript to denote a specific B_K or to indicate values or variables specifically related to this B_K value, such as t_K being the centre of the time window. Defined e.g in equations 2.3.5, 2.3.12.
m	mismatch between SplInter and heterodyne values, indication of drop in SNR, defined in equation 4.1.2.
$n(t)$	noise on $h(t)$, defined section 2.3.1.
$\mathrm{p}(A\|B)$	probability of A given B, equation 2.2.1.
r	sample rate, equation 2.3.9.
$s(t)$	signal in detector $= h(t) + n(t)$, defined section 2.3.1.
t_k	time associated with B_k = time halfway through window, defined section 2.3.1, similarly t_K and t_κ for the time windows for B_K and B_κ.
y_k	complex amplitude of the gravitational wave at the detector at time t_k, including initial phase component $y_k := A(t_k)\exp(i\phi_0)$, equation 2.3.8.
α	right ascension, defined in section 2.1.
δ	declination, defined in section 2.1.
$\delta(t)$	time difference from SSB to detector, equation 2.1.7.
Δt	$t_{k/K/\kappa}$ spacing, the length of the time window during which each $B_{k/K/\kappa}$ is calculated.
ε	ellipticity, equation 2.1.4.
$\eta_{\nu\mu}$	flat spacetime metric, equation 1.1.2.
ι	source inclination angle, defined in section 2.1.
λ	Doppler parameters of a continuous-wave signal, defined in section 2.1.
κ	subscript to denote use of the heterodyne routine with half hour t_κ separation, equation 4.1.1.

Continued

Symbol	Meaning
$\mu(f)$	model of the expected GW signal in the frequency domain, amplitude removed $\mu(f) = \frac{H(f)}{y_k}$. $\mu_k(f)$ denotes $\mu(f)$ in the k^{th} SFT. Defined in equation 3.2.1.
$\pi(\mathbf{a})$	the prior $\mathrm{p}(\mathbf{a}\|I)$, as in section 2.2.
ρ	signal to noise ratio, SNR, defined in section 2.2.
σ_k	standard deviation on values of B_k, also used as SplInter estimate of this value, see section 3.2, first used equation 2.3.9.
σ_T	standard deviation of $n(t)$, the noise in the time domain 2.3.9.
σ_F	standard deviation of noise in the frequency domain, $N(f)$, defined in equation 3.2.2.
$\phi(t)$	GW phase $\Phi(t)$ minus the constant ϕ_0, equation 2.3.3.
ϕ_0	the constant offset of GW phase, defined in section 2.1.
$\Phi(t, \lambda)$	GW phase, a function of time and the Doppler parameters, equation 2.1.8.
ψ	GW polarisation angle, defined in section 2.1.

Bibliography

[1] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, et al. Directed search for continuous gravitational waves from the Galactic center. *Phys. Rev. D*, 88(10):102002, November 2013. doi:10.1103/PhysRevD.88.102002.

[2] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, et al. Gravitational Waves from Known Pulsars: Results from the Initial Detector Era. *Astrophys. J.*, 785:119, April 2014. doi:10.1088/0004-637X/785/2/119.

[3] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys. Rev. D*, 87(4):042001, February 2013. doi:10.1103/PhysRevD.87.042001.

[4] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, et al. Searches for continuous gravitational waves from nine young supernova remnants. *ArXiv e-prints, submitted to ApJ*, December 2014.

[5] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, et al. Advanced LIGO. *Class. Quant .Grav.*, 32(7):074001, April 2015. doi:10.1088/0264-9381/32/7/074001.

[6] J. Aasi, B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, et al. A loosely coherent search of the Orion spur using sixth science run LIGO data. 2015. In the LIGO Document Control Center `https://dcc.ligo.org/P1500034`.

[7] J. Aasi, B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, et al. Search for Gravitational Waves From Nearby Globular Clusters. 2016. In Preparation, see also presentation at `https://dcc.ligo.org/LIGO-G1500006`.

[8] J. Abadie, B. Abbott, R. Abbott, M. Abernathy, C. Adams, et al. First search for gravitational waves from the youngest known neutron star. *The Astrophysical Journal*, 722(2):1504, 2010.

[9] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, et al. Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys. Rev. D*, 85(8):082002, April 2012. doi:10.1103/PhysRevD.85.082002.

[10] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, et al. TOPICAL REVIEW: Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(17):173001, September 2010. doi:10.1088/0264-9381/27/17/173001.

[11] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, et al. Beating the Spin-down Limit on Gravitational Wave Emission from the Vela Pulsar. *Astrophys. J.*, 737:93, August 2011. doi:10.1088/0004-637X/737/2/93.

[12] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, et al. Directional Limits on Persistent Gravitational Waves Using LIGO S5 Science Data. *Phys. Rev. Lett.*, 107:271102, Dec 2011. doi:10.1103/PhysRevLett.107.271102.

[13] B. Abbott, R. Abbott, R. Adhikari, A. Ageev, B. Allen, et al. Setting upper limits on the strength of periodic gravitational waves from PSR J1939+2134 using the first science data from the GEO 600 and LIGO detectors. *Phys. Rev D*, 69(8):082004, April 2004. doi:10.1103/PhysRevD.69.082004.

[14] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, et al. Search for gravitational-wave bursts in ligo data from the fourth science run. *Classical and Quantum Gravity*, 24(22):5343, 2007.

[15] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, et al. All-sky search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 77(2):022001, January 2008. doi:10.1103/PhysRevD.77.022001.

[16] B. P. Abbott, R. Abbott, F. Acernese, R. Adhikari, P. Ajith, et al. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, 460:990–994, August 2009. doi:10.1038/nature08278.

[17] B. P. Abbott, R. Abbott, F. Acernese, R. Adhikari, P. Ajith, et al. Searches for Gravitational Waves from Known Pulsars with Science Run 5 LIGO Data. *Astrophys. J.*, 713:671–685, April 2010. doi:10.1088/0004-637X/713/1/671.

[18] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, et al. Search for gravitational waves from low mass binary coalescences in the first year of LIGO's S5 data. *Phys. Rev. D*, 79:122001, Jun 2009. doi:10.1103/PhysRevD.79.122001.

[19] M. Abernathy, F. Acernese, P. Ajith, B. Allen, P. Amaro-Seoane, et al. Einstein gravitational wave Telescope conceptual design study, 2011. ET Science Team [ET-0106C-10] `https://tds.ego-gw.it/ql/?c=7954`, cited: 28 Oct 2014.

[20] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, January 2015. doi:10.1088/0264-9381/32/2/024001.

[21] R. Adhikari, S. Ballmer, E. Campagna, M. Evans, P. Fritschel, et al. Gravitational Wave Interferometer Noise Calculator. `awiki.ligo-wa.caltech.edu/aLIGO/GWINC` [accessed 22 Jan 2015].

[22] O. Aguiar, L. Andrade, J. Barroso, F. Bortoli, L. Carneiro, et al. The Brazilian gravitational wave detector Mario Schenberg: Status report. *Class.Quant.Grav.*, 23:S239–S244, 2006. doi:10.1088/0264-9381/23/8/S30.

[23] B. Allen. Gravitational Wave Detector Sites. *ArXiv General Relativity and Quantum Cosmology e-prints*, July 1996.

[24] P. Astone, L. Baggio, M. Bassan, M. Bignotto, M. Bonaldi, et al. IGEC2: A 17-month search for gravitational wave bursts in 2005-2007. *Phys. Rev. D*, 82(2):022003, July 2010. doi:10.1103/PhysRevD.82.022003.

[25] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba. Coherent search of continuous gravitational wave signals: extension of the 5-vectors method to a network of detectors. *Journal of Physics Conference Series*, 363(1):012038, June 2012. doi:10.1088/1742-6596/363/1/012038.

[26] Australia Telescope National Facility. The ATNF pulsar database. `www.atnf.csiro.au/people/pulsar/psrcat/`.

[27] S. Braccini, C. Bradaschia, R. Del Fabbro, A. Di Virgilio, I. Ferrante, et al. Seismic vibrations mechanical filters for the gravitational waves detector virgo. *Review of Scientific Instruments*, 67(8):2899–2902, 1996. doi:http://dx.doi.org/10.1063/1.1147069.

[28] P. R. Brady and T. Creighton. Searching for periodic sources with LIGO. II. Hierarchical searches. *Phys. Rev. D*, 61:082001, Feb 2000. doi:10.1103/PhysRevD.61.082001.

[29] California Institute of Technology. "LIGO Laboratory Home Page", project homepage . `www.ligo.caltech.edu/`.

[30] N. Christensen, R. J. Dupuis, G. Woan, and R. Meyer. Metropolis-Hastings algorithm for extracting periodic gravitational wave signals from laser interferometric detector data. *Phys. Rev. D*, 70(2):022001, July 2004. doi:10.1103/PhysRevD.70.022001.

[31] N. Copernicus. *De revolutionibus orbium coelestium.* Johannes Petreius, Nuremberg, Holy Roman Empire, 1543.

[32] J. Creighton, B. Allen, W. G. Anderson, D. Ausmus, R. Balasubramanian, et al. LALSoftware Documentation, 2004. `http://www.lsc-group.phys.uwm.edu/lal/lsd.pdf`.

[33] K. Danzmann and the LISA study team. LISA: Laser Interferometer Space Antenna for gravitational wave measurements. *Classical and Quantum Gravity*, 13(11A):A247, 1996.

[34] G. S. Davies, M. Pitkin, and G. Woan. A targeted spectral interpolation algorithm for the detection of continuous gravitational waves. 2015. In Preparation.

[35] A. de Waard, L. Gottardi, J. van Houwelingen, A. Shumack, and G. Frossati. MiniGRAIL, the first spherical detector. *Classical and Quantum Gravity*, 20:143, May 2003. doi:10.1088/0264-9381/20/10/317.

[36] V. Dergachev. On blind searches for noise dominated signals: a loosely coherent approach. *Classical and Quantum Gravity*, 27(20):205017, October 2010. doi: 10.1088/0264-9381/27/20/205017.

[37] V. Dergachev. Private Communication, 2015.

[38] R. Dupuis. *Bayesian searches for gravitational waves from pulsars.* Ph.D. thesis, University of Glasgow, 2004.

[39] R. J. Dupuis and G. Woan. Bayesian estimation of pulsar parameters from gravitational wave data. *Phys. Rev. D*, 72(10):102002, November 2005. doi: 10.1103/PhysRevD.72.102002.

[40] A. Einstein. Approximative Integration of the Field Equations of Gravitation. *Königlich Preußische Akademie der Wissenschaften: Sitzungberichte*, pages 688–692, 1916. Conference Proceedings of Prussian Academy of Sciences. Available online in *Volume 6: The Berlin Years: Writings, 1914-1917 (English translation supplement)* Translated by Alfred Engel `http://einsteinpapers.press.princeton.edu/vol6-trans`.

[41] R. Essick, S. Vitale, E. Katsavounidis, G. Vedovato, and S. Klimenko. Localization of Short Duration Gravitational-wave Transients with the Early Advanced LIGO and Virgo Detectors. *Astrophys. J.*, 800:81, February 2015. doi: 10.1088/0004-637X/800/2/81.

[42] European Space Agency. “Next steps for LISA”, NGO project homepage. `http://sci.esa.int/science-e/www/object/index.cfm?fobjectid=48728`, 2012.

[43] S. Fairhurst. Source localization with an advanced gravitational wave detector network. *Classical and Quantum Gravity*, 28(10):105021, 2011.

[44] F. Feroz, M. P. Hobson, and M. Bridges. MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Monthly Notices of the RAS*, 398:1601–1614, October 2009. doi:10.1111/j.1365-2966.2009.14548.x.

[45] C. Gill. *Searching for Gravitational Waves from Pulsars.* Ph.D. thesis, University of Glasgow, 2012.

[46] E. Goetz and K. Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Classical and Quantum Gravity*, 28(21):215006, November 2011. doi:10.1088/0264-9381/28/21/215006.

[47] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):6–80, 2010.

[48] M. Hewitson, H. Grote, S. Hild, H. Lück, P. Ajith, et al. Optimal time-domain combination of the two calibrated output quadratures of GEO 600. *Classical and Quantum Gravity*, 22:4253–4261, October 2005. doi:10.1088/0264-9381/22/20/007.

[49] G. Hobbs, A. Archibald, Z. Arzoumanian, D. Backer, M. Bailes, et al. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Classical and Quantum Gravity*, 27(8):084013, April 2010. doi:10.1088/0264-9381/27/8/084013.

[50] Y.-M. Hu, M. Pitkin, I. S. Heng, and M. A. Hendry. Glitch or Anti-glitch: A Bayesian View. *Astrophys. J. Lett.*, 784:L41, April 2014. doi:10.1088/2041-8205/784/2/L41.

[51] INFN. “Virgo”, project homepage. `wwwcascina.virgo.infn.it/`.

[52] R. Inta and B. Owen. Private Communication, 2015.

[53] B. Iyer, T. Souradeep, C. Unnikrishnan, S. Dhurandhar, S. Raja, A. Kumar, and A. Sengupta. LIGO-India, Proposal of the Consortium for Indian Initiative in Gravitational-wave Observations, 2011. IndIGO Consortium Council `https://dcc.ligo.org/public/0075/M1100296/002/LIGO-India_lw-v2.pdf`.

[54] P. Jaranowski and A. Królak. Searching for gravitational waves from known pulsars using the $\mathcal{F}$ and $\mathcal{G}$ statistics. *Classical and Quantum Gravity*, 27(19):194015, October 2010. doi:10.1088/0264-9381/27/19/194015.

[55] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58(6):063001, September 1998. doi:10.1103/PhysRevD.58.063001.

[56] N. K. Johnson-McDaniel. Gravitational wave constraints on the shape of neutron stars. *Phys. Rev. D*, 88(4):044016, August 2013. doi:10.1103/PhysRevD.88.044016.

[57] D. I. Jones. Gravitational wave emission from rotating superfluid neutron stars. *Mon. Not. Roy. Astron. Soc.*, 402:2503–2519, March 2010. doi:10.1111/j.1365-2966.2009.16059.x.

[58] D. I. Jones. Parameter choices and ranges for continuous gravitational wave searches for steadily spinning neutron stars. 2015. `http://arxiv.org/abs/1501.05832`.

[59] P. B. Jones. Comment on “constraining hadronic superfluidity with neutron star precession”. *Phys. Rev. Lett.*, 92:149001, Apr 2004. doi:10.1103/PhysRevLett.92.149001.

[60] KAGRA. Kagra project homepage. `http://gwcenter.icrr.u-tokyo.ac.jp/en/`.

[61] S. Kawamura, M. Ando, N. Seto, S. Sato, T. Nakamura, et al. The japanese space gravitational wave antenna: DECIGO. *Classical and Quantum Gravity*, 28(9):094011, 2011.

[62] M. Kramer and B. Stappers. Pulsar Science with the SKA. *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 36, 2015.

[63] P. D. Lasky. Gravitational Waves from Neutron Stars: A Review. 2015. In the LIGO Document Control Center `https://dcc.ligo.org/P1500152`.

[64] LIGO Scientific Collaboration, Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *ArXiv e-prints*, April 2013.

[65] LSC-Virgo Continuous Waves Search Group. Establishing Detection Confidence for a Continuous Wave Signal in the Advanced Detector Era. 2015. In the LIGO Document Control Center `https://dcc.ligo.org/T1500151`.

[66] R. G. Lyons. *Understanding Digital Signal Processing (2nd Edition)*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2004. ISBN 0131089897.

[67] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility Pulsar Catalogue. *Astron. J.*, 129:1993–2006, April 2005. doi: 10.1086/428488.

[68] Max Planck Institute for Gravitational Physics (Albert Einstein Institute). “GEO600: The German-British Gravitational Wave Detector”, project homepage. `www.geo600.org/`.

[69] G. Mendell and K. Wette. Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves. *Classical and Quantum Gravity*, 25(11):114044, June 2008. doi:10.1088/0264-9381/25/11/114044.

[70] C. Messenger. *Gravitational Wave Data Analysis: Searching for Gravitational Waves from Low Mass X-Ray Binaries*. Ph.D. thesis, University of Birmingham, 2006.

[71] S. Miyoki. Large scale cryogenic gravitational wave telescope. *Nuclear Physics B - Proceedings Supplements*, 138:439 – 442, 2005. ISSN 0920-5632. doi: http://dx.doi.org/10.1016/j.nuclphysbps.2004.11.101. Proceedings of the Eighth International Workshop on Topics in Astroparticle and Undeground Physics.

[72] C. J. Moore, R. H. Cole, and C. P. L. Berry. Gravitational-wave sensitivity curves. *Classical and Quantum Gravity*, 32(1):015014, January 2015. doi:10.1088/0264-9381/32/1/015014.

[73] G. E. Moss, L. R. Miller, and R. L. Forward. Photon-noise-limited laser transducer for gravitational antenna. *Applied Optics*, 10:2495–2498, 1971. doi: 10.1364/AO.10.002495.

[74] M. Musha, A. Ueda, M. Ohashi, and K.-I. Ueda. Interferometric gravitational wave detector TAMA-300 and the frequency-stabilized laser with high-finesse optical resonator. In A. V. Kudryashov, editor, *Laser Resonators II*, volume 3611 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 65–72. May 1999.

[75] R. Neuhäuser, M. M. Hohle, C. Ginski, J. G. Schmidt, V. V. Hambaryan, and T. O. B. Schmidt. The companion candidate near Fomalhaut - a background neutron star? *Mon. Not. Roy. Astron. Soc.*, 448:376–389, March 2015. doi: 10.1093/mnras/stu2751.

[76] B. J. Owen. Maximum Elastic Deformations of Compact Stars with Exotic Equations of State. *Physical Review Letters*, 95(21):211101, November 2005. doi:10.1103/PhysRevLett.95.211101.

[77] P. Patel. *Search for Gravitational Waves from a nearby neutron star using barycentric resampling.* Ph.D. thesis, California Institute of Technology, Pasadena, California, 2011. `http://thesis.library.caltech.edu/6030/2/ppatel_thesis_ver2.pdf`.

[78] B. Paul. Astrosat: Some Key Science Prospects. *International Journal of Modern Physics D*, 22:1341009, January 2013. doi:10.1142/S0218271813410095.

[79] M. Pitkin, C. Gill, D. I. Jones, G. Woan, and G. S. Davies. First results and future prospects for dual-harmonic searches for gravitational waves from spinning neutron stars. *ArXiv e-prints*, August 2015. Submitted to MNRAS.

[80] M. Pitkin, C. Gill, J. Veitch, E. Macdonald, and G. Woan. A new code for parameter estimation in searches for gravitational waves from known pulsars. *Journal of Physics Conference Series*, 363(1):012041, June 2012. doi:10.1088/1742-6596/363/1/012041.

[81] M. Pitkin, S. Reid, S. Rowan, and J. Hough. Gravitational Wave Detection by Interferometry (Ground and Space). *Living Rev. Relativity*, 14:5, 2011. [Online Article] `http://www.livingreviews.org/lrr-2011-5`, cited: 17 Oct 2014.

[82] M. Pitkin and G. Woan. Searching for gravitational waves from the Crab pulsar – the problem of timing noise. *Classical and Quantum Gravity*, 21:843, March 2004. doi:10.1088/0264-9381/21/5/069.

[83] M. D. Pitkin. *Searches for continuous and transient gravitational waves from known neutron stars and their astrophysical implications.* Ph.D. thesis, University of Glasgow, 2006.

[84] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C (2nd Ed.): The Art of Scientific Computing.* Cambridge University Press, New York, NY, USA, 1992. ISBN 0-521-43108-5.

[85] R. Prix. Search for continuous gravitational waves: Metric of the multidetector $\mathcal{F}$-statistic. *Phys. Rev. D*, 75(2):023004, January 2007. doi:10.1103/PhysRevD.75.023004.

[86] W. Rindler. *Relativity: Special, General and Cosmological.* Oxford University Press, Oxford, United Kingdom, second edition, 2006.

[87] B. S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Relativity*, 12:2, 2009. [Online Article] `http://www.livingreviews.org/lrr-2009-2`, cited: 17 Oct 2014.

[88] P. R. Saulson. Thermal noise in mechanical experiments. *Phys. Rev. D*, 42:2437–2445, Oct 1990. doi:10.1103/PhysRevD.42.2437.

[89] R. T. Schilizzi, P. E. F. Dewdney, and T. J. W. Lazio. The square kilometre array. *Proc. SPIE*, 7012:70121I–70121I–13, 2008. doi:10.1117/12.786780.

[90] I. Semeniuk. Hubble spies lord of the stellar rings, 2005. `http://www.newscientist.com/article/dn7564` cited: 15 May 2015.

[91] M. Shaltev, P. Leaci, M. A. Papa, and R. Prix. Fully coherent follow-up of continuous gravitational-wave candidates: An application to Einstein@Home results. *Phys. Rev. D*, 89(12):124030, June 2014. doi:10.1103/PhysRevD.89.124030.

[92] X. Siemens, B. Allen, J. Creighton, M. Hewitson, and M. Landry. Making h(t) for LIGO. *Classical and Quantum Gravity*, 21:1723, October 2004. doi:10.1088/0264-9381/21/20/015.

[93] D. S. Sivia and J. Skilling. *Data analysis : a Bayesian tutorial.* Oxford science publications. Oxford University Press, Oxford, New York, 2006. ISBN 0-19-856831-2. Table des matires http://www.loc.gov/catdir/enhancements/fy0627/2006284782-t.html.

[94] J. Skilling. American Institute of Physics Conference Series Vol. 365. In R. Fischer, R. Preuss, and U. Toussaint, editors, *American Institute of Physics Conference Series*, pages 395–405. 2004.

[95] J. W. Tukey. An introduction to the calculations of numerical spectrum analysis. In B. Harris, editor, *Spectral Analysis of Time Series*, pages 25–46. Wiley, New York, 1967.

[96] M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, et al. LOFAR: The LOw-Frequency ARray. *Astron. Astrophys.*, 556:A2, August 2013. doi:10.1051/0004-6361/201220873.

[97] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys. Rev. D*, 81(6):062003, March 2010. doi:10.1103/PhysRevD.81.062003.

[98] S. J. Waldman. The Advanced LIGO Gravitational Wave Detector. *ArXiv e-prints*, March 2011.

[99] M. Wąs, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, et al. On the background estimation by time slides in a network of gravitational wave detectors. *Classical and Quantum Gravity*, 27(1):015005, January 2010. doi:10.1088/0264-9381/27/1/015005.

[100] K. Wette, B. J. Owen, B. Allen, M. Ashley, J. Betzwieser, et al. Searching for gravitational waves from Cassiopeia A with LIGO. *Classical and Quantum Gravity*, 25(23):235011, December 2008. doi:10.1088/0264-9381/25/23/235011.

[101] G. Woan. *The Cambridge Handbook of Physics Formulas*. Cambridge University Press, 2000. ISBN 9780511755828. Cambridge Books Online.

[102] N. Yunes and X. Siemens. Gravitational-Wave Tests of General Relativity with Ground-Based Detectors and Pulsar-Timing Arrays. *Living Rev. Relativity*, 16:9, 2013. [Online Article] `http://www.livingreviews.org/lrr-2013-9`, cited: 29 Jan 2015.