

Dupuis, Réjean J. (2004) *Bayesian searches for gravitational waves from pulsars*. PhD thesis.

<http://theses.gla.ac.uk/5714/>

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

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

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

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

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

Bayesian searches for gravitational waves from pulsars

Réjean J. Dupuis

Department of Physics and Astronomy

University of Glasgow

Presented as a thesis for the degree of Ph.D.

in the University of Glasgow, University Avenue,

Glasgow G12 8QQ.

©Réjean J. Dupuis 12th December 2004



“Physics is like sex: sure, it may give some practical results, but that’s not why we do it.”— Richard P. Feynman

Acknowledgments

First, I would like to thank all those who made this research possible by their financial support. I was funded by scholarships from the Natural Sciences and Engineering Research Council of Canada, the University of Glasgow, and Universities UK. I would like to thank Jim Hough for always finding money when I wanted to go to conferences and meetings. Thanks also go to Chris Fryer for the financial support during my visits to Los Alamos.

Thanks are due to many people who have helped me through these years. Those who have assisted with the work in this thesis include my supervisor Graham Woan, members of the the LIGO Scientific Collaboration Pulsar Upper Limit Group, and the internal LSC PULG reviewers. I would also like to thank many of my collaborators over the years including Bruce Allen, Stuart Anderson, Nelson Christensen, Curt Cutler, Martin Hewitson, Greg Mendell, Michael Landry, Maria Alessandra Papa, Matthew Pitkin, Peter Shawhan, Xavier Siemens, John Veitch, and Uta Weiland.

I would like to thank members of the gravitational wave group and the astronomy and astrophysics group here in Glasgow for their support and assistance. I would like to thank Jim Hough whose scholarship has been a continuing inspiration and for his encouragement and support. I am also indebted to Graeme Stewart for helping me with several computational issues over the last few years.

I would like to thank my friends and colleagues in Los Alamos including Chris Fryer,

Aimee Hungerford, Alex Heger, Gabriel Rockefeller, and Mike Warren. I learned many lessons from Chris Fryer and am grateful to have had the opportunity to work with him over the last few years. I will always have fond memories of the Friday night astrophysics dinners and appreciate his encouragement and patience.

I would like to thank Michael Kramer and others at Jodrell Bank Observatory for providing up-to-date pulsar ephemerides for the S2 analysis.

Thanks to Morag Casey, Jim Hough, Ken Strain, Harry Ward, and Graham Woan for comments on the draft of this thesis and providing valuable feedback.

Finally, I would like to thank my family for their support and encouragement over the years. I am grateful to Ayşe for her support through this demanding period.

Preface

This thesis is an account of work carried out between October 2000 and October 2004 towards the analysis of data from interferometric gravitational wave detectors for signals from pulsars. This work was carried out mainly by myself under the supervision of Graham Woan of the University of Glasgow with input from several members of the LIGO Scientific Collaboration (LSC) Periodic Upper Limits Group unless noted otherwise. Many members of the LSC were responsible for the operation of the detectors and the collection of data in which I did not directly participate.

In the first chapter, I introduce the concept of gravitational waves, how they interact with matter, and possible astrophysical sources. Several gravitational wave detectors are introduced and background information on Bayesian probability theory is given. The work in this chapter is mostly derived from published literature.

In Chapter 2, I present a new data analysis pipeline to search for gravitational waves from periodic sources using interferometric gravitational wave detectors. The method was developed specifically for targeted searches for signals from known radio pulsars and is applied to GEO 600 and LIGO data in subsequent chapters. In this chapter, the developed software is tested on simulated data and several sanity checks are carried out. The work in this chapter was carried out mainly by myself with some input from Graham Woan in the early stages.

In Chapter 3, the data set from the first science run (S1) of GEO 600 and LIGO is analysed

for gravitational wave signals from pulsar B1937+21. The analysis techniques developed in Chapter 2 are applied to the S1 data and an upper limit on the gravitational wave strain from this source is set. Further software checks on the barycentring routines are presented and the calibration procedure for the LIGO data is explained. The work in this chapter was carried out by myself unless noted otherwise.

In Chapter 4, I report on the analysis the data from LIGO's second science run (S2) searching for gravitational waves from 28 isolated pulsars. Hardware injections of two periodic signals into the LIGO interferometers are analysed to provide a robust end-to-end validation of the overall pipeline. A detailed assessment of the normality and stationarity of the data is also given. Upper limits on the gravitational wave strain from those 28 pulsars are presented. The work presented in this chapter was carried out by myself unless noted otherwise.

In Chapter 5, a preliminary analysis of the GEO 600 and LIGO S3 data for the same 28 pulsars studied in S2 is described. The analysis is essentially a repeat of the previous analyses with significantly more sensitive data. I also describe the outcome of the analysis for 10 hardware injections injected into the LIGO interferometers as well as one signal that was simultaneously injected into GEO 600 and LIGO. Preliminary results are provided further constraining the gravitational wave strain from those sources. The work presented in this chapter was carried out by myself unless noted otherwise.

In Chapter 6, brief conclusions are drawn from the drawn from the work presented in this thesis.

The Appendix contains three papers published by myself and colleagues during the course of my graduate studies on subject matters related to this thesis.

Summary

Gravitational waves, ripples in spacetime, are predictions of Einstein's General Theory of Relativity. Since the 1960s, there has been an effort to detect these waves directly; a pursuit so far unsuccessful. However, a new generation of interferometric gravitational wave detectors with unprecedented sensitivity have begun collecting data. The success of our hunt for gravitational waves will depend, in part, on our ability to extract weak signals buried in noise. Several different types of gravitational waveforms are expected to be emitted from various astrophysical sources. In order to make optimal use of the data, specialized data analysis procedures have been developed for each class of signal.

Methods for searching for periodic gravitational wave signals from triaxial pulsars using interferometric gravitational wave detectors have been developed. Since the gravitational wave signals from pulsars are expected to be weak, long stretches of data must be used for any detection. Over the course of a day, and a year, these periodic signals are Doppler shifted due to the motion of the Earth. The response of the interferometers to each polarisation of gravitational waves will also give rise to an amplitude modulation of the periodic signal. These effects are taken into account and an end-to-end Bayesian scheme for making inferences from the data is presented. Several software tests have been performed to validate the core routines, such as barycentring, using independent software.

The GEO 600 and LIGO interferometers had their first scientific data run (S1) for 17 days between 23 August and 9 September 2002. An analysis was carried out to search

for gravitational wave signals from pulsar B1937+21. While no signals were detected, a 95% upper limit of $h_0 < 1.4 \times 10^{-22}$ was determined using S1 data where h_0 is the amplitude of the gravitational waves. Given that pulsar B1937+21 is at a distance of 3.6 kpc, and assuming a moment of inertia of 10^{38} kg m^2 , the corresponding upper limit on the equatorial ellipticity was determined to be $\epsilon \leq 2.9 \times 10^{-4}$. The upper limit on gravitational waves from pulsar B1937+21 using S1 data was over an order of magnitude lower than the previous best limit at the time.

Data from LIGO's second science run (S2) in the spring of 2003 was analysed with the sensitivity of each detector in the network being roughly an order of magnitude better than in S1 across a large range of frequencies. Upper limits were placed on a total of 28 isolated pulsars using the S2 data. The analysis procedure for S2 was more robust to interfering spectral lines and took advantage of the longer stationarity of the S2 data. Two hardware injections of hypothetical pulsars were injected in the LIGO interferometers during S2. The successful extraction of these signals from the LIGO S2 data significantly increased our confidence in the overall data analysis pipeline. For four of the closest pulsars their equatorial ellipticities were constrained to less than $\epsilon \leq 10^{-5}$ with 95% confidence. These limits are beginning to reach interesting ellipticities which some exotic theories suggest could be supported in neutron stars.

The third science run (S3) in which GEO 600 and LIGO participated took place from late October 2003 to early January 2004. Again, the improvement in sensitivity compared to the previous run (S2) was significant. Preliminary multi-detector results were determined for the same previous 28 pulsars using S3 data. The equatorial ellipticities for 11 of these pulsars are constrained to $\epsilon \leq 10^{-5}$ with 95% confidence. With the S3 data, the upper limit on the gravitational wave emission from the Crab pulsar was only approximately a factor of four from the upper limits inferred from the spindown of the pulsar. When this barrier is overcome the prospects of detecting gravitational waves from the Crab pulsar will become more plausible.

Future work based on these implementations will examine a larger set of missing known pulsars including binary systems. Studies in Markov Chain Monte Carlo techniques may also allow the expansion this method to a larger parameter space.

Contents

Acknowledgments	ii
Preface	iv
Summary	vi
1 Introduction	1
1.1 Background	1
1.2 Nature of gravitational waves	3
1.3 The detectors	5
1.4 Generation of gravitational waves	8
1.4.1 Coalescing binaries	11
1.4.2 Supernovae	13
1.4.3 Gamma-ray bursts	13
1.4.4 Neutron stars	14
1.4.5 Stochastic background	15
1.4.6 Cosmic strings	16

1.4.7 Unknown sources	16
1.5 Bayesian methods in data analysis	17
1.6 Outline of the thesis	19
2 Bayesian methods to search for periodic gravitational waves	21
2.1 Introduction	21
2.2 Nature of the gravitational wave signal	23
2.2.1 Response of interferometric detectors	23
2.2.2 Gravitational waves from triaxial neutron stars	27
2.2.3 Pulsar timing: phase of the gravitational wave signal	29
2.3 Searching for signals from pulsars	30
2.3.1 Data processing	31
2.3.2 Bayesian analysis	34
2.3.3 Gaussian model – known variance	37
2.3.4 Gaussian model – unknown variance	38
2.3.5 Combining data from a network of detectors	43
2.3.6 Marginalisation over parameters and setting upper limits	44
2.4 Performance of algorithm on simulated data	45
2.4.1 Comparison of Gaussian and Student's t likelihoods	46

2.4.2 Effects of changing the observation period	47
2.4.3 Effects of changing the noise level	49
2.4.4 Combining results from several detectors	52
2.4.5 Parameter estimation	52
2.4.6 Sensitivity to pulsar orientation	56
2.4.7 Repeated experiments with no signal	57
2.5 Discussion	58
3 Analysis of the first data from GEO 600 and LIGO	60
3.1 Introduction	60
3.2 The detectors	62
3.2.1 Instrument configurations	63
3.3 Expected sensitivity to periodic signals	64
3.4 Gravitational waves from pulsar B1937+21	67
3.5 Accuracy of timing software	70
3.6 Calibrating the LIGO data	72
3.7 Data analysis method	78
3.8 Characterisation of the data	80
3.9 Results	88

3.10	Conclusions	90
4	Analysis of LIGO S2 data for 28 pulsars	91
4.1	Introduction	91
4.2	Pulsar selection	94
4.3	Analysis pipeline	103
4.4	Hardware injections	104
4.5	Characterisation of S2 data	108
4.5.1	Assessing Gaussianity and stationarity	108
4.5.2	Issues with applying S1 pipeline on S2 data	109
4.5.3	Modified reduced chi-square test	110
4.5.4	Kolmogorov-Smirnov (Lilliefors) normality test	115
4.5.5	F -test for equal variances	121
4.6	Results	123
4.7	Conclusions	127
5	Preliminary analysis of GEO 600 and LIGO S3 data	136
5.1	Introduction	136
5.2	Hardware injections	138
5.3	Characterisation of the S3 data	144

5.4 Preliminary results for 28 pulsars	148
5.5 Conclusions	152
6 General conclusions	153
Bibliography	155
A Related publications	165

List of Figures

1.1 Effect of gravitational wave on a ring of free falling particles.	5
1.2 Schematic layout of basic gravitational wave interferometer.	6
1.3 Target sensitivities for GEO 600, LIGO, Advanced LIGO, and Virgo.	8
1.4 Target noise level for LIGO and Advanced LIGO with several potential astrophysical sources.	12
2.1 Sensitivity of an interferometer to plus polarisation of a gravitational wave.	25
2.2 Sensitivity of an interferometer to cross polarisation of a gravitational wave.	26
2.3 Sensitivity of an interferometer to an unpolarised gravitational wave.	26
2.4 Simple diagram of a triaxial neutron star.	27
2.5 Comparison of Gaussian distribution and Student's t -distribution with varying degrees of freedom.	41
2.6 Comparison of posterior pdfs for h_0 between Gaussian likelihood and Student's t likelihood.	47
2.7 Comparison of posterior pdfs for h_0 for different observation periods.	48
2.8 Time series of B_k for a simulation checking the effect of changing noise levels.	50

2.9 Average posterior pdfs for three scenarios with different noise levels.	51
2.10 Multi-detector posterior pdfs with simulated data.	53
2.11 Posterior pdfs for artificial Signal A.	54
2.12 Posterior pdfs for artificial Signal B.	55
2.13 Two joint posterior pdfs for Signal B.	56
2.14 Contours of constant probability, for h_0 and $\cos \iota$, for Gaussian noise with no signal.	57
2.15 Distribution of the most probable value of h_0 when there is no signal present.	58
2.16 Distribution of the 95% upper limit on h_0 when there is no signal present. .	59
3.1 Typical sensitivity of the three LIGO interferometers and GEO 600 during the S1 data run.	63
3.2 Distribution of 95% upper limits on h_0 for 4000 sources covering the whole sky with $T = 1$ and $S_n(f) = 1$	65
3.3 Typical S1 sensitivities $\langle h_{95} \rangle$ along with spindown based upper limits for several pulsars.	67
3.4 Radio timing residuals for pulsar B1937+21 over several years including the duration of the S1 run.	69
3.5 Beam pattern of GEO 600, LLO, and LHO, for pulsar B1937+21 during S1.	69
3.6 Diagram with Earth, Sun, Solar System Barycenter, and an incoming gravi- tational wave.	70

3.7 Roemer and Shapiro delay for a source in the direction of pulsar B1937+21.	72
3.8 Einstein delay for a source in the direction of pulsar B1937+21 and histogram of the timing residuals between our barycentring routines and TEMPO. . . .	73
3.9 Simple model of the LIGO length sensing and control feedback system. . .	74
3.10 Overall open loop gain factor $\alpha(t)\beta(t)$ over the course of S1 for the three LIGO IFOs.	76
3.11 Reference open loop gain $G_0(f)$ taken at GPS time 715 156 759 for the three LIGO IFOs.	77
3.12 Reference sensing function $C_0(f)$ taken at GPS time 715 156 759 for the three LIGO IFOs.	77
3.13 Reference response function at GPS time 715 156 759 for the three LIGO IFOs.	78
3.14 Single-sided power spectral density of a band of 4 Hz (starting at 1 281 Hz) using the S1 data set analyzed from GEO 600, H1, H2, and L1.	82
3.15 Spectrogram of GEO 600 data during S1 for 4 Hz band centered on 1284 Hz.	82
3.16 Time series of $ B_k $ using GEO 600 data from S1.	83
3.17 Time series of $ B_k $ using H1 data from S1.	84
3.18 Time series of $ B_k $ using H2 data from S1.	85
3.19 Time series of $ B_k $ using L1 data from S1.	86
3.20 Histogram of $B/\sigma = \Re(B_k)/\sigma_{\Re(B_k)} + \Im(B_k)/\sigma_{\Im(B_k)}$ for each interferometer using the S1 data sets.	87

3.21 Marginalised posterior pdfs for h_0 for pulsar B1937+21 using S1 data. . . .	89
4.1 Typical sensitivities of the three LIGO interferometers during the S2 data run.	92
4.2 Antenna pattern of LLO detector at the beginning of S2.	101
4.3 Antenna pattern of LHO detectors at the beginning of S2.	102
4.4 Average antenna pattern of LIGO IFOs over the length of S2.	102
4.5 Envelope of amplitude of the S2 hardware injections in both interferometers.	106
4.6 Recovered parameters for signal P1 of the S2 hardware injections.	107
4.7 Recovered parameters for signal P2 of the S2 hardware injections.	108
4.8 Reduced chi-square for all isolated pulsars when the noise is estimated from a 4 Hz bandwidth with the S1 IIR filter cutoffs.	110
4.9 Fourier transform of a 4 Hz band around the signal frequency for pulsar J1910-5959B in the H2 data.	111
4.10 Reduced chi-square for each detector when the noise level is estimated every 30 minutes from the B_k 's.	116
4.11 Kolmogorov-Smirnov statistic defined as the vertical distance D between the two cumulative distribution functions.	118
4.12 Cumulative distribution of first 30 minute segment of S2 data for pulsar B1937+21 compared to expected cumulative Gaussian distribution.	118
4.13 Cumulative distribution of first 30 minute segment of S2 data for the Crab pulsar B0531+21 compared to expected cumulative Gaussian distribution. .	119

4.14	Percentage of data which is rejected at a 0.05 significance level with a Kolmogorov-Smirnov (Lilliefors) test for normality.	120
4.15	Distribution of F -statistics for data for pulsar B1937+21 compared to expected F -distribution.	122
4.16	Distribution of F -statistics for data for the Crab pulsar B0531+21 compared to expected F -distribution.	122
4.17	Percentage of 30 minute segments that are rejected at a significance level of 0.05 when the sample variances from the first and second 15 minutes are compared using an F -test.	123
4.18	Percentage difference in h_0 upper limits if the non-normal data are excluded from the analysis.	126
4.19	Marginalised probability distributions functions for pulsars B0021–72C, B0021–72D, B0021–72F, and B0021–72G.	129
4.20	Marginalised probability distributions functions for pulsars B0021–72L, B0021–72M, B0021–72N, and B0531+21.	130
4.21	Marginalised probability distributions functions for pulsars B1516+02A, B1820–30A, B1821–24, and B1937+21.	131
4.22	Marginalised probability distributions functions for pulsars B1951+32, J0030+0451, J0711–6830, and J1024–0719.	132
4.23	Marginalised probability distributions functions for pulsars J1629–6902, J1721–2457, J1730–2304, and J1744–1134.	133
4.24	Marginalised probability distributions functions for pulsars J1748–2446C, J1910–5959B, J1910–5959C, and J1910–5959D.	134

4.25 Marginalised probability distributions functions for pulsars J1910–5959E, J1913+1011, J2124–3358, and J2322+2057.	135
5.1 Typical sensitivities of the three LIGO and the GEO 600 interferometers during the S3 data run.	137
5.2 Posterior probability distributions for the injected signals P0 and P1.	139
5.3 Posterior probability distributions for the injected signals P2 and P3.	140
5.4 Posterior probability distributions for the injected signals P4 and P5.	140
5.5 Posterior probability distributions for the injected signals P6 and P7.	141
5.6 Posterior probability distributions for the injected signals P8 and P9.	141
5.7 Posterior probability distributions for the injected signal P10.	143
5.8 Percentage of data rejected at a significance level of 0.05 when a K-S test is applied to S3 data.	145
5.9 Percentage of data rejected at a significance level of 0.05 when an F -test is applied to S3 data.	146
5.10 Improvements in S3 upper limits on h_0 for the 28 pulsars compared to the S2 limits.	148
5.11 Posterior probability distributions for pulsar B1937+21 using S3 data from GEO 600 and the three LIGO IFOs.	149
5.12 Percent change in the S3 upper limits on h_0 if the non-normal data are ex- cluded from the analysis.	151

List of Tables

2.1 Source parameters for both artificial signals extracted from data.	52
3.1 Total periods where the interferometers were in-lock and the duty cycles for each detector during the S1 run.	62
3.2 Parameters for pulsar B1937+21.	68
4.1 Positional information for 28 pulsars selected for S2 analysis.	97
4.2 Spin information for 28 pulsars selected for S2 analysis.	98
4.3 Second spin frequency derivative for the 7 pulsars for which it is measured.	99
4.4 Cumulative phase error from uncertainties in spin and position.	100
4.5 Signal parameters for the S2 hardware injections.	105
4.6 Upper limits on h_0 for the 28 selected pulsars using the S2 data.	125
5.1 Signal parameters for the 11 periodic signals injected into the LIGO inter- ferometers in S3.	138
5.2 Percentage of data rejected due to F -test and K-S test for each pulsar with S3 data.	147

5.3 Upper limits on for the 28 selected pulsars using the S3 data.	150
--	-----

Chapter 1

Introduction

1.1 Background

The General Theory of Relativity developed by Albert Einstein [1] predicts the existence of gravitational waves; ripples in the fabric of space-time. The proposition that a gravitational field travels at a finite speed led to this phenomenon that was not previously relevant in Newtonian physics, where gravity was thought to act instantaneously. For several decades after the publication of Einstein's theory of gravitation in 1916, the existence of gravitational waves was either ignored or dismissed as a mathematical oddity. It would be the 1960's before Joseph Weber at the University of Maryland began his pioneering work in gravitational wave detection [2].

Although no direct detections have been confirmed to date, there is strong indirect evidence for the existence of gravitational waves. In fact, the 1993 Nobel Prize for Physics was awarded to Joseph Taylor and Russell Hulse for their joint discovery of the first binary pulsar which provided such evidence [3]. They discovered this binary neutron star system in 1973 and accurately measured the time-of-arrivals of the radio pulses from one of the stars. Whereas Newtonian physics predicts that the orbital period should remain

constant over time, General Relativity predicts that the period should decrease due to a loss of orbital energy to gravitational waves. From these radio observations, Hulse and Taylor were able to confirm that the 8 hours orbital period of this system was decreasing within 1% of the rate predicted by General Relativity. With an orbital period of 8 hours, however, we will not be able to detect these gravitational wave signals with ground based detectors.

In addition to confirming one of the main predictions of General Relativity, the direct observation of gravitational waves could open a new window on the universe; a new field of observational astronomy. Gravitational waves could provide information on astrophysical phenomenon not seen via electromagnetic radiation. Even for sources already seen electromagnetically, gravitational waves could provide a different view; whereas electromagnetic waves are usually emitted incoherently from astronomical sources, gravitational waves are emitted coherently by the bulk motion of these objects.

Today there is a worldwide effort dedicated to the direct detection of gravitational waves (see [4] for recent overview). The technological difficulties in constructing gravitational wave detectors of sufficient sensitivity make it one of the most ambitious goals of present day experimental physics. But the difficulties with establishing the field of gravitational wave astronomy do not end there. Since the gravitational waves are expected to be weak, and it takes a long time for the detectors to improve their sensitivities, a whole field of research has emerged in trying to extract weak gravitational wave signals buried in noise. Indeed, the bulk of this thesis focuses on data analysis techniques for one potential source of gravitational waves: rapidly rotating neutron stars.

In this chapter, a brief introduction to gravitational waves is given including a description of their nature, the detectors, and potential astrophysical sources. Further discussion concentrates on Bayesian methods in data analysis and finally an overview of the thesis is provided.

1.2 Nature of gravitational waves

The General Theory of Relativity explains gravitation as curvature in spacetime [5]. Gravitational waves are ripples in spacetime and propagate at the speed of light. In this section, we explain very briefly the basics of relativity, the nature of gravitational waves, and their effect on inertial masses.

The Einstein gravitational field equations describe the relationship between matter and curvature. The form of the field equations is given by [5]

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu}, \quad (1.1)$$

where $R_{\mu\nu}$ is the Ricci tensor, $g_{\mu\nu}$ is the spacetime metric, R is the scalar curvature, G is the gravitational constant, and $T_{\mu\nu}$ is the energy momentum tensor. For the purposes of this description, it is sufficient to note that in a vacuum the energy momentum tensor $T_{\mu\nu}$ is null, so the field equations reduce to a simple wave equation [6]

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)h_{\mu\nu} = 0, \quad (1.2)$$

with waves propagating at the speed of light c .

We will begin by describing the concept of proper distance between two events in flat spacetime and then introduce a small perturbation to examine the weak field gravitational waves that result. In flat spacetime we can use Einstein's Theory of Special Relativity to define the invariant, ds^2 , between two events as

$$ds^2 = -c^2dt^2 + dx^2 + dy^2 + dz^2. \quad (1.3)$$

This can be expressed more generally as

$$ds^2 = \eta_{\mu\nu}dx^\mu dx^\nu, \quad (1.4)$$

where summation is assumed over each index and the flat Minkowski metric $\eta_{\mu\nu}$ is given by

$$\eta_{\mu\nu} = \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (1.5)$$

In the presence of a gravitational field the generalised spacetime metric $g_{\mu\nu}$ becomes necessary. Here we use coordinates that are nearly Minkowski with a small perturbation $h_{\mu\nu}$ as a first approximation. The flat metric plus a small deviation can be expressed as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (1.6)$$

With the appropriate coordinate transformation, called transverse-traceless, it can be shown that for a wave propagating in the z direction the metric takes the form

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{xx} & h_{xy} & 0 \\ 0 & h_{yx} & h_{yy} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.7)$$

where $h_+ = h_{xx} = -h_{yy}$ and $h_\times = h_{xy} = h_{yx}$ (transverse since there is no z component and traceless since $h_{xx} + h_{yy} = 0$). The two distinct polarisations of gravitational waves, 'plus' and 'cross', are separated by 45 degrees. A passing gravitational wave will have a stretching and squeezing effect on matter depending on the phase and polarisation of the gravitational wave. A gravitational wave of amplitude h will give a total strain for each polarisation of

$$h = \frac{2\delta L}{L}, \quad (1.8)$$

where, for example, L is the original radius of a ring of particles and δL is its deviation due to a passing a gravitational wave. The effect of this oscillating field on free falling




Figure 1.1: Effect of gravitational wave, of amplitude $h = 0.2$, on a ring of free falling test particles as function of time, from [7]. The upper part shows the effect of the ‘plus’ polarisation and the bottom part shows the effect of the ‘cross’ polarisation. The propagation direction of the wave is orthogonal to the plane in which the particles lie.

particles arranged in a circle is shown in Figure 1.1.

1.3 The detectors

There are two main types of detectors used in ground-based gravitational wave detection: resonant bars and laser interferometers.

The first resonant bar detector was constructed by Joseph Weber in the 1960’s [8]. The main resonant bar detectors currently in operation around the world are Explorer [9], Allegro [10], Niobe [11], Nautilus [12], and Auriga [13]. These detectors are generally large metallic cylinders with sensitivity to gravitational waves around their fundamental resonant frequencies. If a gravitational wave is incident on such a detector it will cause the bar to vibrate which can, in principle, be measured accurately. The main drawback to this type of detector is that it is only sensitive to gravitational waves in a narrow band of a few tens of Hz.

Laser interferometric gravitational wave detectors can be sensitive to gravitational waves in a wide bandwidth, from a few tens of Hz to several kHz. The nature of gravitational




Figure 1.2: Schematic layout of basic gravitational wave interferometer, from [18], where L_1 and L_2 are the lengths of the arms, and R shows the location of the power recycling mirror which recycles photons back into the interferometer.

waves make interferometric detectors ideal to measure the effects of gravitational waves incident on free test masses. These interferometric detectors were first proposed in the 1970s with some details in [14]. Several prototypes constructed subsequently including a 10m interferometer in Glasgow [15], a 30m interferometer in Garching [16], and a 40m interferometer at CIT¹ [17]. For these detectors, a laser beam is incident on a beamsplitter which transmits the light to mirrors suspended at the end of two orthogonal arms. The light is then recombined and since a gravitational wave signal will affect the light travel time in each arm differently it will influence the measured interference pattern. Figure 1.2 provides a simple schematic layout for an interferometric gravitational wave detector. The major interferometric gravitational wave detectors currently operating in the world are:

GEO 600 A British and German collaboration with the detector located in Hannover, Germany. This detector is a dual-recycled Michelson interferometer with 600m arms, an advanced optical layout including signal recycling, and fused-silica mirror suspensions. The automatic alignment systems in GEO 600 have so far permitted very high duty cycles during science runs [19].

¹California Institute of Technology

LIGO An American collaboration operated by CIT and MIT² with two detectors located in Hanford, WA, and one detector in Livingston, LA. The two detectors in WA are located at the same site and are 2 km and 4 km in length, and the detector in LA has arms of 4 km in length. These are Michelson interferometers and make use of Fabry-Perot cavities to increase the effective length of the arms. After about 2 years of operation at design sensitivity they will likely be upgraded to Advanced LIGO [18].

Virgo A French and Italian collaboration with the detector located near Pisa, Italy. This is a 3 km Michelson interferometer with Fabry-Perot cavities similar to LIGO. The novel seismic isolation system in Virgo will give it the best sensitivity among current detectors at low frequencies, from ~ 10 Hz to 40 Hz [20].

TAMA A Japanese detector located near Tokyo, Japan. This is a 300 m arm length Michelson interferometer with Fabry-Perot cavities in each arm. The main aim of the project is research for the eventual construction of a future kilometer sized interferometer in Japan [21].

Several noise sources will limit the sensitivity of current ground-based interferometers. At low frequencies, below about 50 Hz, the main source of noise is of seismic origin. Thermal noise from the mirror masses and the suspensions wires limits the sensitivity of the interferometers from about 50 Hz to 250 Hz. The sensitivity of these detectors at high frequencies is limited by photon shot noise which is due to statistical fluctuations in the number of photons detected. A full review of the fundamental noise limits on ground-based interferometers is given in [22].

Figure 1.3 shows the targeted sensitivities for GEO 600, LIGO, the proposed upgrade Advanced LIGO, and Virgo.

²Massachusetts Institute of Technology




Figure 1.3: Target sensitivities for GEO 600, LIGO, Advanced LIGO, and Virgo.

1.4 Generation of gravitational waves

Gravitational waves are generated when masses are accelerated. This is similar to the emission of electromagnetic radiation when charges are accelerated. However, for matter the lowest order of a mass distribution that can change in an oscillatory manner is the quadrupole moment. Oscillatory changes in the monopole moment of mass distribution are forbidden due to mass conservation, and changes in the dipole moment forbidden due to conservation of momentum.

The amplitude of observed gravitational waves, in the slow-motion approximation, due to changes in the quadrupole moment of the source is [23]

$$h_{\mu\nu} = \frac{2G}{c^4 r} \frac{d^2 Q_{\mu\nu}}{dt^2}, \quad (1.9)$$

where r is the distance to the source and Q is the reduced quadrupole moment defined as

$$Q_{\mu\nu} = \int \rho(\vec{r}) \left(x_\mu x_\nu - \frac{1}{3} \delta_{\mu\nu} r^2 \right) dV, \quad (1.10)$$

where $\rho(\vec{r})$ is the mass density and the integral is over the volume of the source. It is clear from Equation 1.9 that only massive objects with very large accelerations will emit large amplitude gravitational waves. Since it is impossible to create such environments in a controlled experiment we must turn to nature for sources of gravitational waves. Binary neutron star systems lose energy to gravitational waves but even for those extreme systems we can see that the strain at the Earth will be very small.

For example, here we use Equation 1.9 directly to calculate the gravitational wave strain from a binary system (following an example in [23]). If we let the \vec{z} -axis be normal to the plane of the orbit, r_0 be the orbital radius, f_{orb} be the orbital frequency, and M be the mass of both objects, then from Equation 1.10 it follows that

$$Q_{xx} = 2MR^2 \left[\cos^2(2\pi f_{\text{orb}}t) - \frac{1}{3} \right], \quad (1.11)$$

$$Q_{yy} = 2MR^2 \left[\sin^2(2\pi f_{\text{orb}}t) - \frac{1}{3} \right], \quad (1.12)$$

and

$$Q_{xy} = Q_{yx} = 2MR^2 [\cos(2\pi f_{\text{orb}}t) \sin(2\pi f_{\text{orb}}t)], \quad (1.13)$$

with the other terms being either constant or zero. For a system that is a distance of R along the z axis we can calculate the gravitational wave amplitude seen at the origin. By calculating the double time derivatives of the previous equations and using the double angle trigonometric identities³, we find from Equation 1.9 that

$$h_{xx} = -h_{yy} = h \cos(4\pi f_{\text{orb}}t), \quad (1.14)$$

and

$$h_{xy} = h_{yx} = -h \sin(4\pi f_{\text{orb}}t), \quad (1.15)$$

³ $\sin(2u) = 2 \sin u \cos u$ and $\cos(2u) = \cos^2 u - \sin^2 u$

where

$$h = \frac{32\pi^2 G}{Rc^4} M r_0^2 f_{\text{orb}}^2. \quad (1.16)$$

We note that gravitational waves are emitted at twice the orbital frequency of the system. If we take two neutron stars with masses of $1.4 M_\odot$ with an orbital radius of 20 km (almost touching), Newtonian mechanics states that their orbital frequency will be ~ 400 Hz. Assuming the system is in the Virgo Cluster at a distance of 15 Mpc then we have that $h \approx 1 \times 10^{-21}$. The number of binary neutron star systems in the universe is uncertain and this will influence the rate and, thus, the likely distance of such sources. However, arguments such as those given above are why it is believed that with the sensitivity of the current detectors, shown in Figure 1.3, we may soon be able to detect gravitational waves.

In a similar example, which is more relevant to this thesis, we can estimate the amplitude of the gravitational wave signal emitted by a rotating neutron star. We consider a rotating ellipsoid with semiaxes a , b , and c , which is rotating with angular velocity $(0, 0, \Omega)$ around one of the principal axes. Using the definition of the quadrupole moment and a few trigonometric identities, the quadrupole moment can be written as

$$Q_{\mu\nu} = \frac{I_2 - I_1}{2} \begin{pmatrix} \cos 2\varphi & -\sin 2\varphi & 0 \\ -\sin 2\varphi & -\cos 2\varphi & 0 \\ 0 & 0 & 0 \end{pmatrix} + \text{constant} \quad (1.17)$$

where $\varphi = \Omega t$, $I_1 = \frac{M}{5}(b^2 + c^2)$, and $I_2 = \frac{M}{5}(c^2 + a^2)$. We note that the gravitational wave frequency will be at twice the rotation frequency and that there is no gravitational wave emission if the neutron star is axisymmetric about the rotation axis, that is if $a = b$. If we let $\epsilon = \frac{I_2 - I_1}{I_3}$ then from Equation 1.9 we get an expression for the gravitational wave amplitude such that

$$h = \frac{4G\Omega^2}{c^4 r} I_3 \epsilon. \quad (1.18)$$

We cover the waveform of such a signal from a neutron star more extensively in Chapter 2.

The most likely astrophysical sources in the bandwidth for the current ground-based detectors can be divided into three main classes: transients, periodic, and stochastic. Transients sources are burst-type signals of short duration. Periodic signals are long term continuous sources of nearly constant frequency in an inertial reference frame with respect to the source. Stochastic sources, on the other hand, consist of a large number of incoherent signals or a primordial background analogous to the cosmic microwave background.

For each type of signal, different data analysis strategies must be developed. Assuming optimal analysis methods, the estimated signal strengths, h_s , of various astrophysical gravitational waves sources along with the expected sensitivities of LIGO detectors are shown in Figure 1.4 (from Kip Thorne). The signal strengths, h_s , are defined in such a way that if a signal is present at any direction in the sky with random orientation they would be detectable with a false alarm probability of less than one percent.

For a detailed overview of astrophysical gravitational wave sources, see [24]. Below we describe each of the sources in Figure 1.4, and others.

1.4.1 Coalescing binaries

We have already seen that orbiting compact objects inspiral together due to the loss of orbital energy to gravitational waves. In the final stages of coalescence, binary systems emit chirps of gravitational waves, characterized by a rise in frequency and amplitude of the signal in the final seconds before merging. These systems can comprise of neutron stars (NS/NS), a neutron star and a black hole (NS/BH), or two black holes (BH/BH). The three arrowed, long-dashed lines in Figure 1.4 show the strengths of these gravitational wave signals in the last few seconds of inspiralling compared to the sensitivities of the LIGO detectors. The main uncertainty for those signals is determining the expected rate of such inspirals; this is related to the distribution of binary systems in the galaxy and




Figure 1.4: The target noise level for LIGO and Advanced LIGO plotted with estimated signal strength for several sources of gravitational waves (figure from Kip Thorne). The signal strength is defined in such a way that whenever a signal is present at that strength from any direction in the sky, in any orientation, it will be detected with a false alarm probability of less than one per cent.

universe.

In 2003, a highly relativistic binary system containing pulsar J0737-3039 was discovered using the Parkes radio telescope [25]. This system alone increased the previous double neutron star (NS/NS) coalescence rate by a factor of 5 to 7 compared to previous estimates [26]. The current best estimates for the coalescence rate for galactic NS/NS binaries is $83.0^{+209.1}_{-66.1} \text{ Myr}^{-1}$ at a 95% confidence level. That corresponds to detection rates for initial LIGO and advanced LIGO of $34.8^{+87.6}_{-27.7} \times 10^{-3} \text{ yr}^{-1}$ and $186.8^{+470.5}_{-148.7} \text{ yr}^{-1}$, respectively. The detection rates are calculated assuming that initial and advanced LIGO can detect NS/NS inspirals out to $\simeq 20 \text{ Mpc}$ and $\simeq 350 \text{ Mpc}$, respectively.

The coalescence rates of NS/BH systems and BH/BH are more uncertain as no such systems have yet been discovered.

1.4.2 Supernovae

Supernovae are capable of outshining entire galaxies in the electromagnetic domain. If only a small fraction of the explosion energy is converted into gravitational waves then they could be important sources for advanced detectors. Initial detectors are not likely to detect gravitational waves from extra-galactic supernovae since the signals would be too weak. However, the complexity in modelling these sources still leaves doubt about their gravitational wave signatures and strengths. The current consensus is that first generation gravitational wave detectors would require a galactic event (or possibly from the Magellanic Clouds) for detection. For a recent review on supernova physics, see [27]. For reviews on gravitational wave emission from supernovae, see [28] and [29].

1.4.3 Gamma-ray bursts

Gamma-Ray Bursts (GRBs) are intense flashes of gamma-rays that appear to be the most energetic explosions in the universe. They were first observed in 1969 by military Vela satellites meant to monitor nuclear weapon testing from space [30]. GRBs tend to be divided into two classes, ‘long’ and ‘short’, depending on the duration of the bursts. There is convincing evidence that long GRBs are related to supernovae where massive black holes are created (collapsar model, see [31]). If these are indeed the physical origins of these GRBs then the gravitational wave signals from the long bursts could be comparable to those from supernova. Short GRBs, on the other hand, may come from the merger of compact objects. These could emit significant amounts of gravitational waves detectable by advanced detectors [32].

1.4.4 Neutron stars

Gravitational waves can be emitted from rotating neutron stars due to asymmetric mass distribution, normal modes of oscillation of a fluid core, or free precession. Figure 1.4 shows upper limits on the gravitational wave emission from two known pulsars, Vela and the Crab, if we assume that the observed loss in rotational energy is due entirely to gravitational waves. We provide further details and calculate this spindown based upper limit for other pulsars in Section 3.3.

In this thesis, searches are performed to extract signals that would be generated from triaxially shaped neutron stars [33]. The asymmetry could arise from strains in a solid crust left over from its formation [34]. The triaxiality could also be supported by large internal magnetic fields [35]. The black dashed lines on the right hand side of Figure 1.4 give the gravitational wave strengths for equatorial ellipticities of 10^{-7} and 10^{-6} for pulsars at a distance of 10 kpc. More detail on the gravitational wave amplitude and waveform from these systems is given in Chapter 2.

Young, hot, fluid neutron stars may emit copious amounts of gravitational waves due to Chandrasekhar-Friedman-Schutz (CFS) instabilities. CFS instabilities can occur for certain normal modes in rotating neutron stars (see [36] for a review of relativistic rotating stars). These modes can become unstable because of the emission of gravitational waves. Instead of damping them, the gravitational wave emission can actually increase the amplitude of the modes. In recent years, much interest has been given to r -modes as a potential source of gravitational waves from rapidly rotating young hot neutron stars where the saturation amplitude of the waves may be large [37]. The gravitational wave signals from r -modes would be at $\sim \frac{4}{3}$ times the rotation frequency of the neutron star [38].

Perhaps the most likely source of gravitational waves from neutron stars are from sources accreting matter from a companion star, such as Low-Mass X-Ray Binaries (LMXBs). For these systems, the accreting material provides a clear mechanism for how an asymmetry

could arise in the neutron star crust (for example, see [39]). The continuous accretion of matter may also be able to keep the neutron star hot enough for r -modes to be important not only at a young age. There is the added intriguing fact that most LMXBs seem to have rotation frequencies that cluster near ~ 450 Hz, which is far below the maximum rotation rate for neutron stars. This is interesting since without a mechanism limiting neutron star rotation rates one might expect to find some examples with rotation rates nearly breaking up the star. It has been suggested that the clustering of LMXB frequencies is due to an equilibrium forming between angular momentum gained by accretion and that lost via gravitational wave emission [39]. Figure 1.4 shows the corresponding strengths of gravitational waves from known LMXBs if these systems are in such an equilibrium. If the spin-up torque is being balanced by a gravitational wave emission torque then signals from Sco-X1 could be detected with advanced interferometers.

Recent analysis of data from the 76 m Lovell radio telescope at Jodrell Bank has found evidence that pulsar B1828–11 is freely precessing [40]. This particular pulsar, with a period of 405 ms, is rotating too slowly to emit detectable gravitational waves. But the discovery of this system begs the question of whether such systems could be important sources of gravitational waves. Detailed analyses in [41] have looked carefully at the gravitational waves emitted from freely precessing neutron stars and concluded that generally the signals will be too low to be detectable by even second generation detectors.

1.4.5 Stochastic background

There could be a stochastic background of gravitational waves due to i) a cosmological background left over from the early universe analogous to the cosmic microwave background, and/or ii) a number of unresolved gravitational waves signals from astrophysical sources. The spectrum of stochastic gravitational waves is generally characterised by $\Omega_{\text{GW}}(f)$ which is the gravitational wave energy density per logarithmic frequency di-

vided by the critical energy density to close the universe. Specifically, we have

$$\Omega_{\text{GW}}(f) = \frac{f}{\rho_c} \frac{d\rho_{\text{GW}}}{df}, \quad (1.19)$$

where ρ_c is the critical energy density to close the universe⁴ [42]. If the standard inflation theory is correct then $\Omega_{\text{GW}} \leq 10^{-15}$ and the gravitational waves would be too weak to be detectable even with advanced interferometers [43]. The black dotted lines in Figure 1.4 show the different signal strengths assuming a cross-correlation of four months of data and isotropic waves. For a good review on cosmological sources of stochastic gravitational waves, see [44].

1.4.6 Cosmic strings

Cosmic strings may produce gravitational waves in two ways that could be detectable: i) they could contribute to a stochastic background and, ii) they could emit individual burst ‘cusps’ of gravitational waves. The idea that cosmic strings could emit gravitational radiation has been around for a long time [45] (and more recently in [46]) but has recently received increased attention. In [47] it has been argued that certain cosmic strings could produce gravitational bursts strong enough to be detectable by first generation gravitational wave detectors.

1.4.7 Unknown sources

None of the astrophysical or cosmological sources described in this section are ‘guaranteed’ to be detected by the first generation of detectors. Indeed, the current consensus among scientists is that the first generation of detectors *could* possibly make a detection while advanced detectors *should* probably see gravitational waves. But perhaps, with luck, the first discovery will come from an unexpected and/or unknown source.

⁴The critical density is $\rho_c \equiv 3c^2 H_0^2 / 8\pi G$ where H_0 is the Hubble expansion rate.

1.5 Bayesian methods in data analysis

In data analysis, inferences are drawn from the data by the application of probability and statistics. There are, however, different ways of defining the term *probability*. The frequentist, classical, methodology uses probability to mean the *expected frequency* of an event occurring. In a Bayesian framework, on the other hand, when we use the term probability we mean *degree of belief*. In essence, these two frameworks pose different questions and that must be considered when deciding which one to use in statistical problems.

In the searches presented in this thesis, we assess whether gravitational waves are emitted by known pulsars. If no gravitational waves are observed, then we want to place upper limits on emission by these objects. Since we fundamentally want to compare different models, which have varying amplitudes of gravitational waves, the choice of a Bayesian approach is quite natural for this problem. We want to use all of the data available to make the best possible inference on physical properties of these stars.

In [48] (see Appendix), we compare upper limits on gravitational wave emission from pulsar B1937+21 using both a Bayesian and frequentist approach to the analysis. The results for the Bayesian analysis are described in detail in Chapter 3. The frequentist approach used the maximum likelihood principle to derive detection statistics for the signal as described in [49]. Although we would not expect the same numerical value between the numerical results, we did find qualitative agreement between the two methodologies. While this was reassuring, in principle, the two results have completely different meanings.

The merits of Bayesian analyses in astrophysics are discussed further in [50] and specifically for gravitational wave data analysis in [51]. In the rest of this section, the basic rules of Bayesian probability theory are presented.

Two basic axioms of probability theory, the sum rule and the product rule, lead to the

foundations of Bayesian theory. The *sum rule* simply states that if we know how much we believe something is true, then we implicitly also state how much we believe it is false. It can be written in the form

$$p(\theta|I) + p(\bar{\theta}|I) = 1, \quad (1.20)$$

where $p(\theta|I)$ is the probability of θ given our background information I and $p(\bar{\theta}|I)$ is the probability of the opposite. As is usually done, we have implicitly defined probability on a scale of 0 to 1 with a larger number signifying larger probability and 1 meaning certainty of the truth of the proposition. The *product rule* relates the probability of two propositions, $p(\theta, D|I)$, with the probability of one of the proposition, $p(\theta|I)$, and the probability of the other one given the first one, $p(D|\theta, I)$. In equation form we have

$$p(\theta, D|I) = p(\theta|I)p(D|\theta, I). \quad (1.21)$$

Simple manipulation of the product rule gives rise to Bayes' Theorem. Given the interchangeability of θ and D in Equation 1.21, we get

$$p(\theta|D, I) = \frac{p(\theta|I)p(D|\theta, I)}{p(D|I)}, \quad (1.22)$$

which is known as Bayes' Theorem. So we have

$$p(\theta|D, I) \propto p(\theta|I) \cdot p(D|\theta, I) \quad \text{or} \quad (1.23)$$

$$\text{posterior} \propto \text{prior} \cdot \text{likelihood}, \quad (1.24)$$

where we have replaced $p(D, I)$ with a proportionality as it is not a function of θ . This is an important result for data analysis as it describes how to update our prior knowledge on θ from $p(\theta|I)$ to $p(\theta|D, I)$ with the use of the likelihood of data $p(D|\theta, I)$.

The second rule of Bayesian theory can be derived by considering the *sum rule* and the

product rule together. Using these rules, it is easy to show that

$$p(\theta|I) = p(\theta, \beta|I) + p(\theta, \bar{\beta}|I). \quad (1.25)$$

By considering β as a continuous parameter representing a set of mutually exclusive and exhaustive possibilities we can write

$$p(\theta|I) = \int_{-\infty}^{\infty} p(\theta, \beta|I) d\beta, \quad (1.26)$$

which is known as the *marginalisation rule*. Since these represent a set of exhaustive possibilities, we can also write a normalization requirement that

$$\int_{-\infty}^{\infty} p(\theta|I) d\theta = 1. \quad (1.27)$$

The marginalisation rule is particularly useful in data analysis in order to deal with parameters which are part of a model but of no interest to the analyst; we call these *nuisance* parameters.

The use of prior probabilities in Bayesian statistics has caused unease among some who regard them to be less objective than the methods of classical frequentist statistics. Nevertheless, the intuitive appeal of the Bayesian approach as well as the natural interpretation of the results has led us to use Bayesian reasoning for the problems we address in the following chapters. For a detailed discussion on the topic of priors and subjective probabilities, the reader is referred to a set of papers by D'Agostini, especially [52] and [53].

1.6 Outline of the thesis

The field of gravitational wave astronomy is approaching and relevant data analysis techniques are being developed to make optimal use of data from several detectors. This thesis concerns the application of new data analysis methods to recent data from the

GEO 600 and LIGO gravitational wave detectors. In the second chapter, a pipeline is developed to search for periodic gravitational wave signals from targeted sources such as pulsars. The robust method presented makes use of the known rotation phase of these objects and a Bayesian statistical interpretation is employed. In Chapter 3, we place upper limits on periodic gravitational waves from pulsar B1937+21 using data from the first science run of GEO 600 and LIGO. In Chapter 4, we use data from LIGO's second science run to expand the search to 28 isolated pulsars. Two simulated signals injected into the interferometers during the second science run are also analysed to increase our confidence in the overall pipeline. In Chapter 5, the same analysis is carried out on the LIGO data from the third science run. For one of the pulsars, B1937+21, a coherent multi-detector analysis is performed making use of S3 data from GEO 600. Finally, a brief summary of the main results is provided in Chapter 6.

Chapter 2

Bayesian methods to search for periodic gravitational waves

2.1 Introduction

In this chapter, we present a scheme designed to search for weak periodic gravitational wave signals emitted by neutron stars. Specifically, the method is ideal for sources with known sky positions and frequencies such as radio pulsars. The technique and statistical methodology are described in detail and applied to simulated data. The method has been successfully applied to real data from several science runs of GEO 600 and LIGO. These results are presented in detail in the following three chapters.

We have seen that several mechanisms have been proposed that could lead to the emission of continuous gravitational waves from neutron stars. Here we concentrate on data analysis methods for emission from non-precessing triaxial neutron stars rotating about their principal axes. For these neutron stars to be emitting gravitational waves they need to be asymmetric about their rotation axes. The gravitational wave signal will be emitted at twice the rotation frequency of the neutron star.

Since gravitational wave signals from neutron stars are expected to be weak, it is advantageous to integrate over long stretches of data (months to years). A periodic gravitational wave signal being emitted by a neutron star will be Doppler shifted due to the motion of the Earth, and will be amplitude modulated by the response of the interferometers. In the work presented here we take into account both of these effects on the measured gravitational wave signal.

Searches for periodic gravitational wave signals from neutron stars can be divided into three classes: targeted, directed, and blind. If the location and the spin parameters (frequency and frequency derivatives) of an object of interest are well known we refer to a “targeted” search for a signal. This is the case for pulsars where the location and spin parameters can be inferred directly and accurately from radio or X-ray observations. We refer to “directed” searches when these parameters are poorly known, and possibly the position is the only parameter known to any precision. For example, supernova remnants are objects for which the location is approximately known and there may or may not be information about the spin parameters. Finally, blind all-sky searches correspond to the situation where no parameters are known, not even the location of possible sources. Due to the dependence of the signal on sky location and the very large parameter space involved, such searches are typically computationally limited. Hierarchical strategies that take advantage of fast algorithms (like the Fast Fourier Transform, FFT) have been developed and must be used for blind searches. In this work, we will concentrate on a targeted search technique that was specifically developed to search for signals from known pulsars.

The fact that we can monitor radio pulsars make them particularly interesting sources to target. There are several advantages to performing targeted searches when the locations and frequencies are known in advance. First, for unknown sources we add extra parameters into the problem: right ascension, declination, frequency, and frequency derivatives. This results in a decrease in sensitivity since the statistical threshold for detection will need to be increased. For targeted searches with a small number of unknown parameters

the threshold can be lower. Another advantage of targeted searches is that we focus on a narrow frequency band of data. Due to the presence of several interfering spectral lines in interferometric data, the data characterisation problem can be reduced significantly by only analysing narrow bands, and hopefully avoiding contaminating lines. Finally, since we are only searching over a few parameters in a narrow band, targeted searches are not computationally bound. Indeed, the algorithm presented here does not even need to take advantage of fast algorithms like FFTs.

This chapter provides an end-to-end description of a robust and straightforward method to search for periodic signals from pulsars with interferometric gravitational wave detectors. In Section 2.2.2 we describe the nature of the gravitational wave signal from a pulsar and introduce the relevant physical parameters. We explain the data processing and the statistical methodology including how to incorporate coherently the data from several detectors in Section 2.3. The performance of the algorithm on simulated data is described in Section 2.4. We conclude in Section 2.5 by discussing natural extensions to this work including Markov Chain Monte Carlo methods.

2.2 Nature of the gravitational wave signal

2.2.1 Response of interferometric detectors

The response of interferometers to incoming weak plane gravitational waves can be expressed as a linear combination of the strain of each polarisation $h_+(t)$ and $h_\times(t)$ by

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t) \quad (2.1)$$

where $F_+(t)$ and $F_\times(t)$ are called the *beam pattern* functions. The beam pattern functions, also called amplitude response functions, give the sensitivity of an interferometer to each polarisation. Due to the diurnal motion of the Earth these functions are periodic with a

period of half a sidereal day. The beam pattern functions are dependent on the geometry of the interferometer and also the location and orientation of the source with respect to the interferometer. We will use the same angles as in [49] to describe the beam pattern functions here, namely: α as the right ascension of the source, δ as the declination of the source, ψ as the polarisation angle of the wave, λ as the latitude of the detector's site, Ω_r as the rotational velocity of the Earth, ϕ_r as a deterministic phase defining the position of the Earth at time $t = 0$ (so that $\phi_r + \Omega_r t$ is the local sidereal time (LST) of the detector's site¹), γ as the angle between East and the bisector of the interferometer arms, and ζ as the angle between the arms.

The explicit formulae for the beam pattern are [49]:

$$F_+(t) = \sin \zeta [a(t) \cos 2\psi + b(t) \sin 2\psi], \quad (2.2)$$

$$F_\times(t) = \sin \zeta [b(t) \cos 2\psi - a(t) \sin 2\psi], \quad (2.3)$$

where

$$\begin{aligned} a(t) = & \frac{1}{16} \sin 2\gamma (3 - \cos 2\lambda) (3 - \cos 2\delta) \cos[2(\alpha - \phi_r - \Omega_r t)] \\ & - \frac{1}{4} \cos 2\gamma \sin \lambda (3 - \cos 2\delta) \sin[2(\alpha - \phi_r - \Omega_r t)] \\ & + \frac{1}{4} \sin 2\gamma \sin 2\lambda \sin 2\delta \cos[\alpha - \phi_r - \Omega_r t] \\ & - \frac{1}{2} \cos 2\gamma \cos \lambda \sin 2\delta \sin[\alpha - \phi_r - \Omega_r t] \\ & + \frac{3}{4} \sin 2\gamma \cos^2 \lambda \cos^2 \delta, \end{aligned} \quad (2.4)$$

$$\begin{aligned} b(t) = & \cos 2\gamma \sin \lambda \sin \delta \cos[2(\alpha - \phi_r - \Omega_r t)] \\ & + \frac{1}{4} \sin 2\gamma (3 - \cos 2\lambda) \sin \delta \sin[2(\alpha - \phi_r - \Omega_r t)] \\ & + \cos 2\gamma \cos \lambda \cos \delta \cos[\alpha - \phi_r - \Omega_r t] \\ & + \frac{1}{2} \sin 2\gamma \sin 2\lambda \cos \delta \sin[\alpha - \phi_r - \Omega_r t]. \end{aligned} \quad (2.5)$$

¹LST: angle between local meridian and the vernal point.




Figure 2.1: Sensitivity of an interferometer to plus polarisation of a gravitational wave as function of direction.

For a full derivation of these expressions the reader is referred to [49] and references therein. Figure 2.1 and Figure 2.2 show the directional sensitivity of an interferometer to the plus and cross polarisations, respectively. The hypothetical interferometer would be situated at the origin of the plots with perpendicular arms along the x -axis and y -axis. Figure 2.3 shows the response of that interferometer to unpolarised gravitational waves.

We note that the response of interferometers to gravitational waves is nearly omnidirectional. Unlike an optical telescope for which we need to point toward any source of interest, gravitational wave interferometers are sensitive to a very large area of the sky at all times. This is particularly convenient for continuous sources as they are nearly always present in the data. For those signals, their locations can be determined by observing the Doppler shifts due to the motion of the Earth. On the other hand, for short duration signals, the non-directionality of the interferometers requires using times-of-arrival information from several detectors in order to determine the location of the source in the

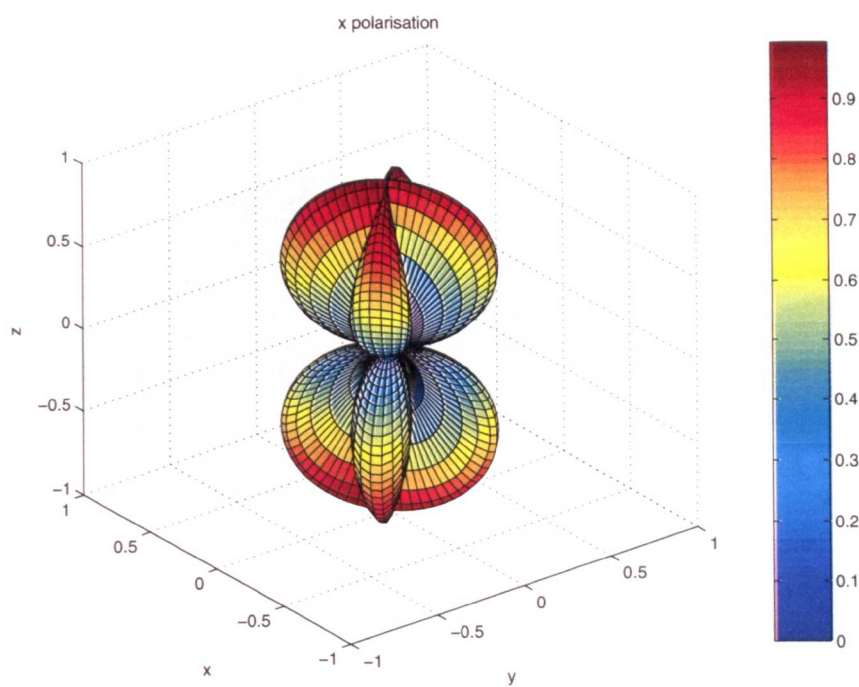


Figure 2.2: Sensitivity of an interferometer to cross polarisation of a gravitational wave as function of direction.




Figure 2.3: Sensitivity of an interferometer to an unpolarised gravitational wave as function of direction.




Figure 2.4: Model of an axisymmetric neutron star emitting gravitational waves at twice the rotation frequency. For the moments of inertia of this star we have $I_{zz} = I_{yy} \neq I_{xx}$.

sky.

2.2.2 Gravitational waves from triaxial neutron stars

A perfectly symmetric rotating sphere will not emit quadrupole gravitational waves. Here we consider the gravitational waves emitted from a rotating rigid axisymmetric body. Figure 2.4 shows an example of such a body with moments of inertia $I_{zz} = I_{yy} \neq I_{xx}$ where z is the axis of rotation. Such a body will emit gravitational waves as long as the axial symmetry is not about the rotation axis. The gravitational waves emitted by an rotating triaxial rigid body about any axis are described in detail in [33]. Here we consider the special case of a triaxial ellipsoid rotating about its principal axis. In this case, it emits gravitational waves solely at twice its rotation frequency. If the star is

precessing then there is an additional component to the gravitational wave signal at the rotation frequency and the precession frequency may need to be taken into account (but it may be small enough to be negligible). The regularity of the time-of-arrivals of signals from the majority of radio pulsars suggests that most of them are not precessing on a short time scale or maybe not at all. We will see an example pulsar in the next chapter for which the timing residuals have been less than $4 \mu\text{s}$ over several years.

The gravitational wave amplitude, h_0 , for a triaxial neutron star is described by

$$h_0 = \frac{16\pi^2 G}{c^4} \frac{I_{zz} f_0^2}{r} \epsilon, \quad (2.6)$$

where r its distance to the pulsar, I_{zz} its moment of inertia about the rotation axis, f_0 is the rotation frequency of the pulsar, and ϵ is its equatorial ellipticity of the pulsars. The equatorial ellipticity is defined as

$$\epsilon = \frac{I_{xx} - I_{yy}}{I_{zz}}. \quad (2.7)$$

The gravitational wave signal seen by the detector will be frequency modulated due to the relative motion of the Earth to the pulsar. As we have seen, the signal will also be amplitude modulated by the strain response pattern of the interferometer to each polarization. Using the model described in [33], we can describe the measured signal, $h(t)$, by

$$h(t) = \frac{1}{2} F_+(t; \psi) h_0 (1 + \cos^2 \iota) \cos 2\Psi(t) + F_\times(t; \psi) h_0 \cos \iota \sin 2\Psi(t), \quad (2.8)$$

where F_+ and F_\times are the beam patterns functions of the detectors², ψ is the polarisation angle, ι is the inclination of the pulsar with respect to the line of sight, and $\Psi(t)$ is the phase of the signal.

²In addition to the polarisation angle the beam pattern functions depend on location and several other fixed angles as described in the previous section.

2.2.3 Pulsar timing: phase of the gravitational wave signal

We choose to express the pulsar phase with respect to the solar system barycentre which is an inertial reference frame. We can model gravitational wave signal phase, neglecting high order terms, as a function of barycentric time, T , by

$$\Psi(T) = \phi_0 + 2\pi \left[f_0(T - T_0) + \frac{1}{2} \dot{f}_0(T - T_0)^2 + \frac{1}{6} \ddot{f}_0(T - T_0)^3 + \mathcal{O}(T^4) \right] \quad (2.9)$$

where ϕ_0 is the phase of the pulsar at a fiducial time T_0 , f_0 is the rotation frequency of the pulsar, \dot{f}_0 is the first frequency derivative, and \ddot{f}_0 is the second frequency derivative. The transformation between the barycentric time (T) and the topocentric time at the detector (t) can be expressed as

$$T = t + \delta t = t + \Delta_{\text{Roemer}} + \Delta_{\text{Shapiro}} + \Delta_{\text{Einstein}} + \Delta_{\text{Binary}} \quad (2.10)$$

where Δ_{Roemer} is the classical Roemer delay, $\Delta_{S\odot}$ is the Shapiro delay due to the curvature of space-time near the Sun, $\Delta_{E\odot}$ is the Einstein delay due to gravitational redshift and time dilation, and Δ_{Binary} contains corrections related to the pulsar's orbit, if applicable. More details on the magnitude of each of these timing effects is provided in Section 3.5. The second term in Equation 2.10, the Roemer delay, is the largest component and due to the motion of the Earth within the Solar System. For isolated pulsars the fifth term in Equation 2.10 is not applicable ($\Delta_{\text{Binary}} = 0$). For pulsars in binary systems this term should include all the classical and relativistic corrections for the shifts in the time-of-arrival of the signal due to the motion of the source within the binary system.

We will not consider binary pulsars in this analysis, but for more details on pulsar timing of binary systems see [54] and references therein.

2.3 Searching for signals from pulsars

Current ground-based interferometric detectors have broadband sensitivity to gravitational waves up to several kilohertz, and thus the sampling rate for data needs to be relatively high. For the GEO 600 and LIGO detectors, the sampling rate has been set to 16 384 Hz. This leads to large data sets when analyzing several months or years of data as advantageous for pulsars. For example for GEO 600, which stores data in double precision, this corresponds to $(16\,384 \text{ samples/s} \times 86\,400 \text{ s/day} \times 8 \text{ bytes/sample})$ 11.3 gigabytes of data per day, or 4.1 terabytes per year. In practice this is a gross underestimate of the actual data stored since a large amount of environmental information is also recorded.

The signal we are trying to extract in a targeted search, however, is actually contained in a very narrow frequency band. The knowledge of the spin parameters of the source from radio observations permits us to reduce the size of the data set considerably.

As explained in more detail in the next sections, we can perform a complex heterodyne of the data and then reduce the sampling rate by a factor of $\sim 10^6$ without loss of relevant information. A similar technique for data compression has been used in previous searches for periodic gravitational waves [55]. Complex heterodyning can be used to remove the expected phase evolution of a signal in the detector output leaving the processed signal near zero frequency. For a pulsar signal recorded by an interferometric detector, the only time varying component remaining in the heterodyned signal would be the response pattern of the interferometer.

Since the beam pattern of the interferometer varies on a time scale much slower than the original periodic signal, after heterodyning we can thus re-sample the data with a much reduced sampling rate. Indeed, since the period of the beam pattern of the antenna is only a sidereal day, in principle from the Nyquist theorem we could re-sample as low as twice per day. However, for practical reasons is it better to have a larger sampling rate than

this. The main reason is that the noise floor of the interferometers changes on time scales of minutes, hours, and days. In order to correctly track the changing noise floor we have found that a sampling rate of 1 sample per minute is quite adequate to deal with current data. Since we keep both the real and imaginary part of the sample for each minute, our effective bandwidth is 1/60 Hz centered on the heterodyning frequency which is the instantaneous frequency of the signal at the detector.

Once the data volume has been reduced to 1 sample per minute, we take a Bayesian approach for the analysis. In the following section, we show explicitly the form of the heterodyned gravitational wave signal in the reduced data set. This signal is dependent on four unknown parameters: h_0 (gravitational wave amplitude), ψ (polarisation angle), ϕ_0 (phase of the signal at time T_0), and ι (the angle between the angular momentum vector of the pulsar and the line of sight). We proceed by calculating the probability of the data given each set of parameters over their respective ranges. Finally, through the application of Bayes' Theorem with suitable prior probabilities and marginalisation, we obtain posterior probabilities for each of these parameters given the data collected.

2.3.1 Data processing

We take the calibrated output of a gravitational wave detector to be

$$s(t) = h(t) + n(t), \quad (2.11)$$

where $h(t)$ is a gravitational wave signal and $n(t)$ is noise. We have already seen the signal, $h(t)$, that we are expecting from pulsars in Equation 2.8. We can rewrite this equation using simple trigonometric identities³ to give

$$h(t) = e^{i2\phi(t)} A_1(t) + e^{-i2\phi(t)} A_2(t), \quad (2.12)$$

³ $\sin z = \frac{e^{iz} - e^{-iz}}{2i}$ and $\cos z = \frac{e^{iz} + e^{-iz}}{2}$

where

$$A_1(t) = \frac{1}{4}F_+(t; \psi)h_0(1 + \cos^2 \iota)e^{i2\phi_0} - \frac{i}{2}F_\times(t; \psi)h_0 \cos \iota e^{i2\phi_0}, \quad (2.13)$$

and

$$A_2(t) = \frac{1}{4}F_+(t; \psi)h_0(1 + \cos^2 \iota)e^{-i2\phi_0} + \frac{i}{2}F_\times(t; \psi)h_0 \cos \iota e^{-i2\phi_0}, \quad (2.14)$$

where $\phi(t) = \Psi(t + \delta t) - \phi_0$ as defined in Equations 2.9 and 2.10.

As we have seen, since the signal bandwidth is small compared to the detector bandwidth no information is lost by reducing the bandwidth of the data around the signal. We do this by complex heterodyning the data, low-pass filtering, and then averaging to reduce the sample rate. We will denote each averaged data point as B_k where k is an index representing the time period for each sample. These binned data points contain all of the information that we need to extract a gravitational wave signal, $h(t)$, from the data.

For the heterodyning step, we multiply the calibrated data from the interferometer by $e^{-i2\phi(t)}$ and get

$$s_{\text{het}}(t) = s(t)e^{-i2\phi(t)} = A_1(t) + A_2(t)e^{-i4\phi(t)} + n(t)e^{-i2\phi(t)}. \quad (2.15)$$

The heterodyning process removes the known phase evolution from the term A_1 . This term will still vary over the day due to the antenna response of the interferometer. The second term, A_2 , will oscillate at nearly twice the gravitational wave signal frequency.

We then apply a low-pass filter to the heterodyned data stream prior to averaging to prevent aliasing noise from outside the signal band into our reduced data set. In practice, we have used a series of third order Infinite Impulse Response (IIR) Butterworth filters for the time domain low-pass filtering. The frequency cutoffs of the IIR filters can be tuned depending on the characteristics of the data. The main requirement is to prevent spectral

disturbances from outside our signal bandwidth being aliased into our calculation of the B_k 's. We want to use a low-pass IIR filters that will attenuate noise outside our 1/60 Hz band. If we wanted to reduce our sampling rate to less than 1 sample per minute, then we could use IIR filters with a tighter cutoff. However, this would correlate the data for longer periods of time and we have seen that this can be disadvantageous due to changing noise levels. The selection of the IIR filters, and the sampling rate, ultimately depends on the stationarity of the noise and whether a narrow band is required to avoid nearby spectral lines

Finally, after applying a suitable low-pass filter, we can bin (average) the filtered data (now s'_{het}) to reduce the number of samples to

$$B_k = \frac{1}{M} \sum_{i=1}^M s'_{\text{het}}(t_i), \quad (2.16)$$

where k is the sample number, the index i is over each sample at 16 384 Hz, and M is the number of samples to average together.

The low-pass filtering and averaging procedure strongly suppress the second term in Equation 2.15. Since the term is periodic and outside our band it is attenuated by the IIR filters and the averaging in Equation 2.16.

There can be advantages in performing the heterodyning and re-sampling process described above in two steps, to prevent having to recalculate the time delay to the solar system barycentre for every sample. This can be done by first complex heterodyning the data at a fixed frequency near signal frequency and reducing the bandwidth to, for example, 4 Hz. A second heterodyne can then be performed to the data to removed the Doppler shifts due to the motion of the Earth. The advantage to this technique is that the time delay from the detector to the solar system barycenter, which can be computationally intensive, can be calculated at 4 Hz instead of 16 384 Hz. In practice, however, we have found that it is not necessary to take advantage of this technique on modern computers.

With the high frequency term in Equation 2.15 suppressed we have

$$B_k = \frac{1}{4}F_+(t_k; \psi)h_0(1 + \cos^2 \iota)e^{i2\phi_0} - \frac{i}{2}F_\times(t_k; \psi)h_0 \cos \iota e^{i2\phi_0} + n(t_k)', \quad (2.17)$$

where $n(t_k)'$ is the heterodyned and averaged complex noise in bin k . For example, if we let $M = 60 \times 16\,384$ in Equation 2.16 we get one binned data point, B_k , every minute for a detector acquiring data at 16 384 Hz. The central limit theorem tells us that we would expect the noise, $n(t_k)'$, to be well described by a Gaussian distribution. However, as we have seen, the width of this distribution may change over time as the detector sensitivity evolves.

2.3.2 Bayesian analysis

We take a Bayesian approach for the statistical analysis. We use Bayes' Theorem to calculate the posterior probability, $p(\mathbf{a}|\{B_k\})$, of a set of parameters \mathbf{a} given the binned data, $\{B_k\}$. For this problem, Bayes' Theorem states that

$$p(\mathbf{a}|\{B_k\}) = \frac{p(\mathbf{a})p(\{B_k\}|\mathbf{a})}{p(\{B_k\})}, \quad (2.18)$$

where \mathbf{a} represents the set of parameters that could produce the set of data, $\{B_k\}$, with likelihood $p(\{B_k\}|\mathbf{a})$. Our prior beliefs in our set of parameters are reflected in the prior probability term, $p(\mathbf{a})$. We will use the least informative priors for most of the parameters in their respective ranges: ϕ_0 uniform over $[0, 2\pi]$, ψ uniform over $[-\pi/4, \pi/4]$, and ι uniform in $\cos \iota$ over $[-1, 1]$, corresponding to a uniform prior per unit solid angle of pulsar orientation. Some comments on these priors are given below.

The principal range for the polarisation angle spans π rad due to the quadrupolar nature of the waves. This is clear from the factor of 2 in front of ψ in Equation 2.2 and Equation 2.3. For example, if a 'plus' mode is rotated by $\pi/2$ rad then it is the opposite of itself. However, for a given signal there are two degenerate solutions for ψ over π rad

since we cannot differentiate between a given mode and its opposite. For this reason, we have limited our range in ψ to $[-\pi/4, \pi/4]$.

It is easy to show why for the orientation angle we need to use a prior probability that is uniform in $\cos \iota$ as opposed to ι . If we consider the probability of a small interval in orientation, $d\iota$, assuming no extra information on the orientation is known, we have

$$p(\iota)d\iota = \frac{2\pi \sin \iota d\iota}{4\pi} = \frac{1}{2}\sin(\iota)d\iota, \quad (2.19)$$

where the factor of 4π can be interpreted as the surface area of a unit sphere and $2\pi \sin \iota d\iota$ as the area of a thin ring. By the transformation property of density functions⁴ we thus have that

$$p(\cos \iota) = \frac{1}{2} \sin \iota \cdot \left| -\frac{1}{\sin \iota} \right| = \frac{1}{2} \quad (2.20)$$

and since $\iota \in [0, \pi]$ we have $\cos \iota \in [-1, 1]$.

For h_0 we will use a flat prior which permits the construction of confidence intervals so that we can determine upper limits. If, on the other hand we used an improper prior, such as a Jeffreys prior ($\propto 1/h_0$) for h_0 , we would have the problem that such a prior is not normalisable. In principle the prior on h_0 should reflect all our initial beliefs on the gravitational wave strength, h_0 . But since the detectors are too insensitive to detect any currently plausible signal, such a realistic prior would overwhelm the posterior probability. We would not learn anything new from the experiment since the posterior probability distribution function (pdf) would largely resemble the prior pdf we chose.

The main caveat in using a flat prior on h_0 is the awkward question that can be asked when we set up a 95% upper limit: Do you really believe that there is 5% probability that the signal is larger than your upper limit? In a strict sense, if we are using a flat prior on h_0 , we should indeed believe what the posterior probability function tells us if we trust our choice of prior probability. However, due to other information which we are

⁴ $p(u) = p(x) \left| \frac{dx}{du} \right|$

not including in the prior like pulsar spindown, which we will see in the next chapter, and neutron star theory, we may actually believe less that a signal is present than our limit suggests (with current detector sensitivities). It is clear, however, that the flat prior on h_0 will simply give us more conservative upper limits (meaning larger) than we would get if we were using realistic priors. With a flat prior like we have chosen, the upper limits on h_0 simply reflect the sensitivity of the detectors. The main reason for not including all of our prior knowledge is simplicity and the fact that we want to focus on the new information the detectors are telling us.

When the posterior probability distribution functions (pdfs) are used to determine upper limits it is important to remember that our choice of a flat prior will favour larger values of h_0 . In fact, when using a uniform prior for h_0 , our resulting posterior distributions are functionally equivalent to normalized likelihoods. If another researcher has different prior beliefs, Bayes' formula can be applied directly on the normalised likelihoods to give a new posterior pdf. The problem with only presenting the normalised likelihoods, however, is that we cannot call the results probability limits without applying a prior (from Bayes' Theorem). That is why we prefer presenting our results as posterior pdfs with a flat prior, and then we can talk about degrees of belief of the value of the parameter h_0 . For more discussion on how to chose a prior in similar cases to this one, in frontier physics when the level of a signal may be below the sensitivity of the experiment or not there at all, see [52]. Ultimately, if there is a strong detection the choice of the prior should not play an important role in the results since the likelihood function would be sufficiently strongly peaked.

In order to calculate the likelihood function we need to have a model of the signal in the processed data. The model of the signal that we are searching for in the data set, $\{B_k\}$, is obtained by processing the original gravitational wave signal $h(t)$ in the same way that we processed the data to give

$$y(t_k; \mathbf{a}) = \frac{1}{4} F_+(t_k; \psi) h_0 (1 + \cos^2 \iota) e^{i2\phi_0} - \frac{i}{2} F_\times(t_k; \psi) h_0 \cos \iota e^{i2\phi_0}, \quad (2.21)$$

where \mathbf{a} is again the parameter vector representing the four unknown parameters h_0, ι, ψ, ϕ_0 . We note that the model is complex and that the only time varying component is the antenna pattern of the interferometers. As we have seen, the Nyquist frequency for this signal is well below our reduced sampling rate of one B_k per minute. In the following two sections we will present two different ways of evaluating the likelihood function depending on whether the variance of the data is known or unknown.

2.3.3 Gaussian model – known variance

Here we give the expression for the likelihood function assuming that we know the variance of the data independently. We assume that the data comprises a signal, $y(t_k)$, which is embedded in Gaussian noise $N(\mu, \sigma)$ of known variance σ_k , that is

$$B_k = y_k + N(0, \sigma_k). \quad (2.22)$$

If the set of B_k are independent, then the likelihood function is simply the product of the individual measurements. In this case, the the likelihood of the data is simply the product of n bivariate normal distributions where n is the number of data points, B_k . Note that the distribution is bivariate because the data are complex numbers after the heterodyning process. The real and imaginary parts of the B_k 's have independent noise though their variances should be of similar amplitude. The expression for the likelihood of the set of data $\{B_k\}$ is given by

$$p(\{B_k\}|\mathbf{a}, \{\sigma_k\}) = (\sqrt{2\pi}\sigma_k)^{-2n} \exp\left(-\sum_{k=1}^n \frac{|B_k - y_k|^2}{2\sigma_k^2}\right), \quad (2.23)$$

where y_k is our model from Equation 2.21, n is the total number of data points B_k , and σ_k is the standard deviation of the noise of each B_k . In this expression, we are assuming that the variance is the same for the real and imaginary parts of the data.

This Gaussian model for the likelihood was used for the first GEO 600 and LIGO analysis

for signals from pulsar B1937+21 (see Chapter 3). As will be explained in detail in Chapter 3, the noise level σ_k was estimated for each B_k by using a 4 Hz band of data around the signal frequency. In the S1 analysis we deemed the uncertainty in the point estimate of σ_k was small enough to be ignored. So we made the approximation that the level of the noise could be taken as a known value. With a 4 Hz bandwidth, if we assumed that the data were stationary for one minute then we would have 240×2 points contributing to our estimate of the variance. Using such a large bandwidth around the signal frequency worked well as long as there were no spectral features distorting the noise floor near the frequency of the signal. This was the problem we encountered when we tried to search for signals from 28 different pulsars in S2 (see Chapter 4).

Although at 1284 Hz (for pulsar B1937+21) there were no significant spectral features in the 4 Hz band around the signal frequency, the same was not true for a many of the other pulsars we analysed in S2 (see Chapter 4). In order to deal with this we decided to reduce the bandwidth from which the noise was being estimated. While this did remove contaminating lines, it also meant that it was necessary to assume stationarity for a longer period to have the same confidence in the noise estimate. Specifically, we moved to calculating the noise floor from a 4 Hz to a 1/60 Hz band around the Doppler shifted signal frequency. We also assumed that the noise level would be stationary for 30 minutes. So we had only 30 points from which to estimate the noise floor for each B_k . We could no longer ignore the uncertainty in our estimates of σ_k . This was taken into account by treating σ_k as another unknown parameter as described in the next section.

2.3.4 Gaussian model – unknown variance

In the previous section, we were evaluating the likelihood of the data assuming that the noise level, σ_k , was known for each B_k . Generally, however, the noise level in the B_k 's may not be known in advance and must be estimated from the data. If σ_k is estimated from the data itself, there can be non-negligible uncertainty in its value if only a few data

values contribute to the estimate. We can correctly take account of this uncertainty by treating the noise level as another nuisance parameter. We will begin by calculating the likelihood of a subset of m_j consecutive data points from $\{B_k\}$ which have a constant noise level σ_j . Once we have that expression, we will calculate the global likelihood simply using the product rule assuming that each segment of data is independent. We will again define n to be the total number of data points B_k and let M be the number of segments of data that we have assumed have the same noise level, so that

$$n = \sum_{j=1}^M m_j. \quad (2.24)$$

Using Bayes' theorem, we can rewrite the likelihood of the j th subset of data as

$$\begin{aligned} p(\{B_k\}_j | \mathbf{a}) &\propto \int_0^\infty p(\{B_k\}_j, \sigma_j | \mathbf{a}) d\sigma_j \\ &\propto \int_0^\infty p(\sigma_j | \mathbf{a}) p(\{B_k\}_j | \mathbf{a}, \sigma_j) d\sigma_j, \end{aligned} \quad (2.25)$$

where $p(\sigma_j | \mathbf{a})$ states the prior knowledge of the noise floor and the likelihood is the same as Equation 2.23. As σ_j is a scale parameter, the least informative prior is uniform with respect to $\log(\sigma_j)$, that is

$$\begin{aligned} p(\sigma_j | \mathbf{a}) &\propto \frac{1}{\sigma_j} \quad (\sigma_j \geq 0), \\ &= 0 \quad (\sigma_j < 0). \end{aligned} \quad (2.26)$$

The conclusions would be essentially unchanged if a uniform prior was used instead of a Jeffreys' prior. This will become clear later in this section. The procedure followed here to integrate out the variance is the same as done by Sivia in [56]. Here we assume that the σ_j associated with each subset $\{B_k\}_j$ is constant for m_j samples. In other words, we assume that the noise level of the interferometer, in a narrow frequency band around the gravitational wave signal, is stationary for m_j samples. In practice, it is very useful to allow the noise floor to change between each subset $\{B_k\}_j$ of data. This allows us to

dynamically track the noise floor seen in real interferometric data which can be changing due to daily variations such as temperature, tides, etc. The length of the period, m_j , over the which the data is assumed to be stationary can also be adjusted to reflect the instrumental performance.

Using Equation 2.25, the prior from Equation 2.26, and a Gaussian likelihood, we can write that the likelihood of the subset $\{B_k\}_j$ of constant noise σ_j is

$$p(\{B_k\}_j|\mathbf{a}) \propto \int_0^\infty \frac{1}{\sigma_j^{2m_j+1}} \exp\left(\frac{-1}{2\sigma_j^2} \sum_{k=1}^{m_j} |B_k - y_k|^2\right) d\sigma_j. \quad (2.27)$$

We can solve this integral by making the following substitutions

$$\begin{aligned} u^2 &= \frac{\sum |B_k - y_k|^2}{2\sigma_j^2}, \\ du &= \frac{\sqrt{2} \sum |B_k - y_k|^2}{\sigma_j^2} d\sigma_j, \end{aligned} \quad (2.28)$$

which we can re-arrange as

$$\begin{aligned} \frac{1}{\sigma_j} &= \frac{\sqrt{2}u}{\sqrt{\sum |B_k - y_k|^2}}, \\ \frac{d\sigma_j}{\sigma_j^2} &= \frac{1}{\sqrt{2} \sum |B_k - y_k|^2} du, \end{aligned} \quad (2.29)$$

which leads to

$$p(\{B_k\}_j|\mathbf{a}) \propto 2^{m_j} \left(\sum |B_k - y_k|^2\right)^{-m_j} \int e^{-u^2} u^{2m_j+1} du, \quad (2.30)$$

$$p(\{B_k\}_j|\mathbf{a}) \propto 2^{m_j-1} \left(\sum |B_k - y_k|^2\right)^{-m_j} m_j!. \quad (2.31)$$

We can drop the constants and the solution is therefore

$$p(\{B_k\}_j|\mathbf{a}) \propto \left(\sum |B_k - y_k|^2\right)^{-m_j}, \quad (2.32)$$

which approximates a Student's t -distribution with $2m_j - 1$ degrees of freedom as is




Figure 2.5: Gaussian distribution (solid line), Student's t -distribution with 5 degrees of freedom (dashed line), and Student's t -distribution with 30 degrees of freedom (dotted line).

shown below. The overall shape of the t -distribution resembles a Gaussian distribution except a bit lower and wider. As the number of degrees of freedom $2m_j - 1$ increases, the t -distribution approaches the Gaussian distribution. Figure 2.5 shows a Gaussian distribution along with two Student's t -distributions with 5 and 30 degrees of freedom.

Recall that the likelihood derived in Equation 2.32 is for a set of m_j data points B_k with the same σ_j . The joint likelihood of all the M stretches of data, taken as independent, is therefore

$$p(\{B_k\}|\mathbf{a}) \propto \prod_j^M p(\{B_k\}_j|\mathbf{a}). \quad (2.33)$$

We note again that both the noise levels, σ_j , and the length of the intervals of constant noise, m_j , can change to reflect the performance of the detectors.

It is straightforward to show that the likelihood expression in Equation 2.32 approximates a Student's t -distribution. The Student's t -distribution applies to estimating a constant signal in the presence of noise with unknown variance. The standard expression for a

Student's t pdf is

$$f_r(t) = \frac{\Gamma\left[\frac{1}{2}(r+1)\right]}{\sqrt{r\pi}\Gamma\left(\frac{1}{2}r\right)\left(1 + \frac{t^2}{r}\right)^{(r+1)/2}}, \quad (2.34)$$

which can be reduced to

$$f_r(t) \propto \left(1 + t^2/r\right)^{-(r+1)/2}, \quad (2.35)$$

where r is the number of degrees of freedom. If we let $n = (r + 1)/2$, we can rewrite Equation 2.35 as

$$f_r(t) \propto \left(1 + t^2/(2n - 1)\right)^{-n}. \quad (2.36)$$

A little algebra and one assumption will make it obvious that Equation 2.32 agrees with Equation 2.36. If we take y_k as constant over the M segments of data then the sum in Equation 2.32 can be rewritten as

$$\sum_{k=1}^{m_j} |B_k - y_k|^2 = \frac{t^2}{2m_j - 1} + V, \quad (2.37)$$

where

$$t^2 = m_j(2m_j - 1) |\bar{B}_k - y_k|^2, \quad (2.38)$$

and

$$V = \sum_{k=1}^{m_j} |B_k - \bar{B}_k|^2. \quad (2.39)$$

\bar{B}_k is the complex number that is the average of the elements in the set $\{B_k\}_j$.

By substituting Equation 2.37 in Equation 2.32 we get

$$p(\{B_k\}_j | \mathbf{a}) \propto \left(\frac{t^2}{2m_j - 1} + V \right)^{-m_j} \quad (2.40)$$

$$\propto \left(1 + \frac{t^2}{(2m_j - 1)V} \right)^{-m_j}, \quad (2.41)$$

which is a Student's t -distribution with $2m_j - 1$ degrees of freedom as noted earlier. It is worth noting, however, that we have made the assumption that y_k was constant over each

j interval. In reality, the model y_k will vary due to the beam pattern of the interferometer. However, if the M segments of data are 30 minutes long, as applied in the S2 analysis (see Chapter 4), this effect is relatively small. For the S2 analysis, only about 12% of the data was rejected due to using blocks of 30 minutes.

2.3.5 Combining data from a network of detectors

Several gravitational wave detectors are currently collecting data. Ideally, we should be able to use the observations from all detectors in a coherent manner in order to draw the best possible inference about the source parameters. In a Bayesian analysis, as presented here, all observations enter via the likelihood function. Assuming that the data from each interferometer are independent, by the product rule the global likelihood is simply the product of the individual likelihoods. For example, by combining observations from GEO 600 and the three LIGO interferometers, we would get

$$p(\{B_k\}_{\text{Joint}}|\mathbf{a}) = p(\{B_k\}_{\text{GEO}}|\mathbf{a}) p(\{B_k\}_{\text{H1}}|\mathbf{a}) p(\{B_k\}_{\text{H2}}|\mathbf{a}) p(\{B_k\}_{\text{L1}}|\mathbf{a}), \quad (2.42)$$

where we denote H1 as the 4 km Hanford interferometer, H2 as the 2 km Hanford interferometer, and L1 as the 4 km Livingston interferometer.

This likelihood embodies all we believe we know about the values of the parameters, optimally combining the data from all the interferometers in a coherent way. Note that the observation periods can be different and so can the sensitivity of the detectors. For detectors with very different sensitivities, this will closely approximate the result from the most sensitive instrument.

2.3.6 Marginalisation over parameters and setting upper limits

The full 4-dimensional posterior pdf contains the complete result of our analysis but is difficult to interpret directly. Using the rules of probability theory we can marginalise over nuisance parameters to focus on the parameter of interest. The marginal distribution can be viewed as a weighted average of all the distributions of one parameter given all the possible combinations of the other parameters. Specifically, for the four unknown parameters in our model, we have

$$p(h_0|\{B_k\}) \propto \iiint p(\{B_k\}|\mathbf{a}) p(h_0) p(\phi_0) p(\psi) p(\cos \iota) d\phi_0 d\psi d\cos \iota, \quad (2.43)$$

$$p(\phi_0|\{B_k\}) \propto \iiint p(\{B_k\}|\mathbf{a}) p(\phi_0) p(h_0) p(\psi) p(\cos \iota) dh_0 d\psi d\cos \iota, \quad (2.44)$$

$$p(\psi|\{B_k\}) \propto \iiint p(\{B_k\}|\mathbf{a}) p(\psi) p(\phi_0) p(h_0) p(\cos \iota) d\phi_0 dh_0 d\cos \iota, \quad (2.45)$$

$$p(\cos \iota|\{B_k\}) \propto \iiint p(\{B_k\}|\mathbf{a}) p(\cos \iota) p(\phi_0) p(\psi) p(h_0) d\phi_0 d\psi dh_0, \quad (2.46)$$

where the integrals are performed over the full ranges of the nuisance parameters and the pdfs must be properly normalized. If a signal is detected, these results are very useful for extraction of the individual parameters. We can also easily calculate the joint pdf for any two parameters. For example, we can calculate

$$p(h_0, \phi_0|\{B_k\}) \propto \iint p(\{B_k\}|\mathbf{a}) p(h_0) p(\phi_0) p(\psi) p(\cos \iota) d\psi d\cos \iota, \quad (2.47)$$

$$p(h_0, \psi|\{B_k\}) \propto \iint p(\{B_k\}|\mathbf{a}) p(\phi_0) p(h_0) p(\psi) p(\cos \iota) dh_0 d\cos \iota, \quad (2.48)$$

$$p(h_0, \cos \iota|\{B_k\}) \propto \iint p(\{B_k\}|\mathbf{a}) p(\psi) p(\phi_0) p(h_0) p(\cos \iota) d\phi_0 dh_0, \quad (2.49)$$

These 2-dimensional slices in the posterior pdf are useful for study of correlations between parameters. Similar slices can be calculated for every other combination of parameters.

Even without a detection, placing upper limits on h_0 can be physically interesting as we are essentially constraining the equatorial ellipticity of the neutron star. This topic will be considered in more detail in the next chapter. We define the upper limit of h_0 bounding

95% of the cumulative probability (from $h_0 = 0$) as the value h_{95} that satisfies

$$0.95 = \int_{h_0=0}^{h_{95}} p(h_0|\{B_k\})dh_0. \quad (2.50)$$

In a similar way upper limits with different levels of confidence, 68%, 90%, or 99%, can be constructed. Note that such a limit can be placed on h_0 even if the pdf is peaked off zero and a strong signal is detected.

2.4 Performance of algorithm on simulated data

In this section we show the performance of the Bayesian part of the analysis on simulated data. Several sanity checks are provided to make sure that our software has no obvious coding bugs. The checks use generated data sets with Gaussian noise with and without simulated signals injected. It is equally important to study the behavior of the method when no signal is present since this will be the main use of this technique until gravitational waves are discovered. We begin by comparing the results between using the Gaussian and Student's t likelihoods. Then we confirm that the width of the posterior distributions behaves as expected with changing observation times and noise levels. We also examine the coherent analysis of artificial data from several detectors and we estimate parameters by generating artificial signals. We comment on the correlation between h_0 and $\cos \iota$ and the information that gives us on pulsar orientation even when h_0 is small due to our flat prior on h_0 . Finally, we show the variability of the posterior pdfs when the same experiment is repeated several times with independent noise.

For all these checks we start with sets of B_k which are generated at a rate of 1 per 60 seconds so that

$$B_k = y_k + n_k \quad (2.51)$$

where y_k is the signal and n_k are independent samples of Gaussian noise. For the signals

injected in this section we assume that the model, y_k , is for a signal from the same direction in the sky as pulsar B1937+21 with a right ascension of $19^h 40^m$ and a declination of $+21^\circ 34'$. As we have seen, the direction of the signal will affect the response of the interferometer. As an example, the beam pattern of pulsar B1937+21 over the course of the first science run is given in Figure 3.5 (next chapter). Note that since we are starting at the stage where we have sets of B_k 's, we assume the Doppler shift has already been removed in the heterodyning process.

2.4.1 Comparison of Gaussian and Student's t likelihoods

Since the likelihood using the Student's t -distribution is wider than the Gaussian distribution, we would expect this to be reflected in the posterior distribution. In order to compare the posterior pdfs from the two likelihoods we have conducted three experiments averaging the marginalised posterior pdfs for h_0 for 1000 data sets using three different likelihoods as explained below. Each data set consisted of 1440 B_k 's at a rate of one per minute equivalent to one day of data. The beam pattern used in the likelihoods function assumed that the data was from a detector at the location of GEO 600.

For each data set the B_k 's were generated from a Gaussian distribution with $\mu = 0$ and $\sigma = 1$ with no signals injected. The experiments were performed using the following three likelihood functions:

1. Gaussian likelihood function given by Equation 2.23,
2. Student's t likelihood given by Equation 2.32 and Equation 2.33 with $m_j = 5$ and thus 9 degrees of freedom,
3. Student's t likelihood given by Equation 2.32 and Equation 2.33 with $m_j = 30$ and thus 59 degrees of freedom.

The averaged posterior pdfs, for $p(h_0|\{B_k\})$, are shown in Figure 2.6 with the solid line




Figure 2.6: Comparison of average posterior marginal h_0 pdfs using the Student's t and Gaussian likelihoods. The solid line represents the average $p(h_0|\{B_k\})$ using the Gaussian likelihood, the dotted line represents the average $p(h_0|\{B_k\})$ using the Student's t likelihood with 59 degrees of freedom, and the dashed line represents the average $p(h_0|\{B_k\})$ using the Student's t likelihood with 9 degrees of freedom.

representing the Gaussian likelihood and the dotted and dashed lines representing the Student's t likelihood with 9 and 59 degrees of freedom, respectively. As expected the Gaussian likelihood and the Student's t likelihood with 59 degrees of freedom essentially give the same result with the latter only insignificantly wider. For the Student's t likelihood with 9 degrees of freedom the tails are significantly broader. In the remainder of this chapter, we will use the Student's t likelihood with 59 degrees of freedom.

2.4.2 Effects of changing the observation period

We expect the width of the posterior distributions to decrease approximately as a function of the square root of observation time. Here we described a test that was carried out comparing the marginalised posterior pdfs for h_0 using data sets of 1 day and 10 days of length. A total of four data sets were analysed with these two observation times. Two data sets had 1 day of data, one of which only contained noise and the other noise plus a




Figure 2.7: The two solid lines represent the posterior density function obtained when there is no signal using observation periods of 1 and 10 days. The two dashed lines represent the same analysis when a signal is injected at $h_0 = 0.5$. In each case, the data set with 10 days of data is the pdf with the smallest width.

signal. The two other data sets had 10 days of data with the same set up. In all cases the noise was generated from a Gaussian distribution with $\mu = 0$ and $\sigma = 1$ and the injected signal had parameters of $h_0 = 0.5$, $\psi = 0^\circ$, $\phi_0 = 0^\circ$, and $\cos \iota = 0$.

Figure 2.7 shows $p(h_0|B_k)$ for the 4 sets of data. The two solid lines represent Gaussian noise only with no signal injected, $h_0 = 0$. The two dashed lines represent the cases where the signal was injected into the noise, with $h_0 = 0.5$.

The wider solid line corresponds to the 1 day data set and the narrower line corresponds to the 10 days set. This agrees, at least qualitatively, with what we expect. The matching 95% upper limits on are $h_0 < 0.58$ and $h_0 < 0.16$ giving a reduction in h_{95} of 3.6 with 10 times more data. Since we expect the width to decrease as the square root of time, we would expect this factor to be $\sqrt{10} = 3.2$. The discrepancy can be explained due to the independent noise in each set. In the next section we conduct a similar experiment and avoid this uncertainty by averaging several posteriors.

The strongly peaked dashed line in Figure 2.7 represents the posterior on h_0 with the 10 days data set and the broader dashed line represents the same result for the 1 day data set. The location of the peak is not exactly at $h_0 = 0.5$ but within the width of the pdf in both cases. Using only one day of data the best fit is $h_0 = 0.24$ with 95% confidence that $0 < h_0 < 0.73$. While 10 days of data improves our inference to a best fit of $h_0 = 0.41$ with 95% confidence that $0.27 < h_0 < 0.57$. These intervals were constructed by choosing the 95% interval with the smallest range in h_0 . Clearly these numbers are only relevant for this particular set of data and would vary with different noise. But again, at least qualitatively the results are as expected.

Note that the widths of the posterior pdfs for h_0 are larger than what we would expect if we were simply estimating the variance of the average of the data. This extra uncertainty comes from the antenna pattern and the marginalisation over the nuisance parameters.

2.4.3 Effects of changing the noise level

We expect the width of the posteriors to be directly proportional to the level of the noise. Here we verify this hypothesis by analysing several sets of data with varying levels of noise. We also let the level of the noise change for each 30 minute stretch. This is very similar to comparing the results using data sets for different time intervals.

We compare three different scenarios and in each case the data sets correspond to 10 days of observations. The first data set contains Gaussian noise with $\mu = 0$ and $\sigma = 1$. For the second data set, the noise level alternates each 30 minutes between $\sigma = 10$ and $\sigma = \frac{1}{2\sqrt{2}}$. For the third data set, the noise level varies each 30 minutes between $\sigma = 100$ and $\sigma = \frac{1}{5\sqrt{2}}$. Two time series plots showing representative stretches of data from the first and second sets are shown Figure 2.8.

For this test, we repeated and averaged the posterior pdfs for 100 generations of the data sets described above. The average marginalised posterior pdfs for h_0 are shown in

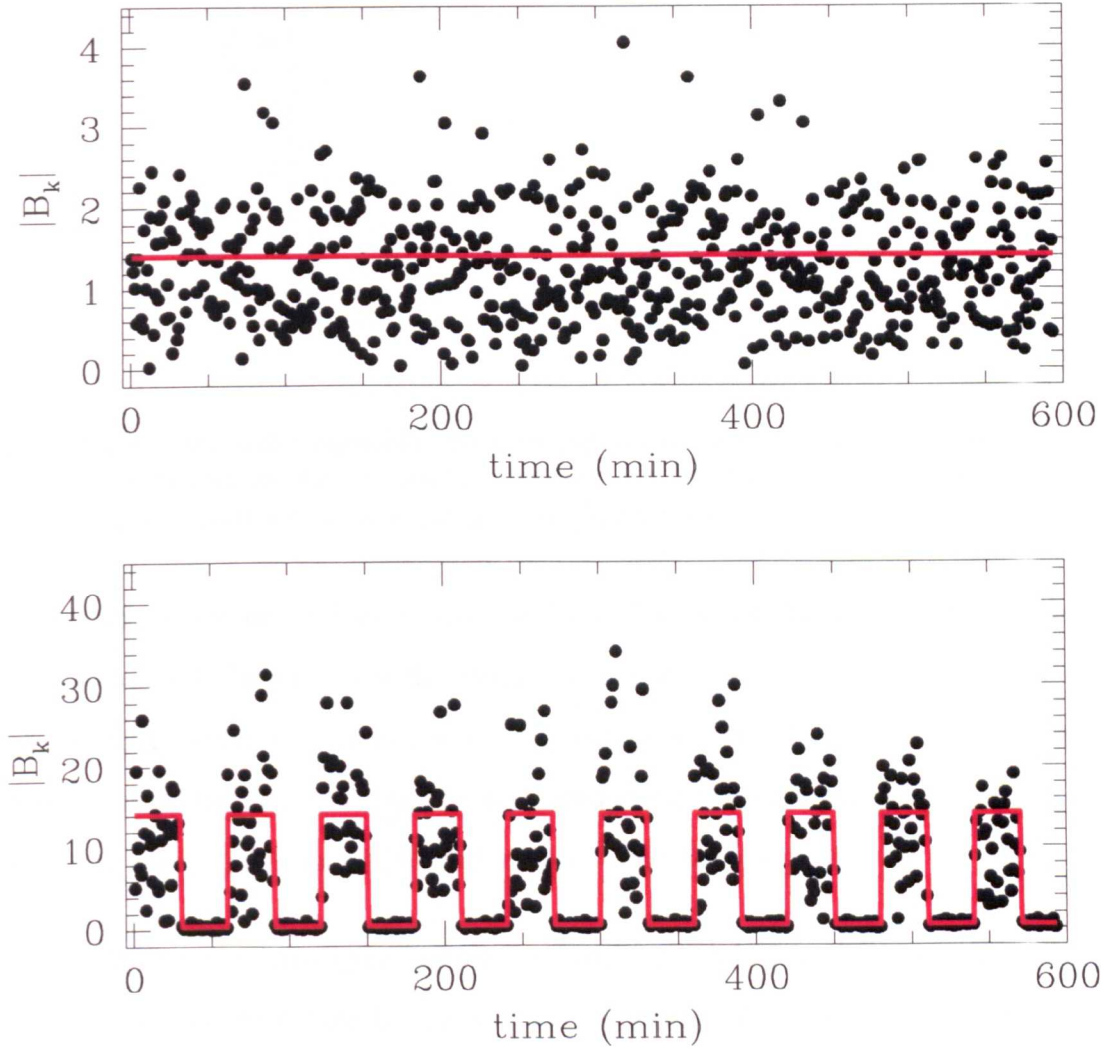


Figure 2.8: The black dots are the $|B_k|$ and the red line represents the variance of the B_k 's. The top figure represents the first data set with constant variance and the bottom figure represents the second data set with alternating noise levels.




Figure 2.9: Averaged marginalised posterior pdfs for h_0 for three scenarios: constant unit variance (solid line), alternating noise level between $\sigma = 10$ and $\sigma = \frac{1}{2\sqrt{2}}$ (dashed line), and alternating noise level between $\sigma = 100$ and $\sigma = \frac{1}{5\sqrt{2}}$ (dotted line).

Figure 2.9. Using the 66% upper limit on h_0 to characterise the width of the pdfs we have $h_0 < 0.095$ for the case with constant noise with $\sigma = 1$, $h_0 < 0.047$ for the noise alternating between $\sigma = 10$ and $\sigma = \frac{1}{2\sqrt{2}}$, and $h_0 < 0.019$ for the noise alternating between $\sigma = 100$ and $\sigma = \frac{1}{5\sqrt{2}}$. So compared to the first case the second case has a narrower width by a factor of 2.02 and the third case a factor of 5.05.

How well do these widths agree with what we would expect? In the cases with alternating noise levels, about half of the data has very low sensitivity. We can assume that half of the data will not play a significant part in the posterior compared to the more sensitive half. Thus compared to the first case with constant noise the two later cases with essentially have half of the observation period. This reduces the sensitivity by a factor of $\sqrt{2}$. Thus compared to the first case with constant noise, we would expect the two sets with alternating noise levels to have widths about a factor of 2 and 5 narrower from the variance of their sensitive halves. There is very good agreement between the results and the predictions.

name	h_0	$\cos \iota$	ϕ_0 (degrees)	ψ (degrees)
A	0.25	0.0	180	0
B	0.25	1.0	180	0

Table 2.1: Source parameters for both artificial signals extracted from data.

2.4.4 Combining results from several detectors

An advantage to our Bayesian approach to this problem is the simplicity with which we can combine data from several detectors. As described in Section 2.3.5 the combined posterior distribution is simply the product of the individual likelihoods from each detectors multiplied by the priors. For illustration, we generated 4 sets of data with Gaussian noise ($\mu = 0$ and $\sigma = 1$) and assumed that they came from GEO, L1, H1, and H2. With the data from the four IFOs having the same sensitivity we would expect the coherent results to be approximately $\sqrt{4}$ times tighter than the individual results. The four posterior pdfs for each detector as well as the joint multi-detector posterior pdf for h_0 are shown in Figure 2.10. The individual 95% upper limits are 0.15 for GEO, 0.16 for H1, 0.18 for H2, and 0.13 for L1 giving an average of 0.155. The combined 95% upper limit, on the other hand, is 0.08, which is indeed approximately a factor of 2 better than the average of the limits from the individual detectors. Since the noise level is the same in each interferometer, adding extra detectors is approximately the same as having a larger observation period. The beam patterns of each detector are different, however, and this can lead to small differences in sensitivity.

2.4.5 Parameter estimation

Here we inject two signals into noise and attempt to extract the signal parameters via the marginalised posterior pdfs. The parameters of the two injected signals are given in Table 2.1. In each case we generated data sets consisting of 10 days of data with Gaussian noise with $\mu = 0$ and $\sigma = 1$.

For both signals the phase was chosen to be 180° , and the polarisation angle was chosen




Figure 2.10: The solid line represents the joint marginalised posterior pdf for h_0 using the data from GEO, L1, H1, and H2. The dashed lines are the pdfs of the individual interferometers.

to be 0° . The only parameter differing between both signals is $\cos \iota$ which is set to zero for signal A, and 1.0 for signal B. These are two of the extreme orientations of the pulsar.

Signal A was injected in four data sets mimicking data from GEO 600, H1, H2, and L1. Figure 2.11 shows the marginalised posterior pdfs for each of the parameters. The solid lines represents the marginalised pdfs for data from GEO 600 whereas the dotted line represents the marginalised multi-detector pdfs. For each case, our resulting posterior pdfs are consistent with the injected parameters which are denoted by the vertical dashed lines.

Figure 2.12 shows the marginalised posterior pdfs for Signal B using only simulated data from GEO 600. Again, each of the pdfs is statistically consistent with the parameters of the injected signal.

The shapes of the pdfs for Signal B are quite different from those of Signal A. This is due to the correlations between the parameters that arise when $\cos \iota$ is not zero. Specifically, h_0 and $\cos \iota$, as well as ψ and ϕ_0 , are strongly anti-correlated. These can be appreciated




Figure 2.11: Marginalised posterior pdfs for Signal A for each source parameter. The dotted vertical line show the location of the actual injected signal. Solid line represents posterior pdfs from simulated data from one detector (GEO 600) and the dotted line represents posterior pdfs using data from four detectors.




Figure 2.12: Marginalised posterior pdfs for Signal B for each source parameter. The dotted vertical line show the location of the actual injected signal.




Figure 2.13: Contours of constant probability density for joint posterior probability of $\cos \iota$ and h_0 on the left and ψ and ϕ_0 on the right for Signal B.

by the contour plots of equal probability density in Figure 2.13.

2.4.6 Sensitivity to pulsar orientation

The correlation between h_0 and $\cos \iota$ can provide some information about the inclination of the pulsar when h_0 is small (or zero). This is a consequence of our uniform prior probabilities for h_0 . If we used a more conservative prior, such as only allowing values of h_0 less than the detector sensitivity, then obviously we do not get any information on pulsar orientation. The peculiarity presented here is only relevant to the extent that we believe that a flat prior for h_0 is appropriate.

If the data are consistent with a small h_0 this is also indicating that $\cos \iota$ is likely to be small. The reason for this is that the orientation of the pulsar is not as likely to be optimal if h_0 is small. The strain produced by a signal is largest when the pulsar rotation axis is aligned to our line of sight. When $\cos \iota = 0$ the pulsar is oriented with the poles perpendicular to our line of sight. In this case we are only irradiated with the plus polarisation of the gravitational wave signal. On the other hand, when the pulsar is oriented with the angular




Figure 2.14: Contours of constant probability, for h_0 and $\cos \iota$, for Gaussian noise with no signal using a flat prior for h_0 .

momentum vector pointing toward us ($\cos \iota = 1$) we are irradiated by both polarisations of the signal at maximum amplitude.

Figure 2.14 is a contour plot of equal probability density for h_0 and $\cos \iota$ when no signal is present (just Gaussian noise) using a flat prior for h_0 .

2.4.7 Repeated experiments with no signal

It is instructive to repeat the same analysis on different noise of the same amplitude to get a better understanding of the expected spread of pdfs in repeated experiments. Here we use 10 000 data sets of 10 days long with noise from a Gaussian distribution with $\mu = 0$ and $\sigma = 1$. For each of the marginalised posterior pdfs for h_0 the location of the peak, the most probable value, was recorded as well as the 95% upper limit.

Figure 2.15 shows the distribution of the most probable values of h_0 from each of the experiments. In 62% of the experiments the best estimate for h_0 was less than 0.005 (the width of the bins in the histogram). So even when no signal is present we would expect the peak in the pdf for h_0 to be off zero in 38% of the experiments.

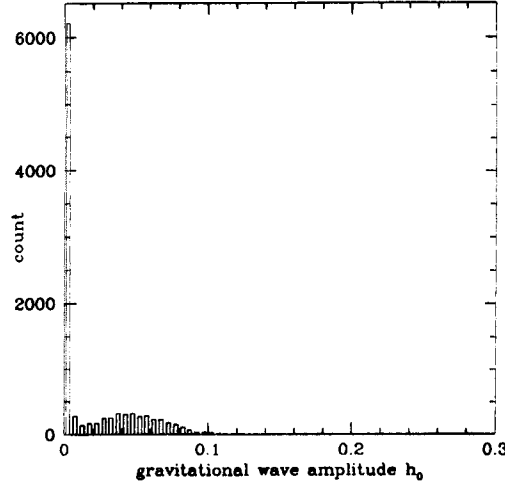


Figure 2.15: Distribution of the most probable value of h_0 when there is no signal present (10,000 data sets).

Figure 2.16 shows the distribution of 95% upper limits obtained using the same 10,000 data sets. Note that the width of this distribution is quite large; about 50% of the size of the upper limits.

2.5 Discussion

In this chapter we have presented an end-to-end Bayesian method of searching for gravitational waves from known pulsars. The method involves processing the data to reduce the number of samples required in the analysis. We calculate the likelihood function for given model parameters from the decimated data thus reducing computational requirements. In order to validate the algorithm and test how well it can retrieve signal parameters, we have analysed simulated data with signal injections. We have also shown that it is easily adapted to deal with a network of detectors.

The methodology presented in this chapter was initially developed for targeted searches with known locations and spin evolutions of the sources. Some work has been done studying the feasibility of expanding the numbers of parameters by taking advantage

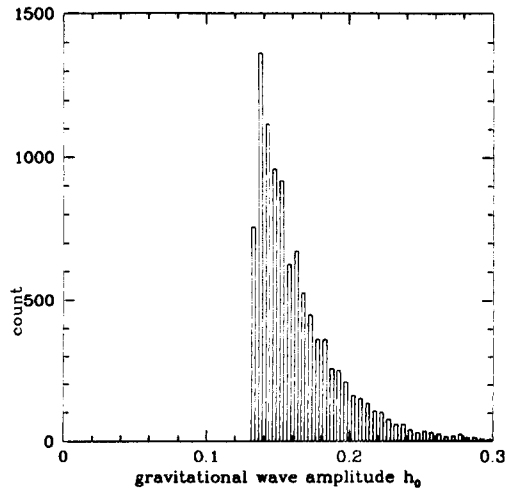


Figure 2.16: Distribution of the 95% upper limit on h_0 when there is no signal present (10,000 data sets).

of Monte Carlo Markov Chain (MCMC) techniques. Such a technique is required for increasing the number of unknown parameters since the method presented here would be too computationally intensive for a large number of parameters. More information on these studies of MCMC techniques can be found in the Appendix.

Chapter 3

Analysis of the first data from GEO 600 and LIGO

3.1 Introduction

The GEO 600 and LIGO gravitational wave detectors conducted their first coincident science run (S1) for 17 days between 23 August and 9 September 2002. The main objective of this short observation run was to collect scientifically interesting data to be analysed for gravitational waves of astrophysical origin. The detectors, however, were still in a commissioning phase and so their sensitivities were not yet optimal. In addition to providing the first results from GEO 600 and LIGO, the data collected in S1 provided an opportunity to improve our understanding of the detectors and the analysis methods. The analysis of the S1 data was divided into four groups within the LIGO Scientific Collaboration:

- a search for binary neutron star inspirals [57]
- a search for unmodelled bursts [58]

- a search for a stochastic background [42]
- a search for periodic signals from pulsar B1937+21 [48]

In all cases, the upper limits derived for each gravitational waves source were lower than from previous broadband experiments. For the binary neutron star inspiral search, an upper limit on their coalescence rate was set to less than 170 per Milky Way Equivalent Galaxy per year with a 90% confidence. This analysis concentrated on binary neutron star systems in which each component had a mass in the range of $1 - 3M_{\odot}$. The search for gravitational wave bursts was able to constrain the event rate to less than 1.6 events per day, at 90% confidence level. This result focused on short bursts where an event had a duration from 4 ms to 100 ms with significant power in the sensitivity band of 150 to 3000 Hz. The stochastic background analysis established a 90% confidence limit of $\Omega_{GW}h_{100}^2 \leq 23$ in the frequency band 40 to 314 Hz, where h_{100} is the Hubble constant in units of 100 km/sec/Mpc and Ω_{GW} is the gravitational wave energy density per logarithmic frequency interval in units off the closure density (as defined in Chapter 1). In this chapter, we present an analysis that was conducted to constrain the level of gravitational waves from pulsar B1937+21 at twice its rotation frequency.

The main result from this chapter, the upper limit on gravitational waves from pulsar B1937+21, was published in [48] (see Appendix). In that paper, two independent analysis methods were used: a form of the time-domain Bayesian methods presented in Chapter 2, and a frequency domain technique using a statistic from the method of maximum likelihood [49]. In this chapter a detailed description of the application of the time-domain method to the S1 data is presented. In Section 3.2 we describe the configuration and the performance of the GEO 600 and LIGO interferometers during the S1 run. We describe the expected sensitivity of the data to periodic signals and astrophysically motivated existing limits on gravitational wave emission from these systems in Sections 3.3. We then proceed to introduce the some properties of the targeted pulsar, B1937+21, in Section 3.4. In Section 3.5 we address the issue of pulsar timing and validation of our timing software.

Detector	Total science data (hours)	Duty cycle
GEO 600	401	98%
LIGO H1	235	57.6%
LIGO H2	298	73.1%
LIGO L1	170	41.7%

Table 3.1: Total periods where the interferometers were in-lock and the duty cycles for each detector during the S1 run.

The calibration procedure for the LIGO detectors is explained and relevant measurements are given in Section 3.6. The data analysis method is reviewed in Section 3.7. In Section 3.8 we characterise the data in the frequency band of interest for this analysis, near 1284 Hz. The results are presented in Section 3.9 and we conclude in Section 3.10.

3.2 The detectors

The GEO 600 and LIGO interferometers had reached a sufficient level of sensitivity by the summer of 2002 to warrant a short coincident science run (S1). A previous coincident run, E7, had taken place in 2001 but no astrophysical results were published. The sensitivities reached during the S1 run were comparable or better than previous broadband gravitational wave detectors. The locked times and duty cycles of the four interferometers (IFOs) for the S1 run are given in Table 3.1. Note that these numbers are slightly higher than the observation periods that were analysed for periodic signals later in this chapter. This reflects the fact some data had to be dropped from the analysis because the detectors were not in ‘science-mode’. For example, this could be due to periods when the calibration lines (required for calibration) were not being measured properly but the detector was in-lock. The high duty cycle (98%) of GEO 600 is partly due to the control systems that were specifically designed for automatic operation. The automatic alignment system in GEO 600 was so stable and reliable during S1 that no operator presence was required overnight at the detector site [59]. For LIGO, the Livingston site had the lowest duty cycle at 41.7%. Most of the time lost at L1 was due to the interferometer falling out-of-lock during the daytime. This was mostly caused by environmental noise sources




Figure 3.1: Typical sensitivity of the three LIGO interferometers and GEO 600 during the S1 data run.

due to humans (e.g. logging).

Figure 3.1 shows typical amplitude spectra of equivalent strain noise for GEO 600 and each of the LIGO IFOs during the S1 run. The broadband sensitivity of the interferometers did not change significantly over the course of the run. Note that at frequencies above 1 kilohertz (where we searched for periodic signals) the sensitivities of all four detectors were comparable. The Livingston IFO was the most sensitive, about a factor of two better than the other LIGO IFOs at 1284 Hz.

3.2.1 Instrument configurations

Both GEO 600 and LIGO were still in commissioning mode during the S1 run and neither were running at design sensitivity. GEO 600 ran as a power-recycled Michelson interferometer during S1 without a signal recycling mirror. The signal recycling mirror was installed in the autumn of 2003 and improved the sensitivity of GEO 600 at high frequencies for the S3 run [60]. The laser power build up in the power recycling cavity was limited to 300 during S1 (with a target of 2000) due to the poor quality of the power

recycling mirror installed at that time. This contributed to poorer GEO sensitivity at high frequencies where it was limited by shot noise. For the LIGO IFOs, the noise at high frequencies can also be attributed to using lower laser power than the design configuration. For more details and a complete description of the detector configurations and performances in S1, see [59].

The IFOs are calibrated by injecting a series of sinusoidal forces of known amplitude to the test masses. These signals are then recovered and used to construct a frequency dependent transfer function to calibrate the raw output. For GEO 600, the data are calibrated in a similar fashion in the time-domain in advance and then stored in the Frame¹ data format. Details of the GEO 600 calibration procedure are available in [62]. For LIGO, the data are stored in their raw form (error-signal) and must be calibrated for interpretation. The procedure to follow to calibrate the LIGO data is described in Section 3.6. The main uncertainty in the calibration process comes from the imprecision in the electromechanical coupling coefficients of the force actuators [48]. For the analysis presented in this chapter we only looked at data near 1284 Hz for signals from pulsar B1937+21. For the S1 run, the total amplitude calibration uncertainty at this frequency was estimated at $\pm 4\%$ for GEO, and $\pm 10\%$ for each of the LIGO interferometers.

3.3 Expected sensitivity to periodic signals

From typical detector sensitivities we can determine the average 95% upper limit $\langle h_{95} \rangle$ expected as a function of observation time and noise power spectral density. We characterise the noise level by its single-sided power spectral density $S_n(f)$ where f is the frequency of the signal. To calculate $\langle h_{95} \rangle$ we calculated the upper limits in 4000 simulations varying the location in the sky and detector site² for each set. The locations of the sources were picked so that they would be uniform per unit solid angle in the sky. For

¹A common data format for interferometric gravitational wave detectors, see [61]

²Detector sites were between GEO 600, H1, H2, and L1




Figure 3.2: Distribution of 95% upper limits on h_0 for 4000 sources covering the whole sky with $T = 1$ and $S_n(f) = 1$.

each of the hypothetical sources, however, we do not inject any signal into the Gaussian noise. By averaging the upper limits obtained we found that

$$\langle h_{95} \rangle = 15.3 \sqrt{S_n(f)/T}, \quad (3.1)$$

where T is the observation time in seconds. Figure 3.2 shows the distribution of h_{95} when $T = 1$ and $S_n(f) = 1$. Note that the width of the distribution is relatively large so the actual upper limits could easily be a factor of two larger than $\langle h_{95} \rangle$.

The coloured upper curves on Figure 3.3 show $\langle h_{95} \rangle$ for the LIGO and GEO 600 interferometers during S1. The observation times used correspond to the actual up-time of each interferometer. The black lower curves represent $\langle h_{95} \rangle$ corresponding to the design sensitivity of GEO 600 and LIGO. For the target sensitivities of the detectors, an observation period of one year was assumed.

The amount of gravitational waves potentially emitted by pulsars can be limited by letting the total loss in rotational energy be due to gravitational wave dissipation. In such a model, we ignore the spin-down contribution from electromagnetic and particle emission.

Equating the rotational energy loss to the gravitational luminosity of a rotating triaxial star, we get [63]

$$\frac{d}{dt} \left(\frac{1}{2} I \omega^2 \right) = -\frac{32G}{5c^5} I^2 \epsilon^2 \omega^6, \quad (3.2)$$

where $\omega = 2\pi f_0$ is the pulsar angular frequency and ϵ is its equatorial ellipticity. Equation 2.6 in Chapter 2 relates the ellipticity ϵ to the gravitational wave amplitude h_0 , so that

$$\epsilon = \frac{h_0 c^4 r}{16\pi^2 G I f_0^2}. \quad (3.3)$$

From Equation 3.2 and Equation 3.3, it is easy to show that a spindown based upper limit on the gravitational wave amplitude h_0 can be expressed as

$$h_0 = \sqrt{\frac{5G}{2c^3} \frac{I}{r^2} \frac{\dot{P}}{P}}, \quad (3.4)$$

where r is the distance to the source, P is the rotational period of the pulsar, and \dot{P} is the first time derivative of the period.

The solid circles on Figure 3.3 show the constraints on the gravitational wave amplitude from the measured spindown rates for a selection of radio pulsars. For each pulsar represented on the plot, a standard value of 10^{38} kg m^2 was used for the moment of inertia. Many of the known pulsars are clearly rotating too slowly to be detected by ground-based interferometers. However, the number of known millisecond pulsars should continue to increase with new radio surveys. A number of globular cluster pulsars are not included in the figure since their measured spindown rate is positive! The observed spin-up of these sources is attributed to proper motions of the pulsars within the clusters.

It is clear from the limits in Figure 3.3 that we did not expect to detect gravitational waves from the known radio pulsars with the S1 data. However, this does not exclude a potential class of neutron star undiscovered that may be seen in all-sky gravitational wave surveys. For known sources, the Crab pulsar (B0531+21) is the only pulsar where we will be able to set limits tighter than those set by the spindown argument even with one year of data




Figure 3.3: The coloured upper curves represent average 95% upper limits $\langle h_{95} \rangle$ expected using typical GEO 600 and LIGO sensitivities during S1. The black lower curves represent the expected $\langle h_{95} \rangle$ for the design sensitivities of GEO 600 and LIGO using an observation period of one year. The arrow points to the solid circle representing pulsar B1937+21.

at design sensitivities.

3.4 Gravitational waves from pulsar B1937+21

For illustration we chose to set upper limits on pulsar B1937+21 during S1. This is the fastest rotating millisecond pulsar rotation frequency of 642 Hz and so we searched for gravitational waves near 1284 Hz.

There are a few reasons why this specific pulsar was chosen as the first test case. In previous commissioning runs of GEO 600 and LIGO, we had done preliminary testing of our analysis methods on this pulsar. We had chosen B1937+21 because the detectors had their best sensitivity above 1 kilohertz. This was argument was still true for GEO 600 during S1 but not for the LIGO IFOs. Another reason for choosing this pulsar is that the behaviour of the noise at such high frequency was relatively stationary and Gaussian compared to lower frequencies. Since B1937+21 is the fastest rotating known pulsar we

right ascension (J2000)	19 ^h 39 ^m 38 ^s .560 210(2)
declination (J2000)	+21° 34' 59".141 66(6)
RA proper motion	−0.130(8) mas yr ^{−1}
dec proper motion	−0.464(9) mas yr ^{−1}
period	0.001 557 806 468 819 794(2) s
period derivative	1.051 193(2) × 10 ^{−19} s s ^{−1}
epoch of period and position	MJDN 47 500

Table 3.2: Parameters for pulsar B1937+21 also known as pulsar J1939+2134. Numbers in parentheses indicate uncertainty in the last digit.

thought it would also provide one of the tightest limits on ellipticity. In retrospect, there are a few nearby isolated pulsars with signal frequencies of a few hundred Hertz that would have given us tighter limits on ellipticity due to their proximity and the sensitivity of the LIGO IFOs in that band.

It was verified through radio observations that B1937+21 did not glitch during S1 and that the source parameters we used were still valid [64]. Table 3.2 gives the parameters of pulsar B1937+21 taken from the literature [65]. Figure 3.4 shows the timing residuals from pulsar B1937+21 over several years including the duration of the S1 run. The residuals are between the actual time-of-arrival of the radio signals and a standard timing model as we use. These are small enough not to affect our analysis. However, timing delays as small as a hundred microseconds could affect a multidetector analysis.

The declination of pulsar B1937+21 is relatively low at +21° 34' 59" (as indicated by its name). Given the location of the GEO 600 and LIGO interferometers on the Earth, one might expect the Louisiana site to have the best view of the pulsar. Figure 3.5 shows the beam profile for a source in the direction of pulsar B1937+21 for GEO 600, LLO (L1), and LHO (H1 and H2) during S1. The quantity plotted is $(F_x^2 + F_y^2)^{1/2}$ which is the overall sensitivity of the interferometers to both polarisations. In Figure 3.5 the average sensitivities of each IFO to signals from the position of B1937+21 during S1 are shown (dashed lines) and they are: 0.54 for GEO 600, 0.60 for LLO, and 0.56 for LHO. As expected the Livingston IFO is slightly more sensitive to signals from B1937+21. However, one must keep in mind that the LLO IFO had the worst duty cycle among the




Figure 3.4: Radio timing residuals for pulsar B1937+21 over several years including the duration of the S1 run, from [64].




Figure 3.5: Sensitivity of GEO 600, LLO (L1), and LHO (H1 and H2), to gravitational waves from the direction of pulsar B1937+21 during S1. The dotted line shows the average sensitivity for each IFO and has a value of 0.54 for GEO 600, 0.60 for LLO, and 0.56 for LHO. The better sensitivity of LLO is due to its low latitude compared to LHO and GEO 600.




Figure 3.6: Diagram of the Earth, Sun, Solar System Barycenter, and an incoming gravitational wave from direction \mathbf{k} . The vector \mathbf{r}_d indicates the position of the detector with respect to the SSB. The goal of the timing software is to provide the transformation from topocentric time t at the detector to time T at the SSB.

four IFOs during S1. Nonetheless since the Livingston IFO had the best sensitivity near 1284 Hz it was that detector that provided the best upper limit.

3.5 Accuracy of timing software

A crucial part of the time domain analysis method is the demodulation of the phase of the gravitational wave signal. In order to heterodyne the data with the appropriate phase we need to use a reference frame which is inertial with the pulsar. For this purpose we use time as measured by a clock at the Solar System Barycentre (SSB) but with the absence of a massive object. Figure 3.6 is a schematic of the objects of interests including the SSB which is situated just outside the surface of the Sun towards Jupiter. The vector \mathbf{r}_d is the position of the detector with regard to the SSB. The unit vector \mathbf{k} is pointing toward the source of gravitational waves. The timing software used must transform the time-of-arrival of a signal at the detector to the time-of-arrival of that signal at the SSB. We want to use the time at the SSB so we must correct for the different gravitational

potential between the detector and the SSB.

There are three main terms that must be taken into account for the Barycentric corrections. For an isolated pulsar such as B1937+21, the time difference δt between the topocentric time at the detector, t , and the time as measured by a clock at the SSB, T , can be expressed as

$$\delta t = \Delta_{\text{Roemer}} + \Delta_{\text{Shapiro}} + \Delta_{\text{Einstein}}. \quad (3.5)$$

The largest of these corrections is the Roemer delay Δ_{Roemer} which can be as large as ~ 8.5 minutes. The Roemer delay can be expressed as

$$\Delta_{\text{Roemer}} = \frac{\vec{r}_d \cdot \vec{k}}{c} + \frac{(\vec{r}_d \cdot \vec{k})^2 - |\vec{r}_d|^2}{2cd}, \quad (3.6)$$

where \vec{r}_d is the position of the detector with regard to the solar system barycentre, \vec{k} is a unit vector in the direction of the pulsar, and d is the distance from the detector to the pulsar. In order to calculate the Roemer delay we obviously need accurate knowledge of the position of the Earth with regard to the SSB. For our Barycentric software we use the Solar System ephemerides published by the Jet Propulsion Laboratory [66]. The second term in Equation 3.6 is known as the timing parallax. It is a measure of the curvature of the wavefronts emitted from the source and is only important for nearby pulsars. Figure 3.7 shows the magnitude of the Roemer delay during and after S1 for a source in the direction of B1937+21.

The second term in Equation 3.5 is the Shapiro delay Δ_{Shapiro} and is a relativistic correction due to the curvature of space-time near the SSB. Since this curvature is not negligible there will be an extra time delay in the arrival of a signal. In principle this delay can be as large as $120 \mu\text{s}$ for signals passing near the edge of the Sun. The maximum contribution from Jupiter is 200 ns and so not important for us. Figure 3.7 shows the Shapiro delay in the direction of B1937+21 during and after S1.

The last term in Equation 3.5 is the Einstein delay Δ_{Einstein} . This term describes the




Figure 3.7: Roemer delay (right) and Shapiro delay (left) for a source in the direction of pulsar B1937+21 for a period of 1 year beginning at the start of S1.

combined effect of gravitational redshift and time dilation due to the motion of the Earth. The Einstein delay correction takes into account the varying gravitational potential experienced by a clock on the Earth as it follows its elliptical orbit around the Sun. Figure 3.8 shows the effect of the Einstein delay for the S1 analysis (aside from a constant offset).

The timing routines that we used are available in the LIGO Analysis Library and were written by Curt Cutler [67]. We independently checked this software against a widely used radio astronomy package called TEMPO for accuracy [68]. This comparison, see Figure 3.8, confirmed a timing accuracy better than $\pm 4 \mu\text{s}$. This is sufficiently accurate since this ensures no more than 0.01 rad phase mismatch between a putative signal and its template.

3.6 Calibrating the LIGO data

The gravitational wave strain that we want to measure with the interferometers can be described as

$$s(t) = [L_x(t) - L_y(t)] / L_0 \quad (3.7)$$




Figure 3.8: Left: Einstein delay within a constant offset for a source in direction of B1937+21 for 1 year beginning at the start of S1. Right: Histogram of timing residuals between our barycentring routines and TEMPO, derived by comparing the phase evolution of test signals produced by the two software packages. 156 locations in the sky were chosen at random and the residuals calculated once an hour for the entire year 2002. The maximum timing error is $< 4 \mu\text{s}$.

where $L_{x,y}$ are the lengths of the orthogonal arms of the interferometer and L_0 is the effective length of the cavities in the absence of an external strain. The raw output of the GEO 600 and LIGO interferometers, the so-called error-signal $q(t)$, is not calibrated. In this section, the procedure for calibrating the LIGO data, going from $q(t)$ to $s(t)$, will be describe. The data stored in the Frames files for the GEO 600 interferometer are pre-calibrated by the method in [62].

The raw output, $q(t)$, from the interferometers is derived from light from the anti-symmetric port which is held close to a dark fringe. After going through the servo loop (feedback mechanisms) it is stored in arbitrary units of counts. The response function, $R(f)$, is defined as the function that will reconstruct the gravitational wave strain, $s(t)$, from the raw output, $q(t)$. In the frequency domain, we have

$$\tilde{s}(f) = R(f)\tilde{q}(f), \quad (3.8)$$

where $\tilde{s}(f)$ and $\tilde{q}(f)$ represent the Fourier transforms of $s(t)$ and $q(t)$, respectively.

The so-called error signal $\tilde{q}(f)$ differs from $\tilde{s}(f)$ because of the control forces that keep

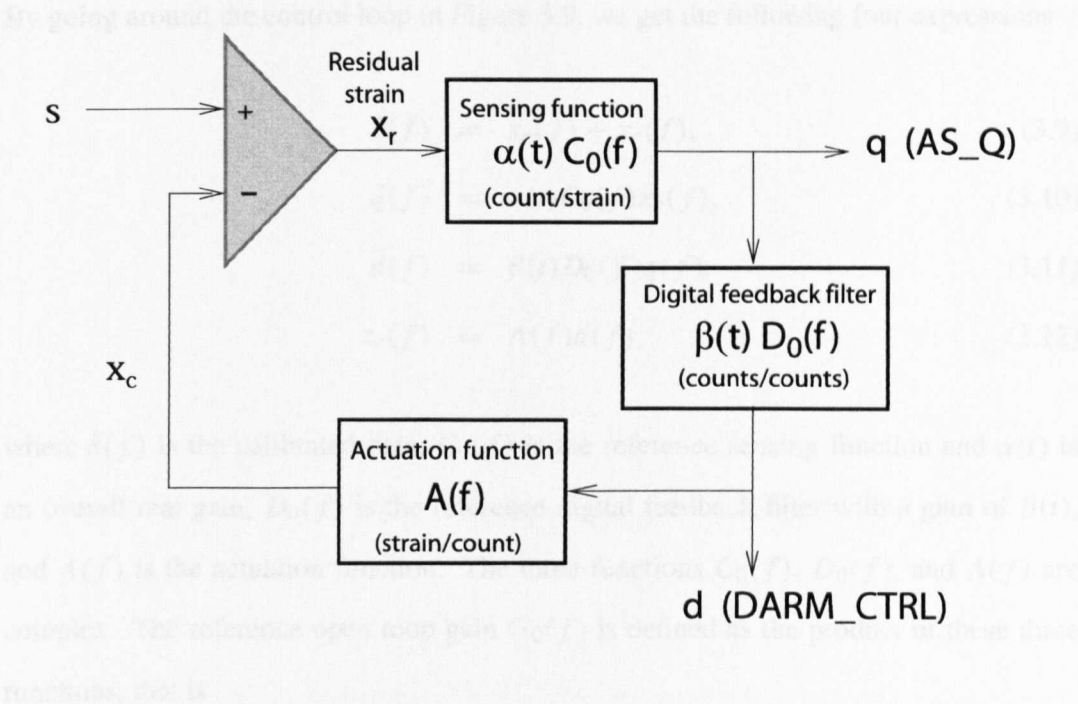


Figure 3.9: Simple model of the LIGO length sensing and control feedback system. Two digital signals, $q(AS_Q)$ and $d(DARM_CTRL)$, are recorded by the data acquisition system with the names in parentheses representing their names as stored in the Frames files.

the motion of the mirrors small and the optical cavity in resonance. In order to reconstruct the strain signal $\tilde{s}(f)$ we need to undo these effects introduced by the feedback loop into $\tilde{q}(f)$. This is done via the response function $R(f)$ which can be built from the feedback loop equations.

The LIGO length sensing and control system subtracts a control signal x_c from the external strain s to minimize the residual strain x_r that enters the optical cavity. In the simple model shown in Figure 3.9 two digital signals are recorded by the data acquisition system: $q(t)$ which is the digitized error signal, and $d(t)$ which is a control signal that is fed back to an actuator to adjust the length of the cavity.

By going around the control loop in Figure 3.9, we get the following four expressions

$$\tilde{s}(f) = x_r(f) + x_c(f), \quad (3.9)$$

$$\tilde{q}(f) = \alpha(t)C_0(f)x_r(f), \quad (3.10)$$

$$\tilde{d}(f) = \beta(t)D_0(f)q(f), \quad (3.11)$$

$$x_c(f) = A(f)d(f), \quad (3.12)$$

where $\tilde{s}(f)$ is the calibrated data, $C_0(f)$ is the reference sensing function and $\alpha(t)$ is an overall real gain, $D_0(f)$ is the reference digital feedback filter with a gain of $\beta(t)$, and $A(f)$ is the actuation function. The three functions $C_0(f)$, $D_0(f)$, and $A(f)$ are complex. The reference open loop gain $G_0(f)$ is defined as the product of these three functions, that is

$$G_0(f) = A(f)C_0(f)D_0(f). \quad (3.13)$$

Using the equations above, along with Equation 3.8, it is easy to show that

$$R(f) = \frac{1 + \alpha(t)\beta(t)G_0(f)}{\alpha(t)C_0(f)}. \quad (3.14)$$

Figure 3.10 shows the variation of the overall open loop gain ratio, $\alpha(t)\beta(t)$, during S1 for all three LIGO IFOs. The gain on the digital feedback filter, $\beta(t)$, was set to one for H1 and H2 for the duration S1. $\beta(t)$ only deviated from one for a short period of time in L1 during S1. The gaps in Figure 3.10 pertain to periods when no data were available. There were no data available on the first day of S1 as that day was used for calibration studies. The values of $\alpha(t)$ were calculated every 60 seconds and derived from measurements of the amplitude of the calibration lines. A key assumption here is that the gain varies on a time scale longer than 60 seconds. Figure 3.11 shows the reference open loop gain functions used in for the calibration of the S1 data. Figure 3.12 shows the sensing functions used for the S1 calibration. The measurements for H1 and H2 were done at GPS time 715 156 759. For L1, however, a model was used made up of known




Figure 3.10: Overall open loop gain factor $\alpha(t)\beta(t)$ over the course of S1 for the three LIGO IFOs. The gap of data on the first day of S1 is due to calibration studies that interfered with taking measurements of the gain that day. The other gaps in the data correspond to periods where the detector was out-of-lock or the calibration lines were not being measured.

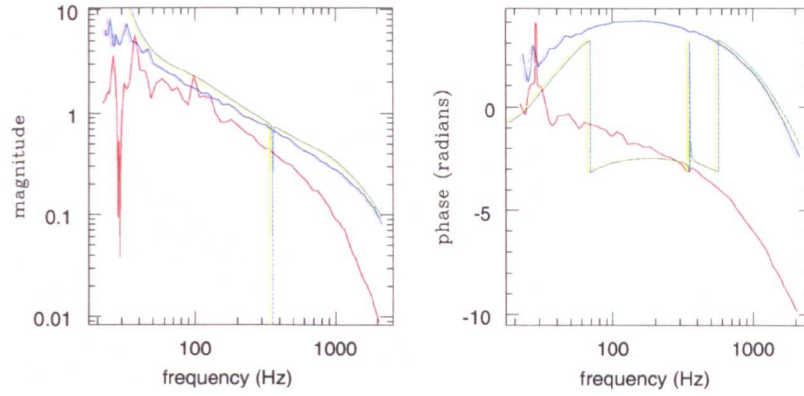


Figure 3.11: Reference open loop gain $G_0(f)$ taken at GPS time 715 156 759 for the three LIGO IFOs: H1 in red, H2 in blue, and L1 in green.

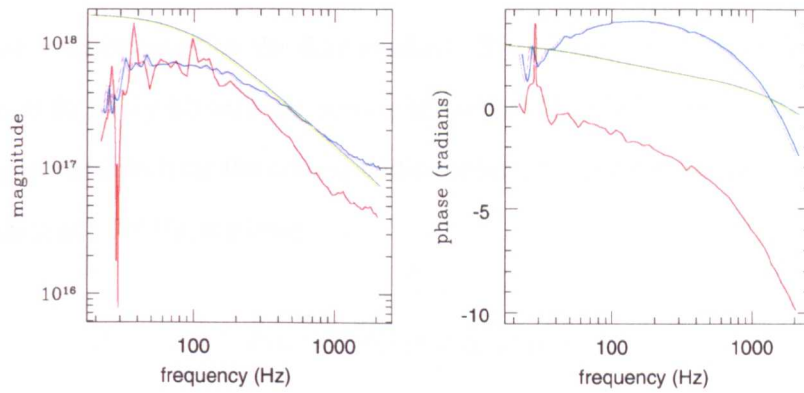


Figure 3.12: Reference sensing function $C_0(f)$ taken at GPS time 715 156 759 for the three LIGO IFOs: H1 in red, H2 in blue, and L1 in green.

components in the loop.

The S1 LIGO data were calibrated keeping in mind that we were just interested in a narrow bandwidth of data around 1284 Hz. Instead of calibrating the full 16 384 Hz band of data we only calibrated the data of interest. Since the frequency of the signal we are seeking is 1284 Hz, and the maximum Doppler shift is 10^{-4} , the band we need to calibrate is less than 1 Hz wide. By examining the example response functions in Figure 3.13 it is clear that it is reasonable to assume that a 1 Hz wide band near 1284 Hz is quite flat. In fact, the change in the magnitude and phase of the response function at 1284 Hz is less than 1% over a range of 1 Hz. So we can assume that we only need one frequency bin of the response function to calibrate the data. We simply used a point estimate of the response function, in the center of our band, to calibrate the data. The width of the bandwidth in




Figure 3.13: Reference response function at GPS time 715 156 759 for the three LIGO IFOs: H1 in red, H2 in blue, and L1 in green.

this analysis is determined by the data products B_k described in Chapter 2. For S1 we produced one B_k every 60 seconds corresponding to a bandwidth of $1/60$ Hz. We used the gain measurements from the corresponding minute to calibrate the data. Since we are analysing data at 1284 Hz, we have

$$B(t_k) = R(1284; t_k)B'(t_k), \quad (3.15)$$

where $B(t_k)$ are the calibrated data, $R(1284; t_k)$ is the response function at 1284 Hz at time t_k , and $B'(t_k)$ is the uncalibrated data.

A more detailed description of the calibration of the LIGO IFOs during S1 is available in [69].

3.7 Data analysis method

The results presented here make use of the data analysis methods presented in Chapter 2. The data were heterodyned and Doppler demodulated, low-passed filtered, and resampled in order to generate one data product, B_k , per minute. For efficiency the complex heterodyning was done in two steps for S1. The data set was first heterodyned at a fixed heterodyne frequency $2f_0$, near the signal's true frequency, and then down-sampled to

a reduced data rate of 4 Hz. We will denote as $V(t_i)$ this intermediate 4 Hz time series centered on the signal frequency $2f_0$. This permits us to do the Doppler shift corrections at a much smaller computational cost; a useful strategy in case the analysis needs to be repeated.

The time series $V(t_i)$ was generated by initially heterodyning the data at a frequency of $2f_0$ (multiplying the data by $e^{-4\pi f_0 t}$). Then we applied a low-pass third order IIR Butterworth filter to the real and imaginary components of the heterodyned data (at 16384 Hz) to remove the rapidly oscillating term and prevent aliasing from outside the band of interest. We then averaged each consecutive 4096 samples to reduce the sampling frequency from 16384 Hz to 4 Hz. This was done in two steps, with two IIR filters to improve the performance of the filters. The data were first low-pass filtered with a cutoff frequency of 128 Hz and averaged. This 128 Hz data stream was then low-pass filtered by a second IIR filter with a cutoff frequency of 4 Hz and averaged once more to give 4 samples per second.

We used essentially the same Bayesian approach as described in Chapter 2 for parameter estimation. For the likelihood we used a bivariate Gaussian distribution so that

$$p(\mathbf{a}|\{B_k\}) \propto p(\mathbf{a}) \exp \left[- \sum_k \frac{\Re\{B_k - y(t_k; \mathbf{a})\}^2}{2\sigma_{\Re\{B_k\}}^2} \right] \times \exp \left[- \sum_k \frac{\Im\{B_k - y(t_k; \mathbf{a})\}^2}{2\sigma_{\Im\{B_k\}}^2} \right] \quad (3.16)$$

where $\sigma_{\Re\{B_k\}}^2$ is the variance of the real parts of B_k and $\sigma_{\Im\{B_k\}}^2$ is the variance of the imaginary parts of B_k .

Since the frequency bins near the signal were relatively flat (except for GEO 600), we used data from our existing 4 Hz band ($V(t_i)$) (centered on the signal frequency $2f_0$) to estimate the noise in the B_k 's. The B_k 's were constructed from the 4 Hz time series, $V(t_i)$, as follows

$$B_k = \frac{1}{240} \sum_{j=1}^{240} V(t_j) e^{-i\zeta(t_j)} \quad (3.17)$$

where

$$\zeta(t) = 4\pi \left[f_0(\delta t - T_0) + \frac{1}{2} \dot{f}_0(t + \delta t - T_0)^2 + \frac{1}{6} \ddot{f}_0(t + \delta t - T_0)^3 \right] \quad (3.18)$$

where δt is the time delay from the detector to the SSB, f_0 is frequency of the pulsar, \dot{f}_0 is the first time derivative of the frequency, \ddot{f}_0 is the second time derivative of the frequency, T_0 is our reference epoch, and t is the topocentric time measured at the detector.

We estimated the noise variance in the bin values, B_k , from the sample variance of $V(t_i)$ for the corresponding time interval. We estimated the variance of the real and imaginary components of B_k separately as follows

$$\sigma_{\Re\{B_k\}}^2 \simeq \frac{\sum_{j=1}^{240} [\Re\{V(t_j)\}_k - \langle \Re\{V(t_j)\}_k \rangle]^2}{240(240 - 1)}, \quad (3.19)$$

where

$$\langle \Re\{V(t_j)\}_k \rangle = \frac{1}{240} \sum_{j=1}^{240} \Re\{V(t_j)\}_k, \quad (3.20)$$

and

$$\sigma_{\Im\{B_k\}}^2 \simeq \frac{\sum_{j=1}^{240} [\Im\{V(t_j)\}_k - \langle \Im\{V(t_j)\}_k \rangle]^2}{240(240 - 1)}, \quad (3.21)$$

where

$$\langle \Im\{V(t_j)\}_k \rangle = \frac{1}{240} \sum_{j=1}^{240} \Im\{V(t_j)\}_k. \quad (3.22)$$

The numerical factor in the above equation is 240 since there were 4 samples per second for 60 seconds giving 240 samples per minute. For S1 we did not marginalise over the variance to compensate for the uncertainty in our estimate of the noise.

3.8 Characterisation of the data

For this analysis we only considered segments of at least 300 seconds of data during which the detectors were in lock and science-mode. For the GEO 600, half a day of data

(642 minutes) from before the official start of S1, and one day and a half (2330 minutes) after the official end were also included. The total number of minutes analysed were 26 981 for GEO 600, 12 290 for H1, 11 282 for H2, and 8275 for L1.

Figure 3.14 shows a single-sided power spectral density of a band of 4 Hz using S1 data for all four interferometers. The expected signal frequency during S1, taking into account only the orbital Doppler shift, was 1283.8 Hz. The dashed lines in Figure 3.14 indicate that frequency bin. There is a clear broad spectral feature in the GEO 600 data near the frequency where we expect a signal. The line, however, is not likely to be an astrophysical signal but rather some instrumental artifact. The line has a width of about 0.5 Hz which is broader than would be expected from Doppler modulations due to the motion of the Earth. Perhaps the strongest case that this is not a gravitational wave signal is that it is not visible in the other interferometers. Of course, we would expect a real gravitational wave signal to appear in all detectors. Figure 3.15 shows a time-frequency spectrogram showing the evolution of the line in the GEO 600 data over the course of S1. The time resolution of the spectrogram is one day. The spectrogram covers a band of 4 Hz around the frequency of the signal and the colour is proportional to the signal to noise.

Figure 3.16 shows in the top plot the modulus of B_k (in black) as a function of time as well as the calculated estimate of the noise σ_k (in red) for each sample where we have

$$\sigma_k = \sqrt{\left(\sigma_{\Re\{B_k\}}^2 + \sigma_{\Im\{B_k\}}^2\right)}. \quad (3.23)$$

In the bottom plot the ratio between $|B_k|$ and $|\sigma_k|$ is shown. It is evident that there are variations in the level of the noise over the course of a day. It is comforting to take note, however, that we seem to be correctly estimating these changes in the noise floor since the bottom plots appear much more Gaussian. Figure 3.17, Figure 3.18, and Figure 3.19 show similar plots during S1 for the LIGO IFOs. The gaps in the data are for periods when the detectors were out of lock and/or not in science mode. These plots are useful monitors of the stationarity of the data in the band of interest. These plots are essentially




Figure 3.14: Single-sided power spectral density ($\sqrt{S_n}$) of a band of 4 Hz (starting at 1281 Hz) using the entire S1 data set analyzed from the four interferometers. The noise $\sqrt{S_n}$ is shown in units of $10^{-20} \text{ Hz}^{-1/2}$. The dashed vertical line indicates the expected frequency of the signal received from B1937+21 (this figure is taken from [48]).




Figure 3.15: This is a spectrogram of GEO 600 data during S1 for a 4 Hz band around the signal frequency (1284 Hz). The color of the pixels is proportional to the signal-to-noise of the pixel compared to data of the same period.




Figure 3.16: The time-series in the top plot shows $|B_k|$ in black (one sample per minute) and the estimate of the noise $|\sigma_k|$ in red for the GEO data for all of S1. In the bottom plot the time series shows $|B_k|/|\sigma_k|$ for the duration of S1 for GEO.




Figure 3.17: The time-series in the top plot shows $|B_k|$ in black (one sample per minute) and the estimate of the noise $|\sigma_k|$ in red for the H1 data for all of S1. In the bottom plot the time series shows $|B_k|/|\sigma_k|$ for the duration of S1 for H1.

showing the level of the noise in a band $1/60$ Hz around the signal frequency (in this case ~ 1284 Hz).

It is important to check that the distribution of B_k and σ_k is consistent with a Gaussian distribution since this is an assumption for the likelihood function we are using. The histograms in Figure 3.20 plot $B/\sigma = \Re(B_k)/\sigma_{\Re\{B_k\}} + \Re(B_k)/\sigma_{\Re\{B_k\}}$ which should be Gaussian with a variance of one if we have correctly modeled the noise. We quantitatively assessed the Gaussianity of the data by calculating the χ^2 statistic, assuming a mean of zero, for both the real and imaginary data. That is,

$$\chi^2 = \sum_k \Re(B_k)^2/\sigma_{\Re\{B_k\}}^2 + \Im(B_k)^2/\sigma_{\Im\{B_k\}}^2. \quad (3.24)$$

If the data are consistent with a Gaussian distribution with a mean of zero then we would expect the value of χ^2 to have a mean of N and a standard deviation of $\sqrt{2N}$ where




Figure 3.18: The time-series in the top plot shows $|B_k|$ in black (one sample per minute) and the estimate of the noise $|\sigma_k|$ in red for the H2 data for all of S1. In the bottom plot the time series shows $|B_k|/|\sigma_k|$ for the duration of S1 for H2.




Figure 3.19: The time-series in the top plot shows $|B_k|$ in black (one sample per minute) and the estimate of the noise $|\sigma_k|$ in red for the L1 data for all of S1. In the bottom plot the time series shows $|B_k|/|\sigma_k|$ for the duration of S1 for L1.




Figure 3.20: Histogram of $B/\sigma = \Re(B_k)/\sigma_{\Re(B_k)} + \Im(B_k)/\sigma_{\Im(B_k)}$ for each interferometer using the S1 data sets. The dotted lines represent the expected Gaussian distribution with a mean of zero and variance of one.

N is the number of degrees of freedom. We excluded a small number of outliers with magnitudes of B_k/σ_k larger than 5 from the analysis. For the GEO600 data set we had a total of 53 955 (real + imaginary) data values (B_k) and a corresponding value of $\chi^2 = 61\,699$. This represented a larger value of χ^2 that we would have expected which could be due to a slight underestimation of the noise. This may be due to an excess of transients in the GEO600 data which made our assumption of stationary over 60 seconds not valid for some periods or the nearby spectral line. For the LIGO IFOs we had a set of 24 550 data values for H1, 22 564 for H2, and 16 544 for L1. The calculated values of χ^2 were 24 623 for H1, 22 829 for H2, and 16 684 for L1 with standard deviations $\sqrt{2N}$ of 222, 212, and 183, respectively. The χ^2 values for the LIGO IFOs were all consistent with a Gaussian data set.

3.9 Results

We set upper limits on the gravitational wave strain h_0 from pulsar B1937+21 by calculating the marginalised posterior probability distribution function of h_0 given the data. Similarly to the examples given in Chapter 2 we used flat prior probabilities for $p(\phi_0)$, $p(\cos \iota)$, and $p(\psi)$. We chose the prior for h_0 as constant for $h_0 \geq 0$ and zero for $h_0 < 0$.

The marginalised posterior pdfs for h_0 for each detector are shown in Figure 3.21. The 95% upper limits are 2.2×10^{-21} for GEO600, 1.4×10^{-22} for L1, 3.3×10^{-22} for H1, and 2.4×10^{-22} for H2. Unfortunately, due to timing uncertainties between the interferometers in S1 we were not able to perform a joint coherent analysis.

The dotted line in the GEO600 pdf in Figure 3.21 represents the result obtained if an artificial signal was added to the data with $h_0 = 2.2 \times 10^{-21}$, $\phi_0 = 0^\circ$, $\psi = 0^\circ$ and $\iota = 0^\circ$. If a signal with those parameters was present in the data we would have certainly seen it.

We can focus the upper limits on h_0 as an upper limit on the equivalent ellipticity of the pulsar using Equation 3.3. Taking the distance to pulsar B1937+21 to be 3 kpc [65], and using our knowledge that $h_0 \leq 1.4 \times 10^{-22}$ from L1, we get

$$\epsilon \leq 2.9 \times 10^{-4} \left(\frac{10^{28} \text{ Hz}^2}{f_0^2} \right) \quad (3.29)$$

with 95% confidence. The limit on the ellipticity from previous measurements for this

pulsar is $\epsilon \leq 1.1 \times 10^{-3}$ which is a factor of 4 larger than the limit from this work. If we assume that the ellipticity is due to the presence of a bar in the pulsar, the distance to the pulsar is the distance to the pulsar, for some pulsars, such as the Crab pulsar, the distance to the pulsar is a factor of 10 larger [70].

The uncertainty in the calibration with 95% confidence in the L1 data set is $\pm 10\%$

and the uncertainty in the calibration with 95% confidence in the L1 data set is $\pm 10\%$

and the uncertainty in the calibration with 95% confidence in the L1 data set is $\pm 10\%$

and the uncertainty in the calibration with 95% confidence in the L1 data set is $\pm 10\%$




Figure 3.21: The solid line in each plot represents the marginalised posterior pdf for h_0 for pulsar B1937+21. The 95% upper limits at the edge of the shaded regions are 2.2×10^{-21} for GEO 600, 1.4×10^{-22} for L1, 3.3×10^{-22} for H1, and 2.4×10^{-22} for H2. The dotted line for GEO 600 shows the posterior PDF for h_0 if an artificial signal was injected into the GEO 600 S1 data with $h_0 = 2.2 \times 10^{-21}$, $\phi_0 = 0^\circ$, $\psi = 0^\circ$ and $\iota = 0^\circ$.

We can recast the upper limits on h_0 as an upper limit on the equatorial ellipticity of the pulsar using Equation 3.3. Taking the distance to pulsar B1937+21 to be 3.6 kpc [65], and using our knowledge that $h_0 \leq 1.4 \times 10^{-22}$ from L1, we get

$$\epsilon \leq 2.9 \times 10^{-4} \left(\frac{10^{38} \text{ kg m}^2}{I_{zz}} \right) \quad (3.25)$$

with 95% confidence. The limit on the ellipticity from spindown measurements for this pulsar is $\epsilon \leq 3.80 \times 10^{-9}$ which is several orders of magnitude below our observational limit. It is worth noting, however, that the uncertainty in the moment of inertia of the pulsar and in the distance to the pulsar are not taken into account. For some pulsars, such as the Crab, the current observational limits allow for a moment of inertia up to a factor of two larger [70].

The uncertainty in the calibration adds 4% error-bars to the GEO600 result and a 10% error-bars to the LIGO results (as mentioned in Section 3.2). We did not fold these uncertainties into the Bayesian analysis for clarity, but in principle we could have marginalised over this uncertainty in calibration.

3.10 Conclusions

Although the results are several orders of magnitude away from astrophysically interesting results, the upper limit presented here was still the most stringent direct limit placed by a gravitational wave detector for this pulsar (at the time of S1). As we discussed in [48] two previous direct upper limits had been published on the strain of a signal from pulsar B1937+21. One previous result using data from the Caltech 40 m interferometer had set limits of $h_0 < 3.1 \times 10^{-17}$ and $h_0 < 1.5 \times 10^{-17}$ for the first and second harmonics of the rotation frequency, respectively [71]. The best previous upper limit of $h_0 < 10^{-20}$ was determined using data from a divided bar at the University of Glasgow [72].

Chapter 4

Analysis of LIGO S2 data for 28 pulsars

4.1 Introduction

The second LIGO science run (S2) took place from Feb 14, 2003 16:00:00 UTC to Apr 14, 2003 15:00:00 for a total 59 days. The sensitivity of the LIGO detectors had improved by nearly an order of magnitude over a broad frequency range compared to S1 (see Figure 4.1). With this much improved sensitivity and a quadruple increase in duration compared to S1, the S2 science run had goals of improving the upper limits set in S1, trying out new searches, and with luck detecting gravitational waves. The GEO 600 detector was being upgraded and did not participate in the S2 data run. The upgrades included the installation of the final optics as well as the signal recycling mirror both requirements for GEO 600 to reach its design sensitivity. Similarly to S1, the analysis of the data concentrated on four main signal classes (inspirals, bursts, stochastic, and periodic).

In this chapter, we will focus on placing upper limits on the gravitational wave emission from several radio pulsars. The main results from this chapter were submitted to Physical Review Letters in October 2004 (see Appendix). First, we give a brief summary of the




Figure 4.1: Typical sensitivities of the three LIGO interferometers during the S2 data run. The LIGO SRD (Science Requirements Document) lines represent the target sensitivities of the initial IFOs.

results from the other LIGO Scientific Collaboration (LSC) astrophysical searches with S2 data.

The LSC Inspirals Search Group carried out the search for waveforms associated with neutron star coalescences. More details on the neutron star inspirals analysis are available in [73]. The inspiral analysis for the S2 data run was sensitive to neutron star coalescences within a maximum distance of 1.5 Mpc, which includes the Andromeda Galaxy and other galaxies in the Local Group. In contrast, the S1 analysis for binary neutron stars inspirals would only have detected events within the Milky Way. The analysis pipeline was changed from S1 to add the additional criteria that a signal must be coincident in at least two LIGO sites to be considered a candidate. While no inspiral gravitational wave events were identified, the upper limit improved significantly over the S1 results. An upper limit was set of less than 45 inspiral events per year per Milky Way equivalent galaxy with 90% confidence for binary systems with masses between $1-3 M_{\odot}$.

The LSC Burst Search Group carried out searches in the S2 data for unmodelled bursts in the 150-1100 Hz range. Their analyses were sensitive to gravitational wave bursts with *root-sum-square* ($h_{rss} \equiv \sqrt{\int |h|^2 dt}$) amplitudes in the range $h_{rss} \sim 10^{-20} - 10^{-19}$. No excess of burst events were found when the foreground rate was compared to the estimated background rate. An upper limit on the rate of gravitational wave bursts was set at the level of 0.24 events per day. A search for potential gravitational wave events associated with Gamma Ray Burst GRB030329 was also carried out. The search did not find any excess power which is perhaps not surprising since the source was at a distance of ~ 800 Mpc. Nonetheless, an upper limit on the gravitational flux from this source was set and the tools were constructed to search for events coincident with GRBs in future runs. More details on these analyses carried out by the LSC Burst Group are available in [74] and [75].

The LSC Stochastic Search Group improved the previous direct upper limits for a stochastic background of gravitational waves by several orders of magnitude with the S2 data. In the band between 50–300 Hz an upper limit of $\Omega_{\text{GW}} h_{100}^2 < 0.018$ was set in their preliminary analysis. An analysis making use of data from the ALLEGRO bar detector is also currently in progress.

In addition to the work presented in this chapter, the LSC Periodic Search Group undertook several new analyses for the S2 run. Those include all-sky searches for periodic signals using both coherent [49] and incoherent methods [76] [77], and also a search for gravitational waves from the low-mass X-ray binary Sco X-1 [78].

In this chapter, we report on upper limits placed on the gravitational wave strain from 28 radio pulsars using S2 data. We elaborate on the selected pulsars and explain why some known pulsars were excluded from the analysis in Section 4.2. The details of the statistical methodology are given in Section 4.3. During the S2 run, artificial periodic signals were injected into the LIGO IFO's simulating the effect of gravitational waves. These hardware injections and the extraction of the signal parameters are described in

Section 4.4. In Section 4.5, we characterize the S2 data in the pertinent frequency bands for our analysis by checking the Gaussianity and stationarity of the data. The results, including coherent multi-detector upper limits, are given in Section 4.6. Finally, the significance of these results and future prospects are discussed in Section 4.7.

4.2 Pulsar selection

For the S2 run, we expanded our search for periodic gravitational waves to target all known *isolated* pulsars with a rotational frequency above 20 Hz. We only selected the pulsars with gravitational wave signals above 40 Hz because of the poor sensitivity of the interferometers at low frequencies. We excluded the binary pulsars, even though they correspond to a large fraction of rapidly rotating pulsars, saving them for future analyses. To search for signals from binary pulsars would require an extra time delay in the heterodyning function (as described in Chapter 2) to account for the orbital motion of the pulsars. We plan to extend the search to include binary pulsars with the S3 data.

There were a total of 38 known isolated pulsars with gravitational wave frequencies above 40 Hz listed in the ATNF pulsar database [65] at the time of the S2 run. At the time of writing, two further pulsars with high frequencies have been discovered but these were not included in this analysis. Jodrell Bank Observatory (JBO) was able to provide up-to-date measurements for 18 of the pulsars using radio observations taken during the S2 run [79]. These observations reduced the uncertainty in the timing and positional parameters of the sources which were outdated for some pulsars in the ATNF catalogue. The radio data also verified that none of the pulsars experienced a ‘glitch’ during the S2 run. Certain young pulsars, like the Crab and Vela, occasionally have sudden changes in their rotational frequencies known as glitches. Such a glitch would offset the gravitational wave phase and thus could undermine the analysis if not taken into account. Many pulsars also exhibit rotational irregularities known as timing noise which causes long term random deviations in phase and frequency. This timing noise is believed to be related with the rotation rate

of the stars and may be pinned to the gravitational wave phase. As with glitches, timing noise is mostly prominent and pronounced in young pulsars, such as the Crab, and its effect on millisecond pulsars can be ignored. For example, the timing noise measured for pulsar B1937+21 is on the order of $5 \mu\text{s}$ over 8 years [80] (also see [64] and Figure 3.4 in Chapter 3) and thus negligible for our purposes. On the other hand, timing noise could have introduced timing delays up to $\sim 2.1 \text{ ms}$ for the Crab pulsar (B0531+21) over the S2 run. The timing noise in the Crab pulsar was taken into account by including an extra time delay into the heterodyning function when demodulating the signal. The JBO collects radio data from the Crab pulsar on a monthly basis and those ephemerides were used to estimate the timing noise over S2. Details on the interpolation procedure used to estimate the timing noise are available in [81]. The timing noise was responsible for an extra shift in phase of 45° for the Crab pulsar over S2. For longer data runs the effect of timing noise on the Crab will be even more significant on gravitational wave searches.

Sufficiently precise timing data from radio observations were not available for all of the 38 known pulsars. For 10 of these pulsars, the uncertainties in the source parameters were large enough to lead to a timing error larger than one cycle in the phase of the signal over S2. Those pulsars were excluded from the analysis to leave a total of 28 pulsars. The timing solutions for these remaining pulsars are given in Tables 4.1, 4.2, and 4.3. The timing uncertainty came from two separate origins where the radio observations were insufficiently accurate and/or outdated:

- uncertainties in frequency and spindown (f_0 , f_1), and
- uncertainties in position (α , δ).

The dominant source of uncertainty in the spin evolution for most sources came from propagating the uncertainty in the spindown from the measurement epoch to the time of S2. In principle, we could have expanded our analysis method to search over a narrow range of frequencies and positions. However, we decided to only include the pulsars

where the uncertainty in these parameters was negligible with the intention of extending later searches to include these pulsars.

pulsar name	RA	uncertainty	dec	uncertainty	PMRA	PMDEC	epoch	source
B0021-72C	00:23:50.35311616	0.00004556	-72:04:31.4926307	0.0001846	5.1859	-3.2343	51600.0	JB
B0021-72D	00:24:13.87934042	0.00003314	-72:04:43.8405009	0.0001458	3.8431	-2.4021	51600.0	JB
B0021-72F	00:24:03.85392614	0.00005901	-72:04:42.8064817	0.0002622	4.6148	-2.8971	51600.0	JB
B0021-72G	00:24:07.95872864	0.00013286	-72:04:39.6911086	0.0003580	4.2385	-3.6734	51600.0	JB
B0021-72L	00:24:03.77085043	0.00051479	-72:04:56.9128329	0.0020867	0.0	0.0	51600.0	JB
B0021-72M	00:23:54.48766265	0.00039767	-72:05:30.7414945	0.0015327	0.0	0.0	51600.0	JB
B0021-72N	00:24:09.18645016	0.00017787	-72:04:28.8804504	0.0008903	6.7676	-1.0774	51600.0	JB
J1913+1011	19:13:20.33538156	0.00300000	10:11:23.0812785	0.0700000	0.0	0.0	52692.138313	JB
B1937+21	19:39:38.56009563	0.00000339	21:34:59.1355203	0.0000653	-0.0525	-0.4679	52328.0	JB
J1024-0719	10:24:38.68522024	0.00014944	-07:19:19.2273449	0.0017	-47.7687	-62.6836	52675.072021	JB
J1730-2304	17:30:21.66007292	0.00086043	-23:04:30.5942317	0.2298694	19.0723	0.0	52678.012687	JB
J1744-1134	17:44:29.39856726	0.00002294	-11:34:54.6364061	0.0017748	18.0869	-9.8762	52685.613159	JB
B1820-30A	18:23:40.48499977	0.00026602	-30:21:39.8456082	0.0304108	1.7123	5.7159	52684.794128	JB
B1821-24	18:24:32.00799410	0.003	-24:52:11.1180442	0.7	0.0	0.0	52718.906190	JB
B1951+32	19:52:58.16216295	0.008	32:52:39.6086014	0.11	0.0	0.0	52765.339561	JB
J2124-3358	21:24:43.85264275	0.00006	-33:58:44.6470248	0.001	-17.5744	-60.3182	52703.817261	JB
J2322+2057	23:22:22.34905276	0.00007	20:57:02.8559397	0.002	0.0	0.0	52706.584716	JB
B0531+21	05:34:31.973	0.005	+22:00:52.06	0.06	-13	7	40675	JB
J0030+0451	00:30:27.432	0.004	+04:51:39.7	0.1	0.0	0.0	50984.4	ATNF
J0711-6830	07:11:54.21813	0.00006	-68:30:47.5793	0.0004	-15.7	15.3	50425.0	ATNF
B1516+02A	15:18:33.3176	0.0006	+02:05:27.55	0.03	0.0	0.0	48445.0	ATNF
J1629-6902	16:29:08.7706	0.0004	-69:02:45.294	0.003	0.0	0.0	51600.0	ATNF
J1721-2457	17:21:05.496	0.002	-24:57:06.1	0.4	0.0	0.0	51600.0	ATNF
J1748-2446C	17:48:04.54	0.01	-24:46:36	4.0	0.0	0.0	50958.0	ATNF
J1910-5959B	19:10:52.050	0.004	-59:59:00.83	0.03	0.0	0.0	52000.0	ATNF
J1910-5959C	19:11:05.5561	0.0007	-60:00:59.680	0.007	0.0	0.0	51910.0	ATNF
J1910-5959D	19:10:52.417	0.002	-59:59:05.45	0.02	0.0	0.0	51910.0	ATNF
J1910-5959E	19:10:52.155	0.002	-59:59:02.09	0.02	0.0	0.0	51910.0	ATNF

Table 4.1: Positional information for 28 pulsars selected for S2 analysis. PMRA and PMDEC represent the proper motions of the right ascension and declination of the source, respectively.

Pulsar name	f0	uncertainty	f1	uncertainty	epoch	Source
B0021-72C	173.7082189660567053	2.64E-12	1.504341529147E-15	9.45E-20	51600	JB
B0021-72D	186.6516698568375716	3.26E-12	1.194593495604E-16	1.13E-19	51600	JB
B0021-72F	381.1586636565490949	9.02E-12	-9.370659285365E-15	2.65E-19	51600	JB
B0021-72G	247.5015250965163887	1.25E-11	2.582288885648E-15	4.85E-19	51600	JB
B0021-72L	230.0877462914992009	3.38E-11	6.461973858910E-15	1.00E-18	51600	JB
B0021-72M	271.9872287889252789	3.64E-11	2.843475770746E-15	1.00E-18	51600	JB
B0021-72N	327.4443186175468554	2.01E-11	2.344905298897E-15	5.08E-19	51600	JB
J1913+1011	27.8482474207728110	1.63E-10	-2.613648680391E-12	4.40E-17	52692.138313	JB
B1937+21	641.9282445344616690	1.78E-12	-4.331046134327E-14	6.64E-20	52328.000000	JB
J1024-0719	193.7156865914471382	2.05E-11	-6.951780847609E-16	1.64E-19	52675.072021	JB
J1730-2304	123.1102891174491276	6.29E-12	-3.058330823703E-16	3.60E-20	52678.012687	JB
J1744-1134	245.4261236024721313	6.00E-12	-5.395739544302E-16	4.22E-20	52685.613159	JB
B1820-30A	183.8234120543312997	2.74E-11	-1.142463671620E-13	1.33E-19	52684.794128	JB
B1821-24	327.4056222612654210	8.92E-10	-1.736847761699E-13	8.94E-17	52718.906190	JB
B1951+32	25.2955361136870742	5.05E-10	-3.733506994707E-12	5.96E-17	52765.339561	JB
J2124-3358	202.7938970585319680	2.81E-11	-8.451080611299E-16	1.72E-19	52703.817261	JB
J2322+2057	207.9681666656943548	7.00E-11	-4.201235730450E-16	6.40E-19	52706.584716	JB
B0531+21*	29.8102713888	2E-10	-3.736982E-10	5.1E-16	52654	JB
J0030+0451	205.53069927493	5E-11	-4.2E-16	0.9E-16	50984.4	ATNF
J0711-6830	182.117237666390	4E-12	-4.9418E-16	4E-20	50425	ATNF
B1516+02A	180.06362469727	7E-11	-1.3365E-15	2.0E-18	48445	ATNF
J1629-6902	166.64990604074	6E-11	-2.78E-16	0.09E-16	51600	ATNF
J1721-2457	285.9893480119	6E-10	-4.8E-16	0.9E-16	51600	ATNF
J1748-2446C	118.5382530563	1.5E-9	8.52E-15	0.06E-15	50958	ATNF
J1910-5959B	119.6487328457	5E-10	1.144E-14	0.008E-14	52000	ATNF
J1910-5959C	189.48987107019	1.5E-10	-7.9E-17	2.6E-17	51910	ATNF
J1910-5959D	110.6771919840	3E-10	-1.180E-14	0.004E-14	51910	ATNF
J1910-5959E	218.7338575893	4E-10	2.091E-14	0.005E-14	51910	ATNF

Table 4.2: Spin information for 28 pulsars selected for S2 analysis.

Pulsar name	f2	uncertainty	epoch(MJD)
J1913+1011	-1.407664115967E-23	4.55E-24	52692.138313
B1937+21	8.947435804829E-27	9.17E-28	52328
J1744-1134	-8.070333665728E-27	N/A	52685.613159
B1820-30A	5.211122925134E-25	N/A	52684.794128
B1821-24	1.842594529548E-28	2.70E-26	52718.906190
B1951+32	2.340810132146E-22	2.26E-23	52765.339561
B0531+21	1.2426E-20	0.0005E-20	52654.0

Table 4.3: Second spin frequency derivative for the 7 pulsars for which it is measured.

The maximum accumulated phase error that could result from uncertainties in frequency, which we denote as $\Delta\Phi_{\text{freq}}$, and position, which we denote as $\Delta\Phi_{\text{pos}}$, were evaluated separately. The uncertainties for each of the relevant parameters are given in Table 4.1 and Table 4.2. For a few of the timing solutions provided by JBO, there were no uncertainty estimates for the parameters, and in those cases we conservatively used older and larger uncertainties from the ATNF catalogue. The accumulated phase error over S2 due to uncertainties in the frequency and spin down was estimated by adding both errors in quadrature, so that

$$\Delta\Phi_{\text{freq}} = 2T [(\delta f_0)^2 + (T_{S2} - T_0)^2 (\delta f_1)^2]^{1/2}, \quad (4.1)$$

where $T = 59$ days for S2, T_{S2} is the time in the middle of the S2 run (MJD 52714), T_0 is the epoch of the spin parameters from the radio observations, δf_0 is the uncertainty in frequency, and δf_1 is the uncertainty in spindown. The factor of two in Equation 4.1 is because the gravitational wave signal frequency is at twice the rotational frequency. For 10 of the pulsars in the ATNF catalogue the accumulated phase error ($\Delta\Phi_{\text{freq}}$) over the S2 run due to frequency uncertainties was above one cycle. For these pulsars we could miss a signal that was present in the data by looking at the wrong frequency bin if we did not search over frequency. In other words, the phase of the signal would change by more than 360° over S2. The timing error was negligible for the 28 remaining pulsars as shown in Table 4.4. The largest potential discrepancy came from pulsar J0030+0451 where the accumulated phase error over S2 could be as large as 13.7% of a cycle. Such a mismatch would only have a negligible effect (a few percent) on this analysis so we kept

Pulsar Name	Spin (cycles)	Position (cycles)	Source
B0021-72C	0.000097	0.000097	JB
B0021-72D	0.000116	0.000079	JB
B0021-72F	0.000276	0.000290	JB
B0021-72G	0.000493	0.000338	JB
B0021-72L	0.001044	0.001456	JB
B0021-72M	0.001053	0.001297	JB
B0021-72N	0.000539	0.000798	JB
J1913+1011	0.001866	0.005108	JB
B1937+21	0.000029	0.000128	JB
J1024-0719	0.000209	0.005590	JB
J1730-2304	0.000064	0.006638	JB
J1744-1134	0.000061	0.000315	JB
B1820-30A	0.000280	0.003479	JB
B1821-24	0.009102	0.069760	JB
B1951+32	0.005814	0.008969	JB
J2124-3358	0.000286	0.000624	JB
J2322+2057	0.000713	0.001134	JB
B0531+21	0.027031	0.004844	JB
J0030+0451	0.137120	0.049512	ATNF
J0711-6830	0.000090	0.000046	ATNF
B1516+02A	0.007555	0.008163	ATNF
J1629-6902	0.008853	0.001160	ATNF
J1721-2457	0.088527	0.042789	ATNF
J1748-2446C	0.094059	0.067639	ATNF
J1910-5959B	0.050573	0.011851	ATNF
J1910-5959C	0.018477	0.003785	ATNF
J1910-5959D	0.028493	0.006318	ATNF
J1910-5959E	0.035645	0.012486	ATNF

Table 4.4: Cumulative phase error from uncertainties in spin and position.

pulsar J0030+0451.

The accumulated phase error due to using inaccurate positional information was calculated as follows

$$\Delta\Phi_{\text{pos}} = 2f_0 \max [T(\alpha_0, \delta_0) - T(\alpha_0 \pm \Delta\alpha, \delta_0 \pm \Delta\delta)], \quad (4.2)$$

where f_0 is the rotational frequency, $T(\alpha_0, \delta_0)$ is the time delay from the detector to the SSB from position (α_0, δ_0) , $\Delta\alpha$ is the uncertainty in RA, $\Delta\delta$ is the uncertainty in dec, and the error is the maximum between the four possible sky locations separated by $\Delta\alpha$




Figure 4.2: Antenna pattern ($\langle F_+^2 + F_\times^2 \rangle$) of LLO detector at the beginning of S2. The stars represent the position of the 28 pulsars selected for the analysis.

and $\Delta\delta$. The results of this exercise on the 28 pulsars are given in Table 4.4. The largest $\Delta\Phi_{\text{pos}}$ is for pulsar J1748+2446C where it could be as large as 6.7% of a cycle. This is small enough to have only a negligible impact on the analysis.

It is worth noting that Teviet Creighton has performed a similar analysis regarding the uncertainties as part of his role in the LSC review committee. His results also indicated no problems with the accuracy of the frequency and position for the selected 28 pulsars [82].

It is interesting to examine whether the antenna pattern of the detectors will significantly impact the sensitivity to specific pulsars. To do this we can calculate the average power, $\langle F_+^2 + F_\times^2 \rangle$, of the response function over the S2 run for different locations in the sky. In Figure 4.2 and Figure 4.3 the response function of the LLO and LHO interferometers, at the start of S2, are shown along with the location of the 28 pulsars in the sky. However, since we are analysing data over 59 days, the response function will be averaged over RA as the Earth rotates. Figure 4.4 shows $\langle F_+^2 + F_\times^2 \rangle^{1/2}$ as a function of declination along with a histogram showing the number of pulsars in each 10° intervals of declination. The average of the response function is rather smooth over declination and should not have a large impact on the sensitivity of the analysis for any of the pulsars.




Figure 4.3: Antenna pattern ($\langle F_+^2 + F_\times^2 \rangle$) of LHO detectors at the beginning of S2. The stars represent the position of the 28 pulsars selected for the analysis.




Figure 4.4: Average antenna pattern ($\langle F_+^2 + F_\times^2 \rangle^{1/2}$) of LIGO IFOs over the length of S2. The bars indicate the number of pulsars in 10 degree intervals of declination. The red line represents the average antenna pattern for LLO and the black line represents the average antenna pattern for LHO.

4.3 Analysis pipeline

The overall pipeline of the S2 analysis is very similar to the techniques used in S1. The data is heterodyned to demodulate the signal as described in detail in Chapter 2, and then we use a Bayesian approach for statistical inference. However, some changes were implemented to make the pipeline more robust to a wide range of frequencies and to take advantage of the improved stationarity of the instrumental noise.

The data were processed with tighter Butterworth IIR filters than S1 in order to suppress noise sources outside the signal band. Three third order IIR filters with cutoff frequencies of 0.5 Hz were applied after the heterodyning. Similarly to S1, the data sampled at 16384 Hz was then reduced to one estimate for the data each 60 s, which we denote B_k . The index k goes from 1 to N where N is the total number of minutes of data processed. Recall that these data points B_k are essentially a 1/60 Hz band-limited time series centered on the instantaneous frequency of the signal.

For S2, we assumed that the noise floor for each band was constant over periods of 30 minutes. As will be shown later in Section 4.5, this was a reasonable assumption for a majority of the pulsars. The noise floor in the S2 analysis is also *implicitly* calculated over a 1/60 Hz bandwidth as opposed to the explicit estimation of the noise floor from a 4 Hz bandwidth for S1. We say implicit because by using the Student's t likelihood in S2, as explained in Chapter 2, we do not need to know the variance. In accordance with the techniques presented in Chapter 2, we marginalised over the noise level using a Jeffreys prior for each 30 minute segments of data. The likelihood of each 30 minute segment of data, labelled by j , is given by

$$p(\{B_k\}_j | \mathbf{a}) \propto \left(\sum_{k=k_1(j)}^{k_2(j)} |B_k - y_k|^2 \right)^{-m}, \quad (4.3)$$

where y_k is the signal model, B_k are the data points, and $m = k_1(j) - k_2(j) + 1 = 30$ is the number of data points. In principle, we could dynamically adjust m to reflect the

stationarity of the data over the run. By limiting the analysis to continuous 30 minute segments we have excluded shorter stretches of good data. In practice, these shorter segments represented less than 12% of the S2 science mode data and were left out for the sake of simplicity.

Because of increased confidence in the timing between sites we were able to carry out a coherent multi-detector analysis with the S2 data. Using the data from all three LIGO IFOs, the global multi-detector likelihood is

$$p(\{B_k\}_{\text{Joint}}|\mathbf{a}) = p(\{B_k\}_{\text{H1}}|\mathbf{a})p(\{B_k\}_{\text{H2}}|\mathbf{a})p(\{B_k\}_{\text{L1}}|\mathbf{a}), \quad (4.4)$$

from which we can calculate the multi-detector posterior distribution for each for the parameters. Similarly to the S1 analysis, we have used uniform priors for the angles over their respective ranges, and a flat prior for $h_0 \geq 0$. In principle, we could use the results from S1 as a prior probabilities for pulsar B1937+21. In practice, however, the improvement in sensitivity in S2 is large enough to make this irrelevant. The likelihood function with the S2 data will not change significantly whether the S1 results are used as a prior or a flat prior is used.

4.4 Hardware injections

We will use the term *hardware injections* for simulated signals that are added to the interferometer control system to make the instrument behave in the same way as if a gravitational wave signal was present. During the S2 run, two hardware injections were carried out that mimicked periodic signals from pulsars in the LIGO interferometers. These injections provide the only true end-to-end validation of the search codes and the data acquisition pipelines. While *a posteriori* software injections can provide some confidence in the search codes, they can overlook subtle problems such as dynamic range issues.

Name	P1	P2
RA (rads)	5.1471621319	2.345678901234567890
Dec (rads)	0.3766960246	1.23456789012345
Frequency (Hz)	1279.123456789012	1288.901234567890123
Spindown (Hz/s)	0.0	-10^{-8}
Phase ϕ_0 (rads)	0.0	0.0
Polarisation angle ψ (rads)	0.0	0.0
GW amplitude h_0	2×10^{-21}	2×10^{-21}
Inclination angle ι (rads)	$\pi/2$	$\pi/2$
SSB Epoch (GPS)	733967667.026112310	733967751.522490380

Table 4.5: Signal parameters for the S2 hardware injections.

Through periodic hardware injections we can increase our confidence that the timing between sites is consistent. This is especially important for coherent multi-detector searches where a large lag in timing between the detectors could severely reduce the sensitivity of the search. Indeed, during the E7¹ data run, a hardware injection of a stochastic signal in both sites revealed a large *negative* correlation in the data between sites. That analysis, by the Stochastic Upper Limit Group, indicated that there was a relative sign ambiguity between the LIGO detectors [83]. The group was expecting a large positive correlation since the signal injected at both sites was supposed to be equivalent. The Periodic Group injected its first two periodic signals in S2 and did not inject any signals in the E7 run. In this section we describe the signals that were injected and report on the results obtained.

The S2 pulsar hardware injections were carried out during a 12 hour period on 10 April, 2003. The same signals as would be seen from an astrophysical source, with the correct time of arrival offsets between sites, were injected into L1, H1, and H2. Approximately 9 hours of science data for each detector was collected with the detectors in lock most of the time. Two simulated pulsar signals were injected and their parameters are listed in Table 4.5. The first pulsar signal, P1, had a constant frequency of ~ 1279 Hz with an antenna pattern consistent with a source in the direction of pulsar B1937+21. The second signal, P2, had an initial frequency of ~ 1289 Hz but was spinning down at a rate of -10^{-8} Hz/s from an arbitrary direction in the sky. The spindown rate for P2 was

¹LIGO and GEO commissioning run in December 2001 prior to the S1 science run.




Figure 4.5: Envelope of amplitude of the S2 hardware injections in both interferometers. The solid lines represent the envelope for signal P1 and the dash lines represent signal P2.

chosen to be large enough so the signal would shift by at least one frequency bin over the observation period. Envelopes of amplitude of the injected signals into each IFO are shown in Figure 4.5. Since the inclination angle is $\iota = \pi/2$ for both of the signals, the cross polarisation term is zero. So these envelopes of the injected signals are essentially just the absolute value of the plus polarization response of the interferometer properly scaled.

The phase of the signals was properly modulated to simulate a Doppler shifted signal from those directions in the sky. In fact, an initial problem in extracting the signals from the S2 data was due to a mix up of the RA and dec in the analysis software. The observation time is sufficiently long that the Doppler shifts must be taken into account.

The injection data was analysed using both the S1 procedure estimating the noise over a 4 Hz bandwidth, and also using the S2 analysis pipeline. Both methods gave qualitatively similar results with correct extraction of the signal parameters. Figure 4.6 and Figure 4.7




Figure 4.6: Recovered parameters for signal P1 of the hardware injections. The dotted line represents the signal parameter injected. The red curves are for H1 data, the green curves for H2 data, and the blue curves are for L1 data.

show the marginalised probability distribution functions for each of the signal parameters.

The greatest achievement of these hardware injections was to verify that the phase of the signal and the timing of the data acquisition system were consistent between sites. With this assurance in hand we can confidently perform coherent multi-detector analyses for signals in the S2 data. The dashed black lines in Figure 4.6 and Figure 4.7 show the marginalised pdfs using data from all three LIGO IFOs coherently. As expected, the confidence in each of the parameters was improved when the data from all three IFOs were used.




Figure 4.7: Recovered parameters for signal P2 of the hardware injections. The dotted line represents the signal parameter injected. The red curves are for H1 data, the green curves for H2 data, and the blue curves are for L1 data.

4.5 Characterisation of S2 data

4.5.1 Assessing Gaussianity and stationarity

The confidence in results from any analysis is only as good as your knowledge of the behaviour of the underlying noise. It is thus imperative to characterize the data and appreciate potential biases in the results due to wrong assumptions about the noise. The effect of making a wrong assumption, however, can in some cases have negligible consequences on the end result.

This section describes several studies that were performed to assess the Gaussianity and stationarity of the S2 data. Here the S2 data refers to the band-limited 1/60 Hz set of data points, $\{B_k\}$, heterodyned at the gravitational wave frequency. We show in Section 4.6 that even if the segments of data with non-normal characteristics are excluded from the

data set, the final upper limits do not change by more than about 30% (in the worst case).

4.5.2 Issues with applying S1 pipeline on S2 data

In the S1 analysis we assumed that the data were consistent with a Gaussian distribution which was reflected in our choice for the likelihood function of the data. For each sample of demodulated data, B_k , we estimated its variance using a band-limited 4 Hz time series heterodyned around the signal frequency. Using 240 points for the calculation of the variance for each B_k , we assumed that the uncertainty in the estimate of the variance was small. We tested the assumption that the data was consistent with being drawn from a Gaussian distribution, with $\mu = 0$ and $\sigma = \sigma_k$ for each of the B_k , by calculating

$$\chi^2 = \sum_k [\Re(B_k)^2 / \sigma_{\Re\{B_k\}}^2 + \Im(B_k)^2 / \sigma_{\Im\{B_k\}}^2]. \quad (4.5)$$

If the data are from a Gaussian distribution, we would expect χ^2 as defined above to be approximately N and with a standard deviation of $\sqrt{2N}$ where N is the total number of data points (real + imaginary). Figure 4.8 shows χ^2/N for several pulsars, for each detector, if we simply repeat the same analysis procedure as described in the S1 analysis on S2 data. If the data were normal we would expect this statistic to be approximately one with a standard deviation of $\sqrt{2/N}$. It is clear from Figure 4.8, however, that when using the S1 procedure most bands are far from being consistent with being normally distributed. Only in a few bands at high frequencies, such as for pulsar B1937+21, are the data consistent with being Gaussian.

The failure of the S1 procedure to properly track the noise for all frequency bands can be attributed to two main causes. First, the cutoff frequencies for the IIR filters used in S1, 128 Hz and 4 Hz, that were meant to prevent aliasing from outside the band of interest were not adequate. Since the noise floor is very large at low frequencies, the filters were not sufficiently attenuating the noise from outside the band for signals at low frequencies.




Figure 4.8: χ^2/N for all isolated pulsars when the noise is estimated from a 4 Hz bandwidth with the S1 IIR filter cutoffs.

For the bands at high frequencies the noise floor is quite flat and thus those wide low-pass filters were not a problem. As a result, with the S1 filters the variance of the $\{B_k\}$ were not correctly estimated, especially at low frequencies where the gradient in the noise floor is largest. Second, since our tightest IIR filter had a cutoff frequency of 4 Hz, it did not prevent any strong spectral lines within that 4 Hz from biasing our analysis. While the band for B1937+21 was relatively flat at such a high frequency, this was not the case for several other pulsars. Such a strong spectral line in the band of pulsar J1910-5959B in the S2 data is shown in Figure 4.9. These problems were fixed in the S2 analysis by using narrower band to estimate the set of $\{B_k\}$ and their variances $\{\sigma_k^2\}$.

4.5.3 Modified reduced chi-square test

For the S2 analysis we applied three third order low-pass IIR Butterworth filters each with a cutoff frequency of 0.5 Hz after demodulating the phase of the signal. Such a tight bandpass prevented the contamination of our data from nearby spectral lines but also limited our ability to assess the variance on a short time scale (e.g. each minute




Figure 4.9: Fourier transform of a 4 Hz band around the signal frequency for pulsar J1910-5959B in the H2 data. The dashed line indicates twice the spin frequency of J1910-5959B.

for S1) since the sampling rate had to be reduced. For S2, we left out the intermediate 4 Hz sampling rate to calculate the variance of the data and directly reduced the sampling rate from 16384 Hz to one B_k per minute. For the S2 data, we have assumed that these data were Gaussian within that 1/60 Hz band and that the noise floor could be considered constant (stationary) over 30 minutes (instead of each minute as in S1). For the purposes of characterising this data in a similar fashion to S1, we can estimate the variance for each contiguous segment of 30 minutes. The sample variance (assumed to be constant over 30 samples) can be calculated as

$$\sigma_k^2 = \frac{1}{29} \sum_{i=k}^{i=k+30} (B_i - \bar{B}_k)^2. \quad (4.6)$$

where \bar{B}_k is the average of 30 data points. In order to replicate as closely as possible the Gaussianity tests from S1, the quantity we want to calculate is

$$\chi_R^2 = \frac{1}{N} \sum_k [\Re(\bar{B}_k)^2 / \sigma_{\Re\{B_k\}}^2 + \Im(\bar{B}_k)^2 / \sigma_{\Im\{B_k\}}^2]. \quad (4.7)$$

where N is the number of 30 minutes segments. While 240 samples contributed to the estimate of the variance in S1 procedure (from 4 Hz bandwidth each minute), only 30 data points contribute to the sample variance calculated from the data as described above. As a result the goodness-of-fit statistic, χ_R^2 , will be slightly larger than unity for Gaussian data. This is because a Student's t -distribution approximates a Gaussian distribution for large numbers but not for small numbers. The details of the expected value for χ_R^2 are explained below for data consistent with a Gaussian distribution.

Each sample contributing to the sum in Equation 4.7 has the form

$$t^2 = \frac{rW^2}{V}, \quad (4.8)$$

where W is from a standard normal distribution and V is from a chi-square distribution with r degrees of freedom. The average value in the numerator and denominator in Equation 4.7 cancel to give the standard distributions in Equation 4.8. The expectation value of t^2 can be used to calculate the expected value of χ_R^2 . In fact, the expectation value for χ_R^2 is simply the average of the expectation values of t^2 . We can similarly calculate the variance of χ_R^2 by evaluating the variance of t^2 .

We proceed by first calculating the expectation value of t^2 . Using Equation 4.8, and the fact that the expectation value of two independent random variables is multiplicative, we have

$$E(t^2) = E(rW^2/V) = E(W^2)E(r/V). \quad (4.9)$$

The problem is thus reduced to finding the expectation value of $E(W^2)$ and $E(r/V)$. We can calculate the expectation value $E(W^2)$ by using the moment generating function (mgf) of a Gaussian distribution. Since $\mu = 0$ and $\sigma = 1$, we have

$$M(t) = e^{\mu t + \sigma^2 t^2/2} = e^{t^2/2} \quad (4.10)$$

for the moment generating function. By definition, we can calculate the expectation value

of the n^{th} moment of a distribution by taking n derivatives of the mgf, and setting the random variable equal to zero. So we have

$$M''(0) = E(W^2) = 1. \quad (4.11)$$

To find the expectation value of t^2 we must now determine $E(r/V)$. By definition, the expectation value of a function is the integral of that function of a random variable over the probability distribution of that random variable. So we have

$$E(r/V) = \int_0^\infty \frac{r}{v} f(v) dv, \quad (4.12)$$

where v is the random variable and $f(v)$ is a chi-square probability distribution function with r degrees of freedom. We can write out

$$E(r/V) = \int_0^\infty \frac{r}{v} \frac{1}{2^{r/2} \Gamma(r/2)} v^{r/2-1} e^{-v/2} dv, \quad (4.13)$$

and by letting $y = v/2$ and $dy = dv/2$ and simplifying the expression we get

$$\begin{aligned} E(r/V) &= \int_0^\infty \frac{r}{4 \Gamma(r/2)} y^{r/2-2} e^{-y} 2 dy \\ &= \frac{r \Gamma(r/2 - 1)}{2 \Gamma(r/2)} \\ &= \frac{r \Gamma(r/2 - 1)}{2 (r/2 - 1) \Gamma(r/2 - 1)} \\ &= \frac{r}{r - 2}. \end{aligned} \quad (4.14)$$

This is reassuring since this is the well known variance of the Student's t -distribution. Thus the expected value for the reduced chi-square statistic will be

$$\langle \chi_R^2 \rangle = \frac{1}{N} \sum_{i=1}^N E(t^2) = \frac{r}{r - 2}. \quad (4.15)$$

As expected, this value will approach one as the number of degrees of freedom increase and the distribution approaches a Gaussian. The number of degrees of freedom in a 30

minute segment is 29 since we calculate the average, and so the expected value of our reduced chi-square statistic is

$$\langle \chi_R^2 \rangle = \frac{29}{29-2} \simeq 1.074. \quad (4.16)$$

For the S1 analysis, there were 240 points contributing to the estimation of the variance and thus 239 degrees of freedom. So in that case we should have expected $\langle \chi_R^2 \rangle = 1.0084 \approx 1$.

In a similar way, we can derive the expected variance of our reduced chi-square statistic. In this case, we want to evaluate

$$\text{var}(\chi_R^2) = \frac{1}{N} \text{var}(t^2) \quad (4.17)$$

where

$$\text{var}(t^2) = E(t^4) - E(t^2)^2 = E(W^4)E(r^2/V^2) - [E(W^2)E(r/V)]^2. \quad (4.18)$$

From the previous calculation of $E(t^2)$ we know the second term is equal to $(\frac{r}{r-2})^2$. We can calculate $E(W^4)$ by again making use of the moment generating function of a Gaussian distribution. If we calculate the fourth derivative of the mgf from Equation 4.10, we get

$$M^{(4)}(t) = 3e^{t^2/2} + 6t^2e^{t^2/2} + t^4e^{t^2/2}, \quad (4.19)$$

so that

$$M^{(4)}(0) = E(W^4) = 3. \quad (4.20)$$

The last remaining term to calculate to get the variance in Equation 4.18 is $E(r^2/V^2)$. We can calculate the expectation value $E(r^2/V^2)$ by evaluating

$$E(r^2/V^2) = \int_0^\infty \frac{r^2}{v^2} \frac{1}{2^{2/r} \Gamma(r/2)} v^{r/2-1} e^{-v/2} dv, \quad (4.21)$$

which, using the definition of the gamma function in a similar way to Equation 4.14, can

be reduced to

$$E(r^2/V^2) = \frac{r^2}{(r-2)(r-4)}. \quad (4.22)$$

Therefore, from Equation 4.18, we get that the variance of the reduced chi-square statistic can be expressed as

$$\text{var}(\chi_R^2) = \frac{1}{N} \left[\frac{3r^2}{(r-2)(r-4)} - \left(\frac{r}{r-2} \right)^2 \right], \quad (4.23)$$

where N is the total number of 30 minute segments.

Figure 4.10 shows the reduced chi-square values using for all 28 pulsars selected for the analysis. For H1, only one pulsar (B0021–72C) has a χ_R^2 value, 0.80, that lies more than 3σ from the expected value. For H2, pulsar B1516+02A has a reduced chi-square value of 0.95 which is slightly below the 3σ level. And for L1, two pulsars B0531+21 and B1516+02A, are below the 3σ level with reduced chi-square values of 0.90 and 0.77, respectively. Generally having reduced chi-square values that are too small suggests that the noise level is overestimated. This should translate into slightly more conservative upper limits than if we were correctly modelling the noise.

4.5.4 Kolmogorov-Smirnov (Lilliefors) normality test

The reduced chi-square calculations in the previous section provide information regarding the Gaussianity of the set of \bar{B}_k over over the run. Ultimately, however, if we use the Student's t likelihood shown in Section 4.3 we want to know whether the individual B_k 's in each 30 minute block are normal. It is promising that the set of \bar{B}_k appear to be normally distributed, but that is not surprising due to the central limit theorem. Deviations from normality are more likely to reveal themselves in the non-averaged B_k 's.

In this subsection, we apply a Kolmogorov-Smirnov (K-S) test to assess whether the data in contiguous segments of 30 minutes are normally distributed. The K-S test can be used

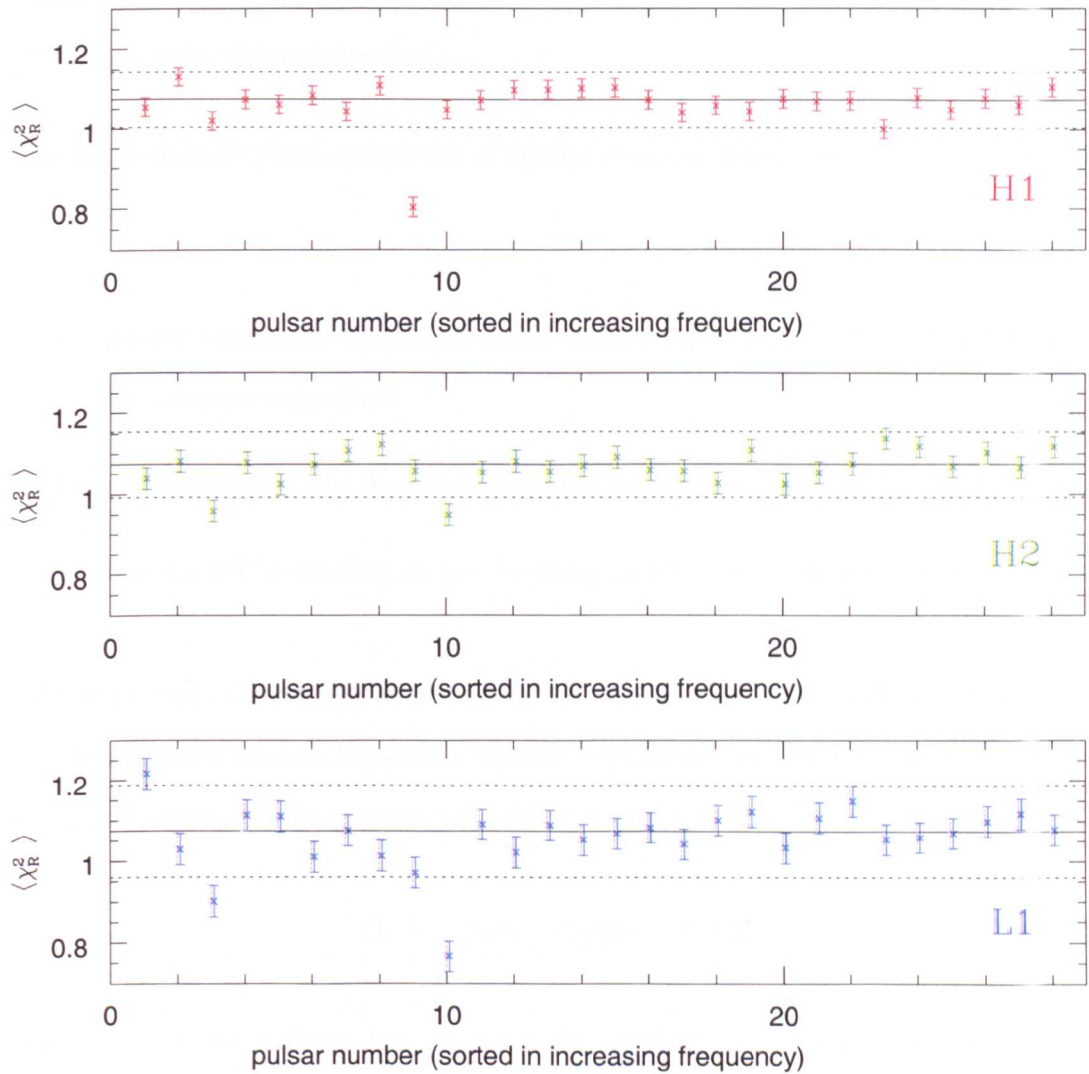


Figure 4.10: $\langle \chi_r^2 \rangle$ for each detector when the noise level is estimated every 30 minutes from the B_k 's ($\frac{1}{60}$ Hz band). The error bars represent one standard deviation, the solid line represents the expected value, and the dashed lines show the 3σ level.

with small sample sizes and is generally more powerful than the chi-square test. The K-S statistic is based on comparing the cumulative distribution functions (cdf) of two distributions. If the two sets of data are from the the same distribution then the deviations between the cdfs should be small. The K-S test is useful because it is a method that quantizes the difference between the two cdfs and provides a means to make statistical statements about the two distributions.

The cumulative distribution function of 30 B'_k s (real and imaginary) can be constructed as follows:

1. estimate variance of the 60 data points - assuming that real and imaginary parts of B_k have same noise level
2. normalize the B_k 's by the standard deviation - we call the normalized data points x
3. construct cdf for the data, $S_N(x)$, by using the fraction of data points to the left of x

Now, we want to compare the experimental cdf, $S_N(x)$, of 60 data points with a Gaussian cdf. The K-S test statistic is a simple method to quantitatively measure the difference of two distributions. It is a measure of the distance between the two cdfs and is defined as

$$D_i = \max_{-\infty < x < \infty} |S_N(x) - P(x)| \quad (4.24)$$

where x is the normalized data, $S_N(x)$ is the cdf from the data, and $P(x)$ is the cdf for a Gaussian distribution. A graphical description of the D statistic from Numerical Recipes [84] is shown in Figure 4.11.

As typically done in hypothesis testing, if the test statistic is larger than a critical value then one will reject a null hypothesis at a given significance level. In this case, the null hypothesis is that the observations are from a normal population. The critical values given for the standard K-S test, however, are for the comparison of two completely known distributions. Here the mean and variance of the original distribution were unknown and




Figure 4.11: Kolmogorov-Smirnov statistic defined as the vertical distance D between the two cumulative distribution functions, from [84].




Figure 4.12: Cumulative distribution of first 30 minute segment of S2 data for pulsar B1937+21 compared to expected cumulative Gaussian distribution (red). The vertical distance between the two distributions is less than the critical value of 0.114 for all three detectors. This period of data for each IFO is thus consistent with a Gaussian distribution.

had to be estimated from the data. When normalized by the sample standard deviation the data points, x , are approximately Gaussian. The critical values, however, will thus be slightly different from the case where both distributions are completely described. The Lilliefors critical values were calculated specifically for this problem, when the mean and variance of a normal distribution are unknown, and are available in [85]. The critical value for a sample size of 60 for a level of significance of 0.05 is $D = 0.114$. The K-S statistic D_i was calculated for every 30 minute segment of S2 data for each pulsar. Figure 4.12 shows a Gaussian cdf compared to the cdf obtained from the first locked 30 minutes of data for pulsar B1937+21. For this specific time period and pulsar, the calculated K-S statistic was less than the critical value in data from all three detectors.




Figure 4.13: Cumulative distribution of first 30 minute segment of S2 data for the Crab pulsar B0531+21 compared to expected cumulative Gaussian distribution (red). For this specific 30 minute segment, the cumulative distribution for H2 and L1 is not consistent with a Gaussian distribution at a level of 0.05.

Thus the observations were consistent with being normal with a significance level of 5%. As a second example, Figure 4.13 represent the cumulative distribution functions for the first 30 minute segment for the Crab pulsar in S2. In this case, for both the H2 and L1 data sets the null hypothesis that the data are from a Gaussian distribution is rejected.

For each pulsar, the complete S2 reduced data set was analysed to test if each contiguous 30 minute segments of data was normal. By definition, at a level of significance of 0.05 we would expect approximately 5% of the data to be rejected if the data was indeed normal. Figure 4.14 shows the distribution of the fraction of 30 minute segments of data rejected for each interferometer. For the majority of the pulsars between 5-10% of the data were rejected. This is quite reasonable and suggests that the data is relatively well behaved for most pulsars. For two pulsars, however, the Crab B0531+21, and B1516+02A, a much larger fraction of the data is rejected in certain detectors. Both of these pulsars are in highly contaminated frequency bands, near 60 Hz for the Crab, and near 360 Hz for B1516+02A. For H2, 66% of the data is rejected for the Crab B0531+21 and 25% of the data for pulsar B1516+02A. For L1, 52% of the data for pulsar B1516+02A is not consistent with a Gaussian distribution.

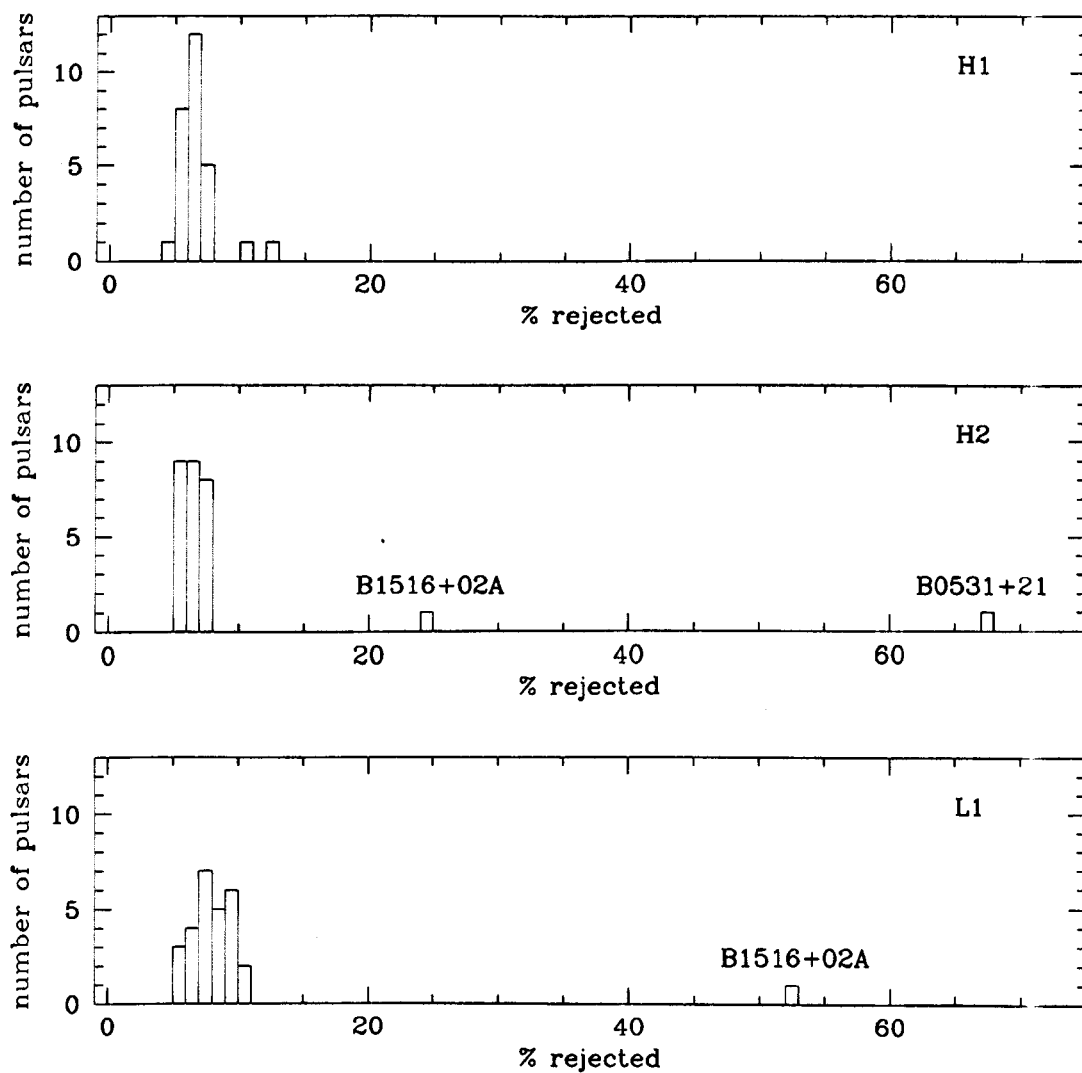


Figure 4.14: Percentage of data which is rejected at a 0.05 significance level with a Kolmogorov-Smirnov (Lilliefors) test for normality.

4.5.5 F -test for equal variances

We assume that the variance of the data is stationary over the course of the 30 minute segments. We can check this assumption by comparing the variance of the first and second 15 minutes of each 30 minute period. If the data is stationary over the course of 30 minutes then we would expect the sample variance not to change by a large amount. We used the real and imaginary components of B_k in the first 15 minutes of each 30 minute segment to calculate the sample variance, s_1^2 . Then, the real and imaginary data points for the corresponding last 15 minutes of each 30 minute segments were used to calculate the sample variance, s_2^2 .

An F -test is a standard test used to determine whether the standard deviations of two populations are equal. The F -statistic is defined as

$$F = \frac{s_1^2}{s_2^2}, \quad (4.25)$$

where s_1^2 and s_2^2 are the sample variances we want to compare. Obviously, the more this ratio deviates from one the less likely the two data sets are to have the same variance. Since the first and second half will have 30 samples contributing to the sample variance, they will both be chi-square variables with 29 degrees of freedom. The F -statistic has a so-called F -distribution with a probability function $f_{n,m}(x)$ given by

$$f_{n,m}(x) = \frac{\Gamma\left(\frac{n+m}{2}\right) n^{n/2} m^{m/2}}{\Gamma\left(\frac{n}{2}\right) \Gamma\left(\frac{m}{2}\right)} \frac{x^{n/2-1}}{(m + nx)^{(n+m)/2}}, \quad (4.26)$$

where n and m are the number of degrees of freedom (both 29 in this case).

The F -statistic was calculated for each 30 minute segment for each pulsar in each detector. Figure 4.15 shows the distribution of the F -statistics for pulsar B1937+21 compared to the theoretical function given by Equation 4.26. The agreement between the distributions is very good and thus the variance within 30 minute segments does not seem to change significantly. For the Crab pulsar, however, the agreement was not very good, especially




Figure 4.15: Distribution of F -statistics for data for pulsar B1937+21 compared to expected F -distribution in red.




Figure 4.16: Distribution of F -statistics for data for the Crab pulsar B0531+21 compared to expected F -distribution in red.

for H2. Figure 4.16 shows the distribution of the F -statistic for the Crab pulsar along with the expected distribution. The last bin to the left and right of the histograms also contains all of the counts smaller or larger than those bins, respectively. The fit for L1 is reasonable, but there is clearly an excess of outliers for H1, and H2. These are due to different noise levels, or maybe just glitches, between the first and second half of the 30 minute segments.

Figure 4.17 shows the distribution for all pulsars of the percentage of segments where the null hypothesis that the variances are equal is rejected at a significance level of 0.05. The F -test was performed on all pulsars and for the majority of those the null hypothesis is rejected for 5-10% of the data. The results seem consistent with the fraction of data rejected with the Kolmogorov-Smirnov tests for normality. We find that a large fraction of data from the Crab and pulsar B1516+02A are rejected. The fraction of data inconsistent with constant variance over 30 minutes segments is 22%, 69%, and 11% for the Crab for

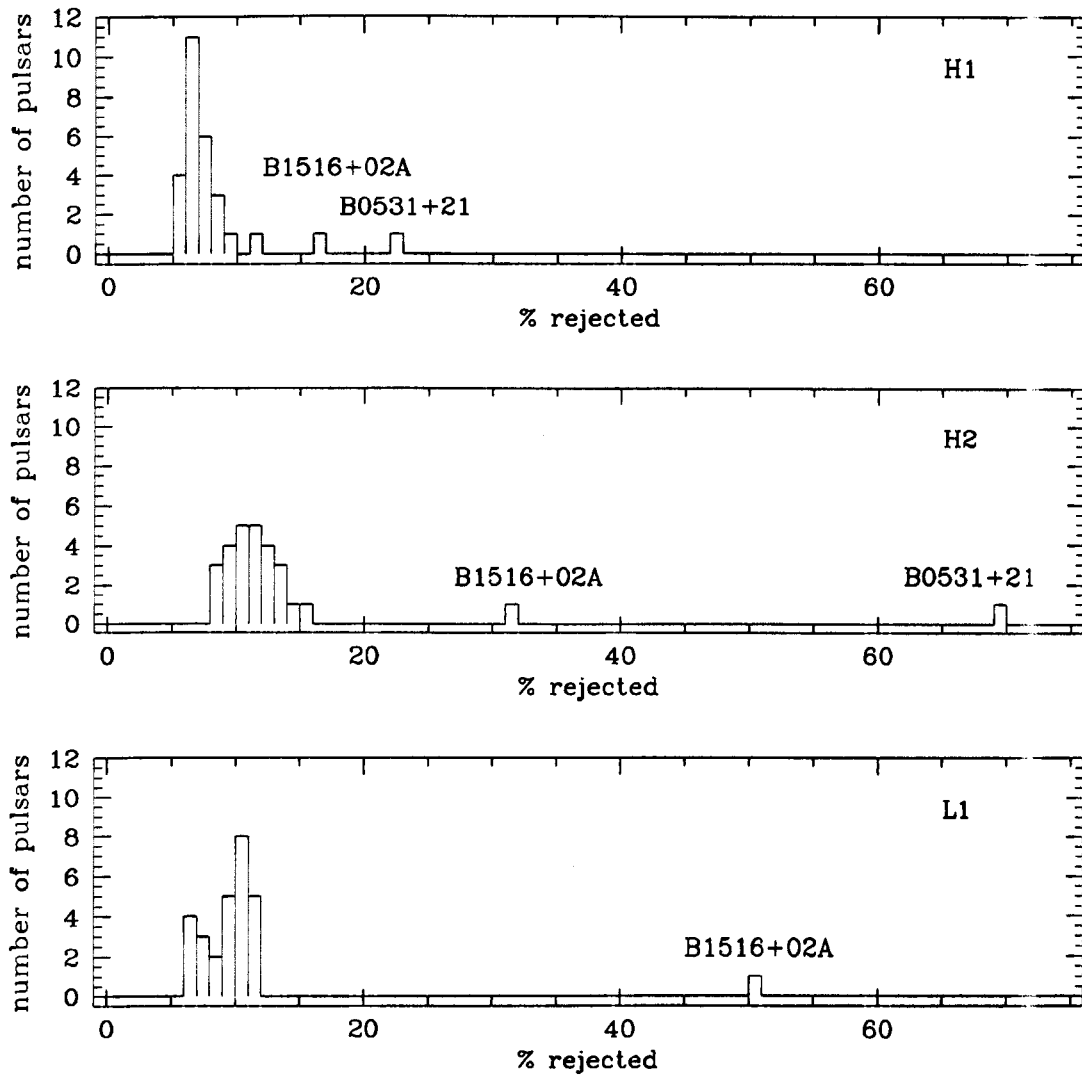


Figure 4.17: Percentage of 30 minute segments that are rejected at a significance level of 0.05 when the sample variances from the first and second 15 minutes are compared using an F -test.

H1, H2, and L1, respectively. The other peculiar pulsar is B1506+02A where 17%, 32%, and 51% of the data is rejected for H1, H2, and L1, respectively.

4.6 Results

The S2 data from H1, H2, and L1 were analysed for the 28 selected pulsars using the previously described search methods. Some data was discarded because: (i) no calibration information was available, (ii) lock stretches were shorter than 30 minutes, (iii) of

the filter impulse response in the first minute of a segment of science data and, (iv) some other data quality flag was highlighted by the operators. The results presented here made use of a total of 910 hours of data from H1, 691 hours from H2, and 342 hours from L1.

The calibration of the S2 data was implemented in the same way as the S1 analysis described in Chapter 3. The systematic errors due to calibration uncertainties were conservatively estimated to 10% and 8% for L1 and H2, respectively. For H1, the maximum uncertainty was up to 18% but typical value for a large part of the run were 6%. More detailed information of calibration studies for the S2 run can be found in [86].

Table 4.6 shows the 95% upper limits on h_0 along with the corresponding limits on equatorial ellipticity assuming a moment of inertia of 10^{38} kg m^2 . The distance to each pulsar is also given in Table 4.6.

In addition to the data vetoes already mentioned, we could also remove the data that failed the Kolmogorov-Smirnov tests in Section 4.5. Recall that the fraction of data failing the K-S test for pulsars B0531+21 and B1516+02A was above 50% for some detectors. However, since that data is non-stationary, and non-Gaussian, we would expect the likelihood of that data to have lower probability than normal data. Thus, for those data sets where we are excluding a large fraction of non-normal data, the reduction in sensitivity would not be as large as if we were excluding normal data. Indeed, the maximum reduction in the sensitivity of an upper limit for all pulsars is $\sim 30\%$ and even less for B0531+21 and B1516+02A. Figure 4.18 shows the percent difference between the upper limits on h_0 when the non-normal data are excluded. The small differences in the width of the posteriors suggest that the non-normal data are not significantly affecting our results. This is likely due to the fact that non-normal data would have low probability when using a Gaussian likelihood.

The marginalised pdfs for each of the signal parameters are shown in Figure 4.19 to Figure 4.25 for each of the pulsars. The tightest limit on gravitational wave strain comes from pulsar J1910-5959D with a 95% upper limit of $h_0 < 1.7 \times 10^{-24}$. At the time

Pulsar	H1	H2	L1	h_{95}	Dist. (kpc)	ϵ_{95}
B0021-72C	12	6.3	6.9	4.3	4.80	1.6×10^{-4}
B0021-72D	3.0	7.1	10	4.1	4.80	1.4×10^{-4}
B0021-72F	7.7	8.2	25	7.2	4.80	5.7×10^{-5}
B0021-72G	7.4	7.7	8.8	4.1	4.80	7.5×10^{-5}
B0021-72L	5.9	8.1	7.2	2.9	4.80	6.1×10^{-5}
B0021-72M	5.4	6.2	7.5	3.3	4.80	5.0×10^{-5}
B0021-72N	5.2	7.9	21	4.0	4.80	4.3×10^{-5}
J0711-6830	3.3	6.0	5.1	2.4	1.04	1.8×10^{-5}
B1820-30A	4.8	9.4	9.8	4.2	8.00	2.4×10^{-4}
J1730-2304	3.7	6.3	5.3	3.1	0.51	2.5×10^{-5}
J1721-2457	6.5	11	9.0	4.0	1.56	1.8×10^{-5}
J1629-6902	3.1	5.2	8.5	2.3	1.36	2.7×10^{-5}
B1821-24	6.5	11	15	5.6	5.80	7.1×10^{-5}
B1937+21	19	27	39	13	3.60	2.7×10^{-5}
J1910-5959E	22	11	7.9	7.5	2.15	7.9×10^{-5}
J2124-3358	5.0	6.5	5.1	3.1	0.25	4.5×10^{-6}
J1910-5959C	4.2	7.1	5.6	3.3	2.15	4.7×10^{-5}
J0030+0451	7.9	13	6.9	3.8	0.23	4.8×10^{-6}
J1024-0719	7.7	7.2	9.1	3.9	0.35	8.6×10^{-6}
J1910-5959D	3.7	7.3	2.9	1.7	2.16	7.2×10^{-5}
J2322+2057	5.5	12	5.7	4.1	0.78	1.8×10^{-5}
B1516+02A	4.1	16	13	3.6	7.80	2.1×10^{-4}
B1951+32	44	600	640	48	2.50	4.4×10^{-2}
J1748-2446C	4.7	5.9	4.3	3.1	4.56	2.4×10^{-4}
J1910-5959B	7.4	5.7	3.5	2.4	2.15	8.5×10^{-5}
J1913+1011	50	230	170	51	4.48	6.9×10^{-2}
B0531+21	33	440	99	41	2.00	2.1×10^{-2}
J1744-1134	6.4	10	8.0	5.9	0.36	8.3×10^{-6}

Table 4.6: Upper limits on h_0 ($\div 10^{-24}$) for the 28 selected pulsars using the S2 data. The column h_{95} represents the multi-detector upper limit obtained when the data from all three detectors is included in the likelihood.

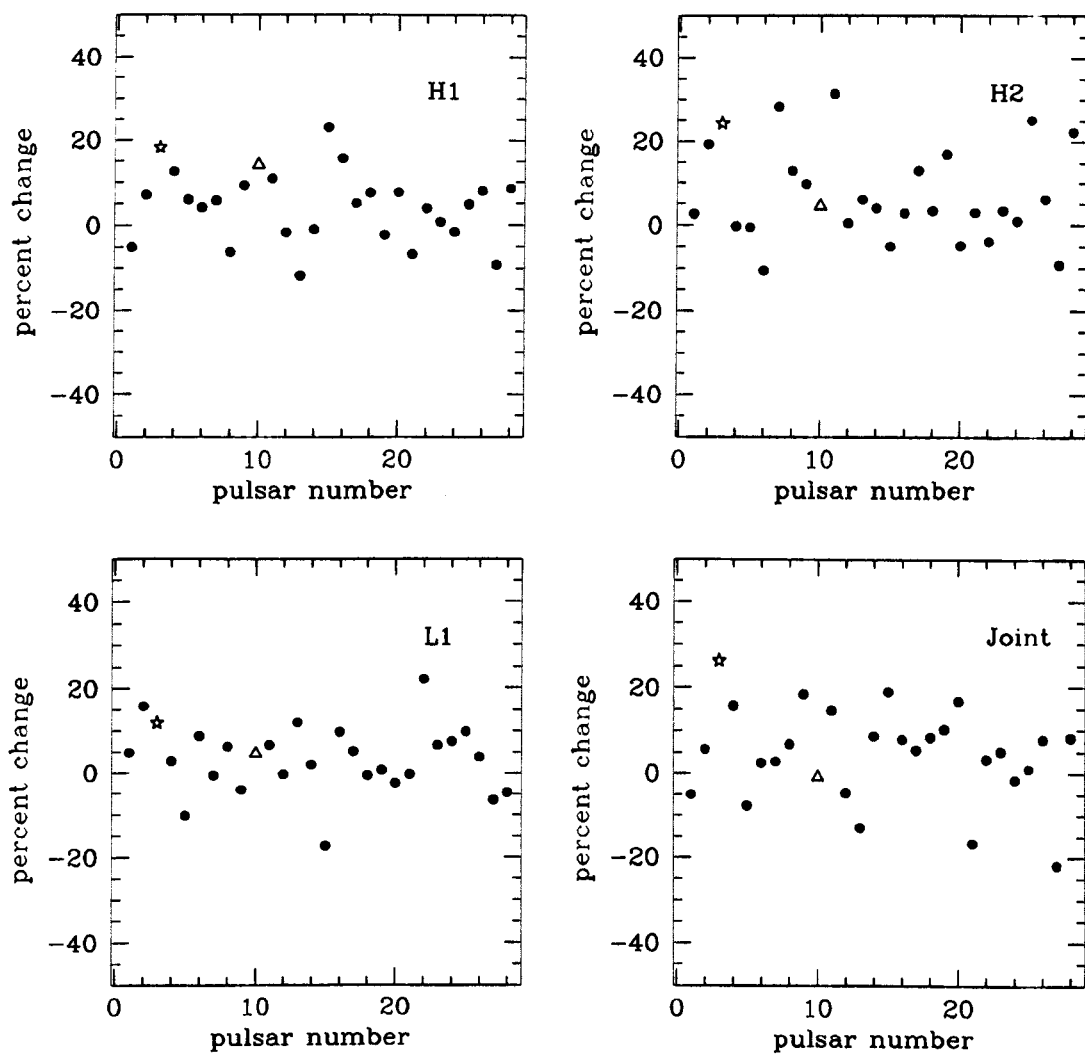


Figure 4.18: Percentage difference in h_0 upper limits if the non-normal data are excluded from the analysis. The star data point represents the Crab pulsar, the triangle represent pulsar B1516+02A, and the pulsars are ordered by increasing frequency.

of S2 this was the lowest upper limit, for an astrophysically motivated source, ever set by an interferometric gravitational wave detector. The four best limits on equatorial ellipticity came from the four closest pulsars J0030+0451, J1024–0719, J1744–1134, and J2124–3358 with 95% upper limits on the ellipticity of 4.8×10^{-6} , 8.6×10^{-6} , 8.3×10^{-6} , and 4.5×10^{-6} , respectively.

As expected the pdfs for each pulsar are consistent with there being no strong periodic gravitational wave signal at those frequencies in the S2 data. However, the limits placed on the pulsar's equatorial ellipticity for a few pulsars are approaching astrophysically interesting levels. Some of the most exotic theories of neutron star matter could actually support such a strain in the crust [87].

The observational limits set, however, are still 2 – 3 orders of magnitude above the energy conservation limits for the pulsars where we have an estimate of their intrinsic spindown. An exception is the Crab pulsar B0531+21 for which we are only a factor of ~ 30 from the spindown upper limit assuming a moment of inertia of $I = 10^{38} \text{ kg m}^2$. The spindown based upper limit for the Crab pulsar is $h_0 \leq 1.3 \times 10^{-24}$ and the observation limit is $h_0 \leq 4.1 \times 10^{-23}$. The previous best observational upper limit on the Crab pulsar was approximately an order of magnitude larger [88].

4.7 Conclusions

For 26 of the 28 pulsars, the upper limits presented here are the first reported limits focusing on these specific objects for gravitational wave emission. For the Crab pulsar B0531+21 and pulsar B1937+21, their upper limits were improved by over an order of magnitude compared to previous searches.

Nine of the pulsars studied have positive frequency derivatives and thus appear to be spinning up. The generally accepted explanation is that these pulsars are being accelerated

towards us. The pulsars in globular clusters can have large proper motions caused by strong gravitational forces. For those pulsars the changing Doppler shift, which alters the frequency more than the intrinsic spindown, gives the appearance that the pulsars are spinning up. We cannot directly measure the intrinsic spindown of those pulsars. For the pulsars in the globular cluster 47 Tucanae, such an explanation seems quite reasonable from a study of ionized gas in the cluster [89]. The pulsars from 47 Tuc are those with the prefix B0021-72, and 6 of them have a positive spindown. The other two pulsars that are spinning up are J1748-2446C and J1910-5959B. For these two pulsars, our results are slightly more interesting since we cannot constrain the gravitational wave flux, in advance, by energy conservation arguments.

In this chapter, we have expanded the search from S1 described in Chapter 3 to include 28 pulsars using the S2 data. The upper limits placed on gravitational wave emission were the most sensitive limits set for those pulsars at the time of S2. With the sensitivity of the detectors continuously improving at this stage of commissioning, we should be able to reduce the limits to physically interesting levels in the near future.

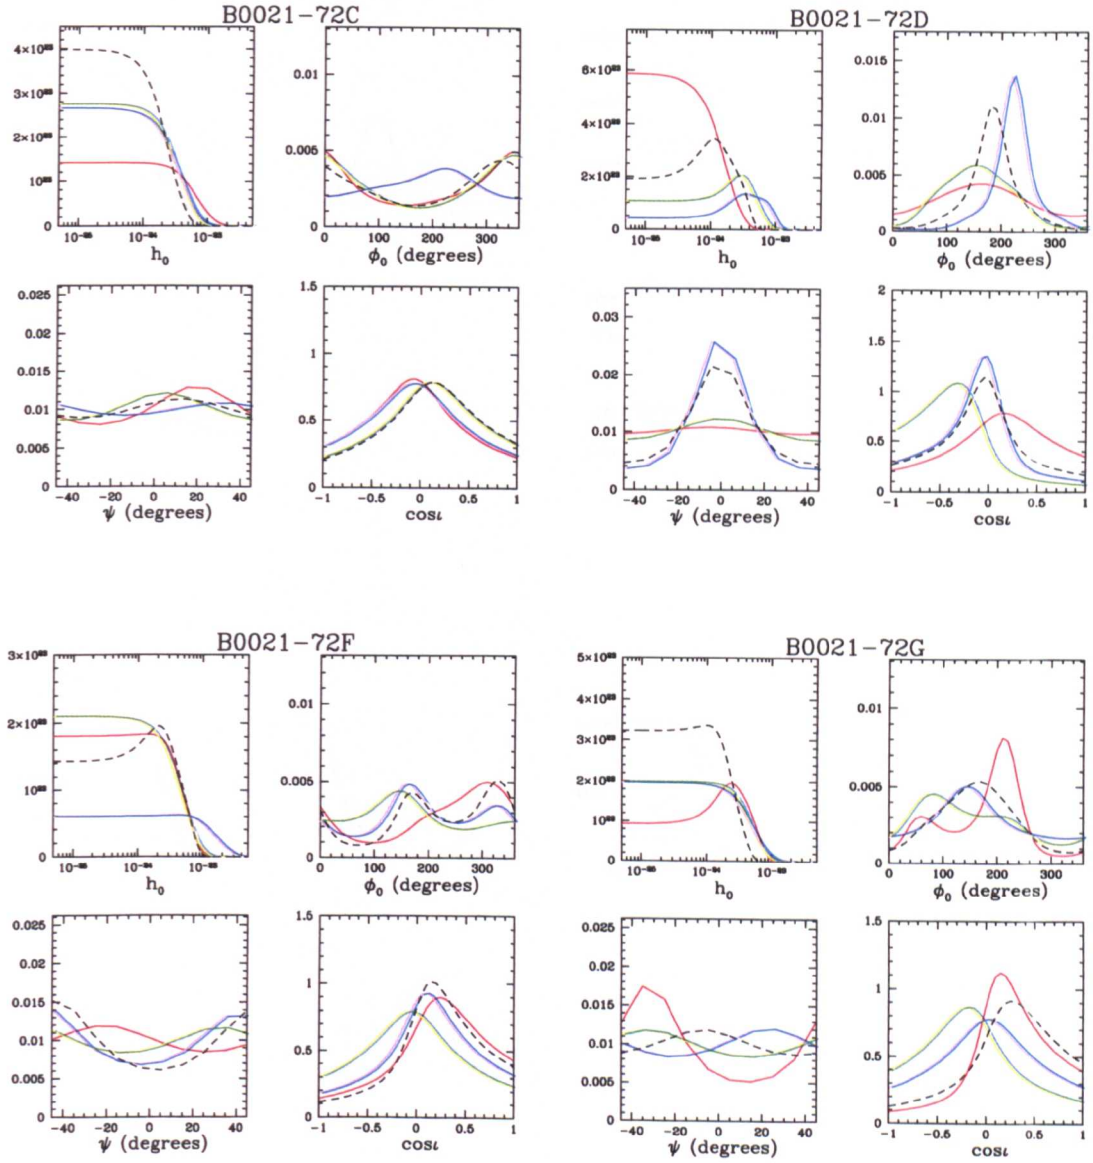


Figure 4.19: Marginalised probability distributions functions for pulsars B0021-72C, B0021-72D, B0021-72F, and B0021-72G.

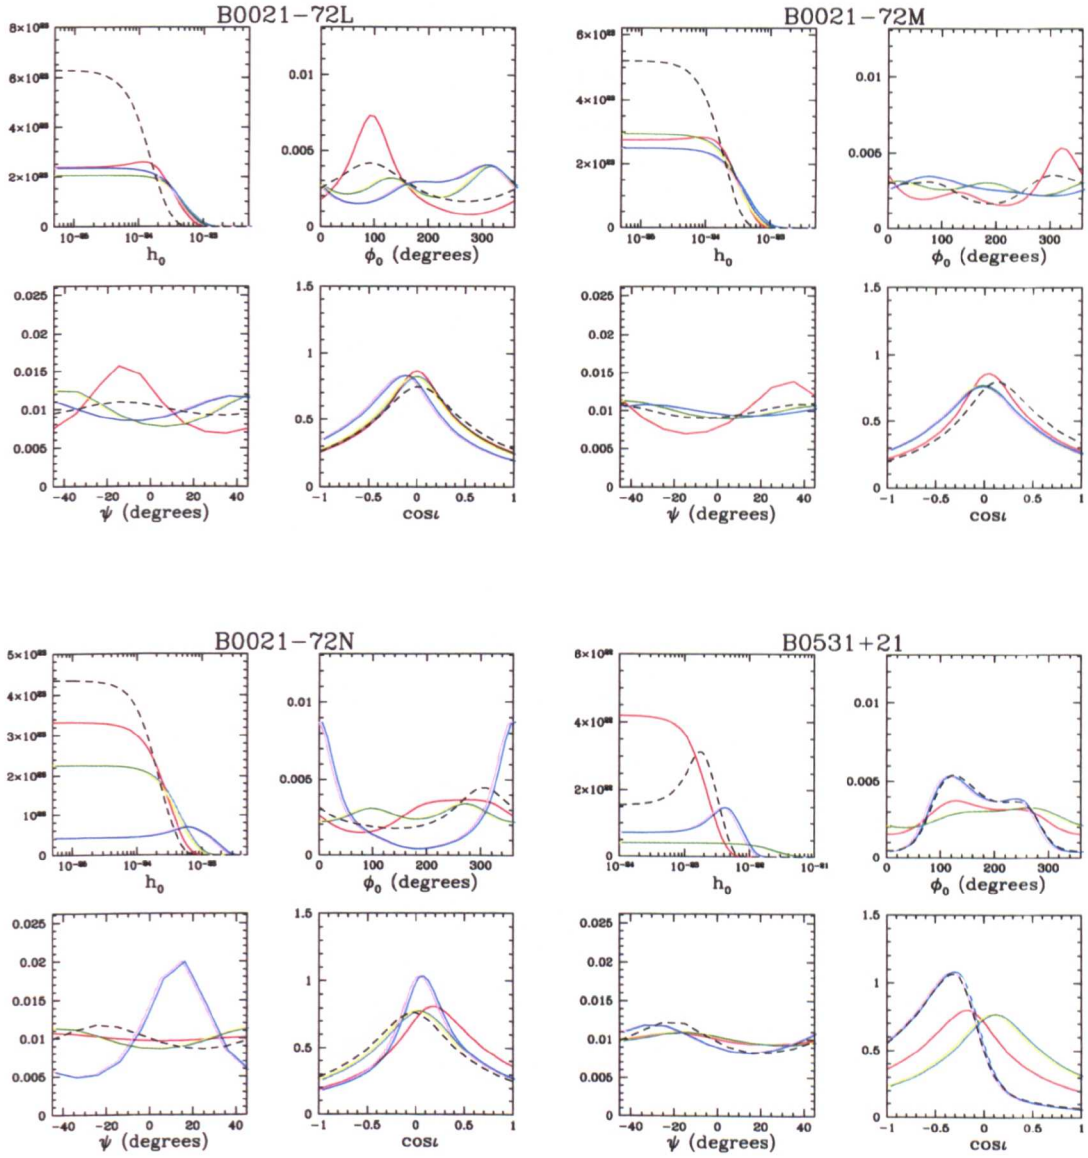


Figure 4.20: Marginalised probability distributions functions for pulsars B0021-72L, B0021-72M, B0021-72N, and B0531+21.

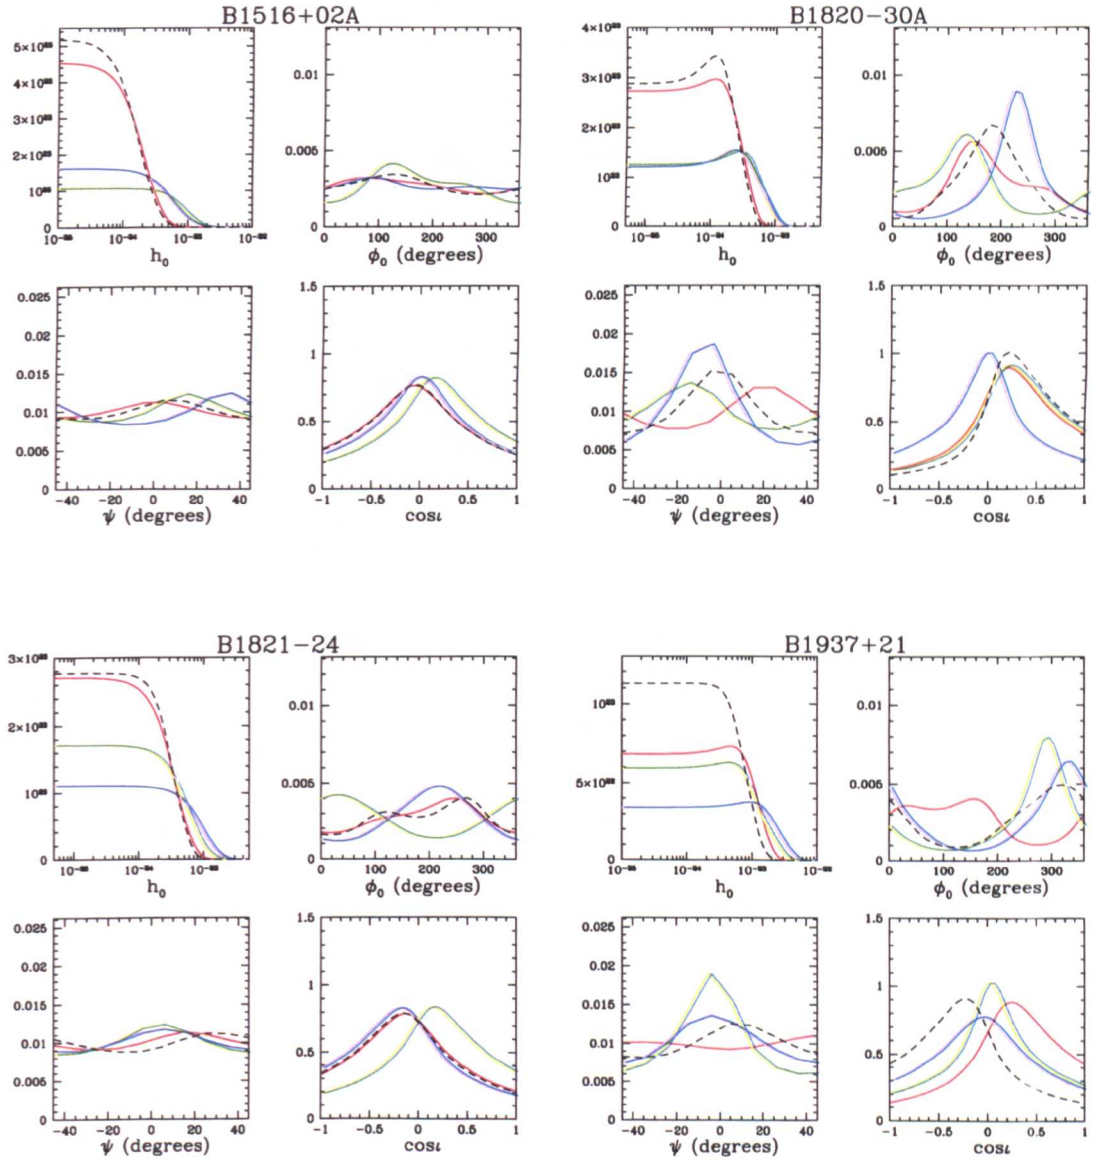


Figure 4.21: Marginalised probability distributions functions for pulsars B1516+02A, B1820-30A, B1821-24, and B1937+21.

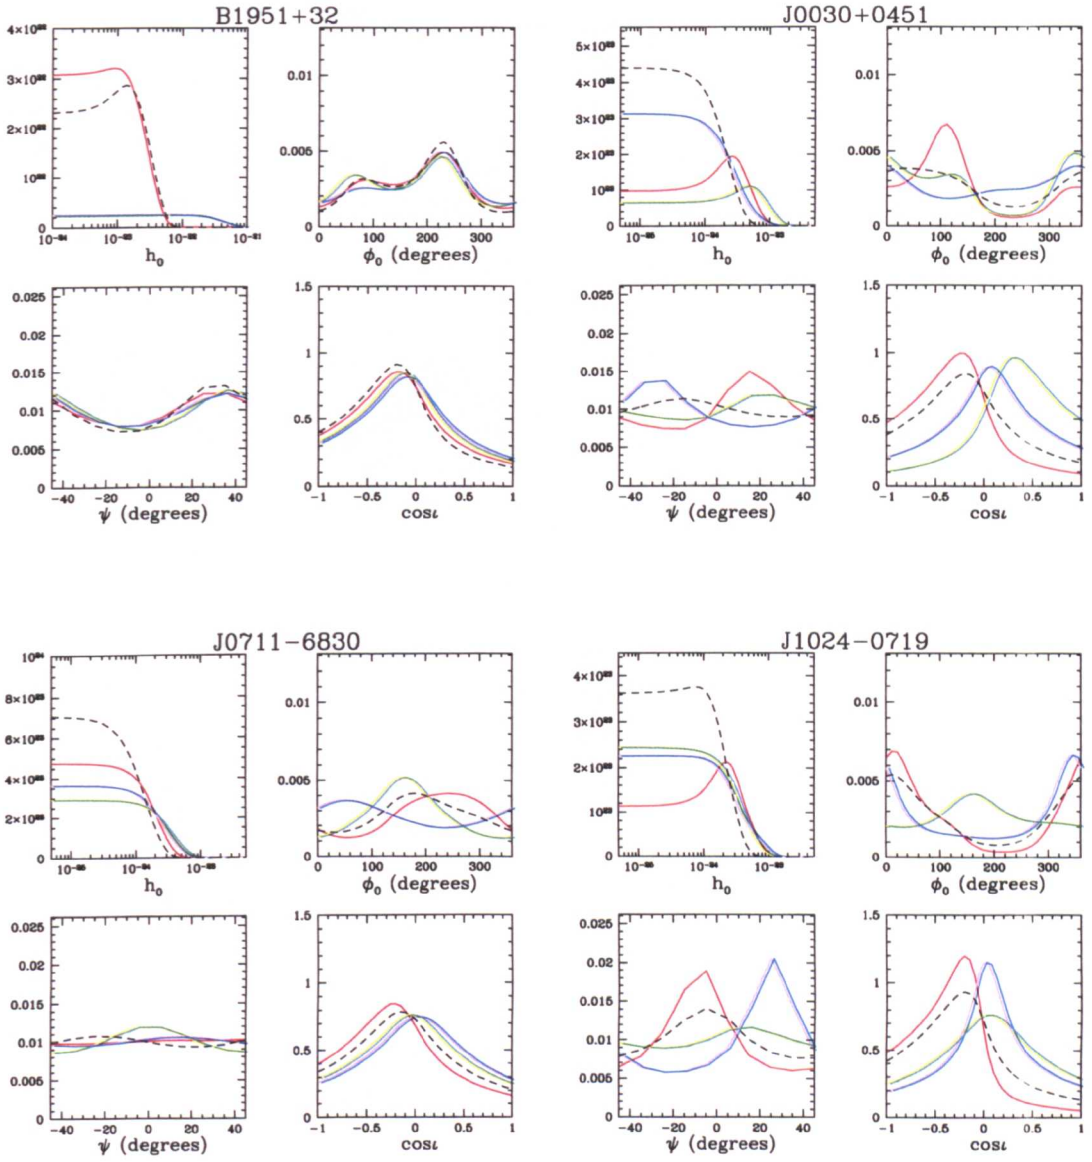


Figure 4.22: Marginalised probability distributions functions for pulsars B1951+32, J0030+0451, J0711-6830, and J1024-0719.

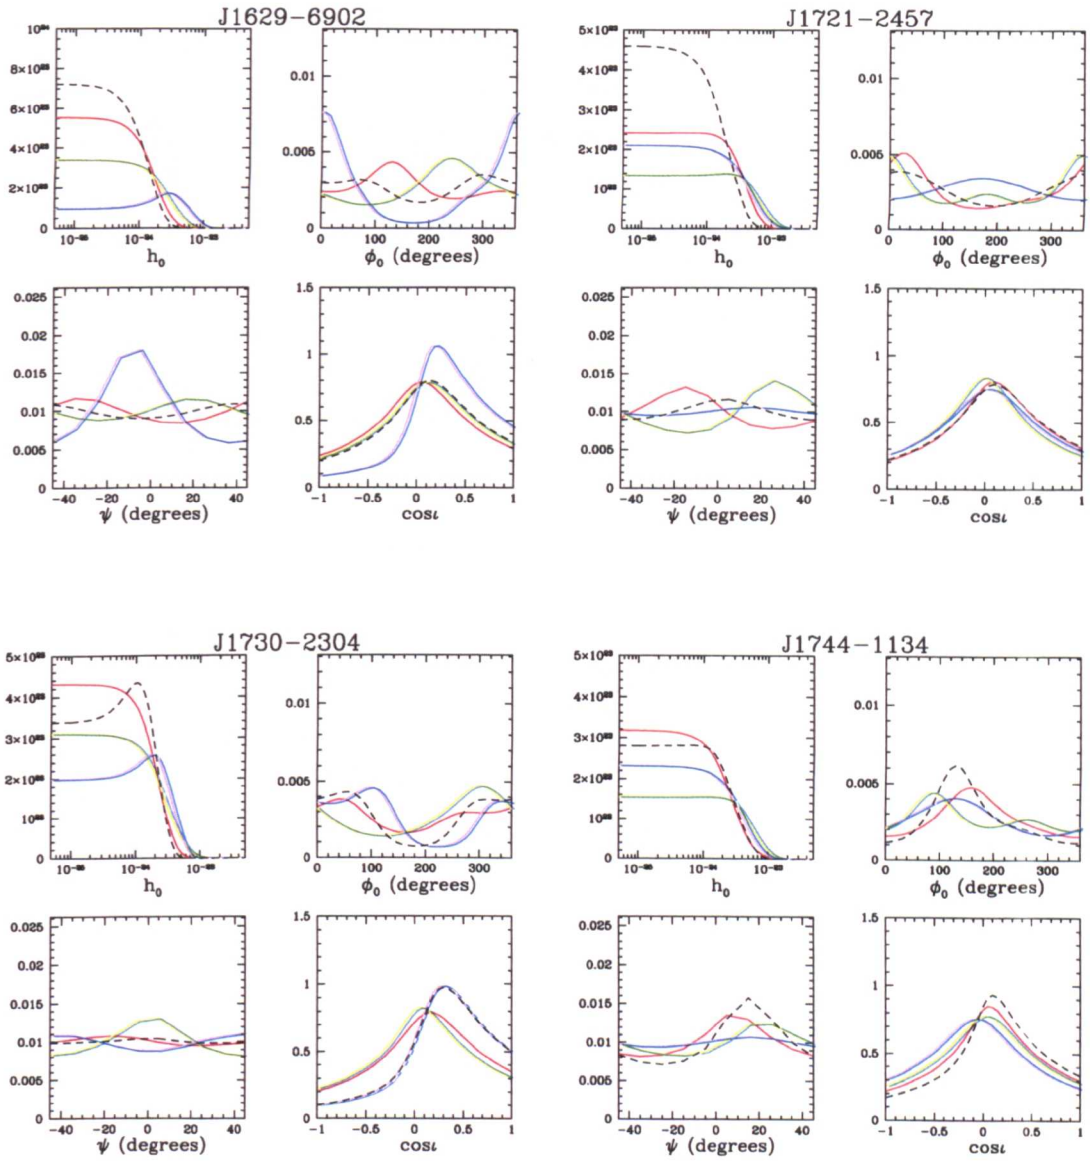


Figure 4.23: Marginalised probability distributions functions for pulsars J1629-6902, J1721-2457, J1730-2304, and J1744-1134.

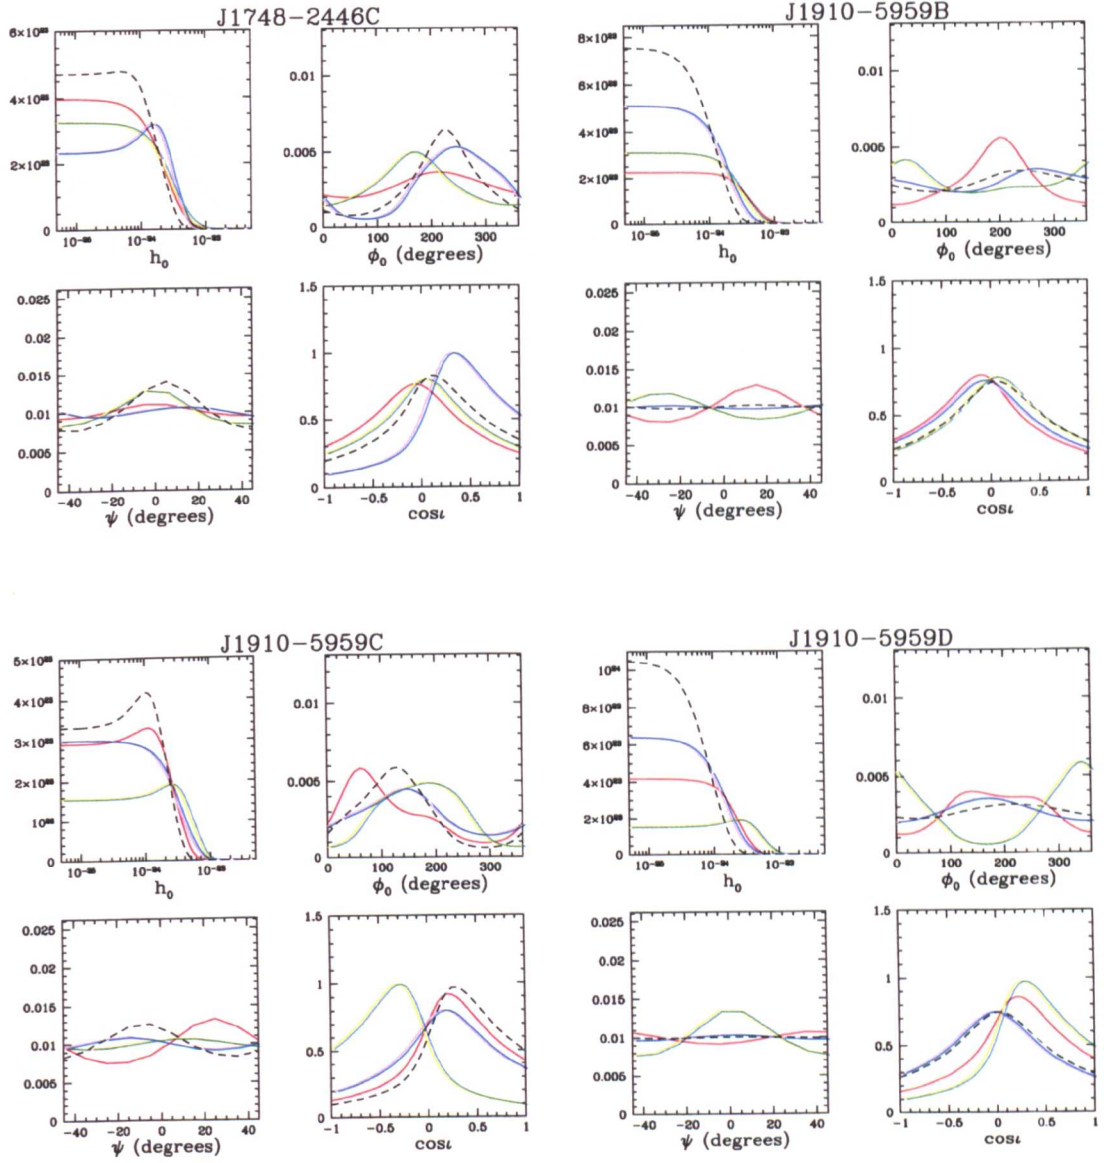


Figure 4.24: Marginalised probability distributions functions for pulsars J1748-2446C, J1910-5959B, J1910-5959C, and J1910-5959D.

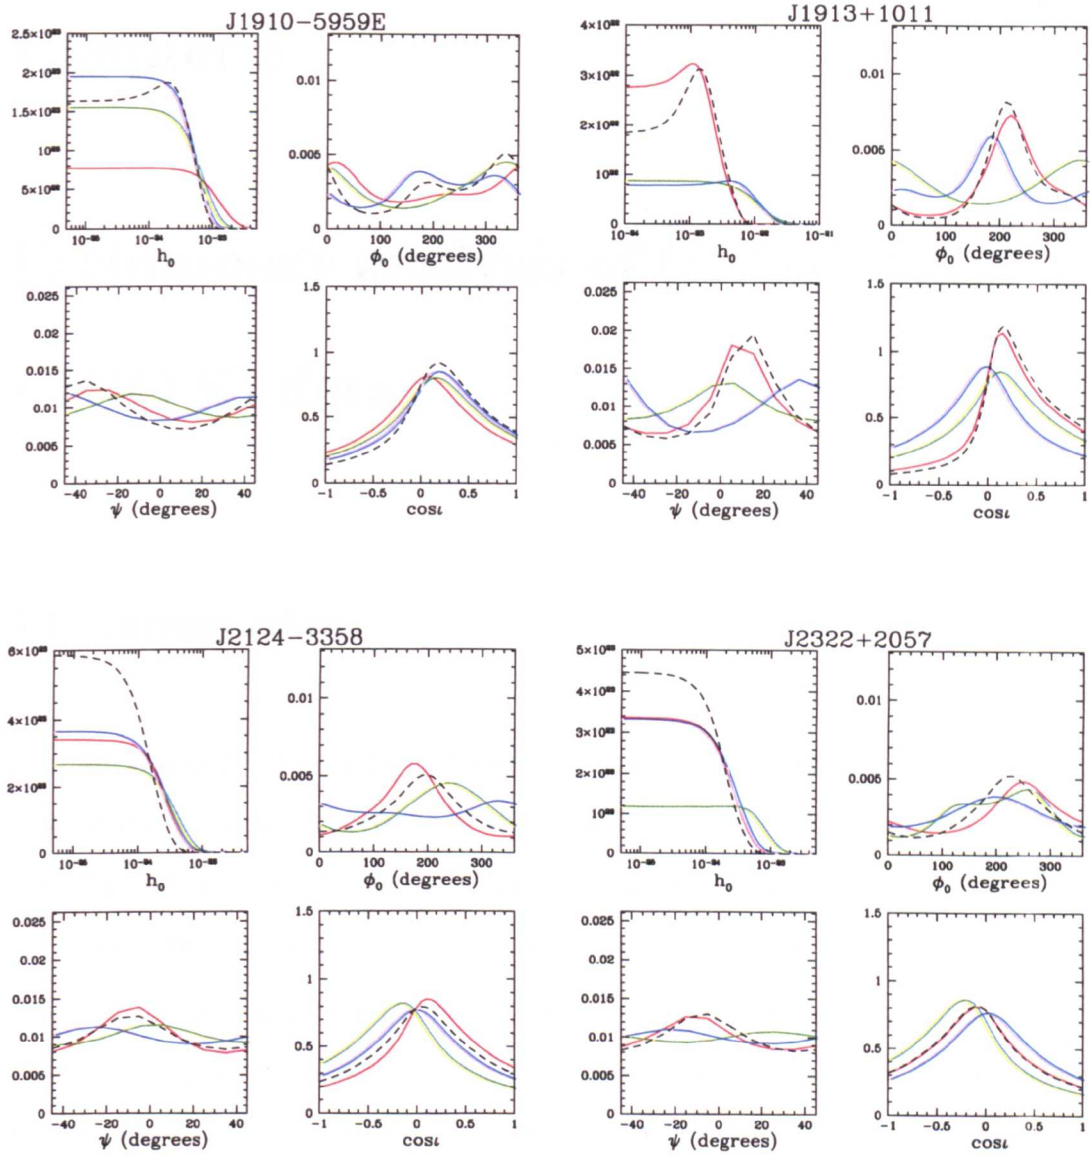


Figure 4.25: Marginalised probability distributions functions for pulsars J1910-5959E, J1913+1011, J2124-3358, and J2322+2057.

Chapter 5

Preliminary analysis of GEO 600 and LIGO S3 data

5.1 Introduction

The third science run (S3) for the LIGO interferometers took place from October 31, 2003 at 08:00 PST to January 9, 2004 at 08:00 PST giving a total of 70 days of observations. The GEO 600 detector collected three weeks of data in coincidence with LIGO during the S3 run. The GEO 600 collaboration decided to label the run as “S3”, even though it did not officially participate in the S2 run. Due to the commissioning schedule of GEO 600, its operation in science mode was divided into two segments: S3 I and S3 II. The first segment, S3 I, started on 5 November, 2003 (13:30:00 UTC) and continued up to November 12, 2003 (15:00:00 UTC). The second segment, S3 II, took place from 30 December, 2003 (15:00:00 UTC) to 13 January, 2004 (16:00:00 UTC). Figure 5.1 shows typical sensitivities for the four interferometers during the S3 run.

This chapter presents a preliminary analysis of the GEO 600 and LIGO S3 data for periodic signals from pulsars. The pulsars selected for the analysis and the analysis methods are




Figure 5.1: Typical sensitivities of the three LIGO and the GEO 600 interferometers during the S3 data run.

the same as those presented in Chapter 4. The only difference with the analysis presented here is the fact that the analysis was carried out on S3 data which were significantly more sensitive than the S2 data. The calibration used for the analyses presented here was still considered preliminary. The S2 timing solutions were used for the heterodyning and thus corrections may be required before these results are considered final. For the majority of the pulsars, the S2 timing solutions provided by Jodrell Bank Observatory should be fine for S3. However, the Crab pulsar may have experienced a glitch during the S3 run [79]. If this is confirmed, then we will need to take this into account as the gravitational wave phase could have been offset by the glitch.

Ultimately, we will use the S3 data to place upper limits on a much larger set of known pulsars. Work is currently in progress to apply the analysis presented in Chapter 4 to several binary pulsars in addition to isolated pulsars.

Name	h_0	ϕ_0	ψ	$\cos \iota$	$2f_0$ Hz	$2f_1$ Hz/s	α	δ
P0	9.4×10^{-25}	152°	44°	0.79	265.6	-4.15×10^{-12}	1.25	-0.98
P1	8.5×10^{-24}	73°	20°	0.46	849.1	-3.00×10^{-10}	0.65	-0.51
P2	1.6×10^{-24}	231°	-13°	-0.93	575.1	-1.37×10^{-13}	3.76	0.06
P3	6.2×10^{-23}	317°	25°	-0.08	108.9	-1.46×10^{-17}	3.11	-0.58
P4	1.0×10^{-21}	277°	-37°	0.28	1403.2	-2.54×10^{-08}	4.89	-0.22
P5	1.8×10^{-23}	128°	-21°	0.46	52.8	-4.03×10^{-18}	5.28	-1.46
P6	5.2×10^{-24}	56°	27°	-0.15	148.7	-6.73×10^{-09}	6.26	-1.14
P7	2.8×10^{-23}	301°	29°	0.76	1221.0	-1.12×10^{-09}	3.90	-0.36
P8	6.0×10^{-23}	337°	10°	0.07	194.3	-8.65×10^{-09}	6.13	-0.58
P9	1.6×10^{-22}	57°	0°	-0.62	763.8	-1.45×10^{-17}	3.47	1.32
P10	7.5×10^{-22}	114°	21°	0.67	1125.6	-2.87×10^{-11}	0.78	-0.62

Table 5.1: Signal parameters for the 11 periodic signals injected into the LIGO interferometers in S3. The signal P10 was simultaneously injected into the GEO 600 detector.

5.2 Hardware injections

A total of 11 hypothetical pulsar signals were injected into the LIGO interferometers during the S3 run. As described in the previous chapter, these hardware injections provide robust end-to-end checks on our data analysis pipeline. Table 5.1 shows the signal parameters for the 11 signals injected in the LIGO IFOs. The ensemble of injected signals includes signals from various location in the sky of varying amplitude. The signals were injected in a wide range in frequency from 53 Hz to 1403 Hz with spindowns as large as -2.54×10^8 Hz/s.

In addition to the LIGO injections, GEO 600 attempted its first pulsar injection during S3 II using the method described in [90]. The intent was to inject the same signal, P10, into the GEO 600 and LIGO interferometers simultaneously. The signal injected into GEO 600, P10, was the strongest of the LIGO injections. With an amplitude of $h_0 = 7.5 \times 10^{-22}$ it was chosen so that it would be visible in GEO 600 as well as the LIGO interferometers.

For the hardware injections it is especially important that the amplitude and the phase of the signals are consistent between the detectors. These are the parameters that are most likely to be affected due to problems with the signal injection mechanism or data acquisition. For this reason and since we know the values of ψ and $\cos \iota$, those parameters

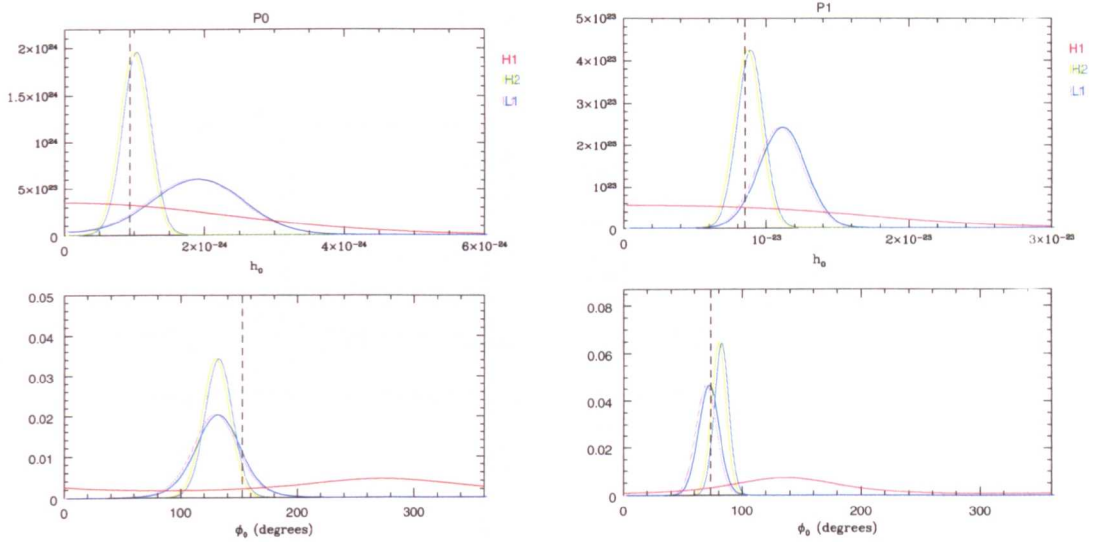


Figure 5.2: Posterior probability distributions for h_0 and ϕ_0 for the injected signals P0 (left) and P1 (right).

are taken as known parameters in the analysis. This means having prior probabilities for those parameters that have the shape of a delta function at the known values. This prevents potentially large correlations between ψ and ϕ_0 from obscuring whether the injected phases are consistent between IFOs and will also improve our inference for weak signals (since the parameters space is reduced).

Figures 5.2 to 5.7 shows the posterior probability distributions for h_0 and ϕ_0 for each of the injected signals. The injected value for each parameter is indicated with the dashed line in each figure. It can be seen from those figures that generally the signals were indeed extracted from the data when the amplitudes were sufficiently large. On the other hand, there are a few anomalous results that require further explanations.

Below we discuss each of the injected signals and attempt to explain some of the unusual results.

P0, P1, P2, P3, P5, P6, and P8 These represent 7 of the 8 weakest signals injected into the LIGO IFOs during the S3 run. These signals were injected from the beginning of S3 and were then switched off on December 1, 2003 for three weeks. All these

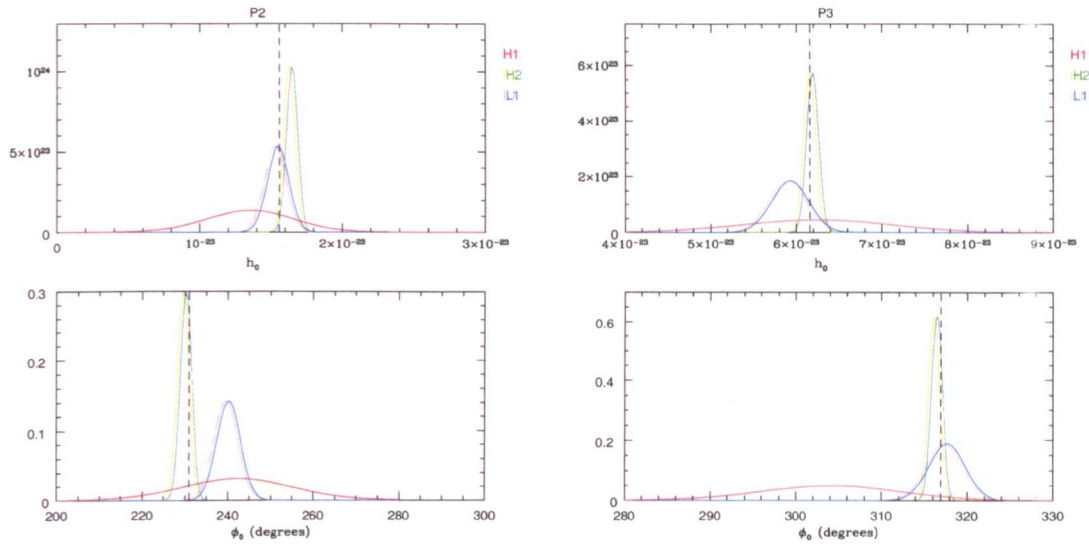


Figure 5.3: Posterior probability distributions for h_0 and ϕ_0 for the injected signals P2 (left) and P3 (right).

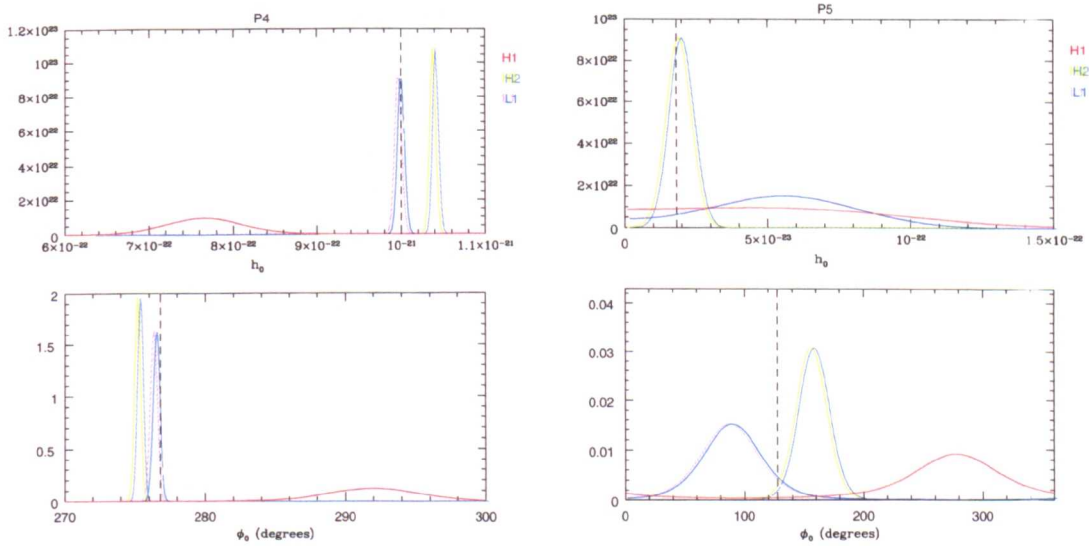


Figure 5.4: Posterior probability densities for h_0 and ϕ_0 for the injected signals P4 (left) and P5 (right).

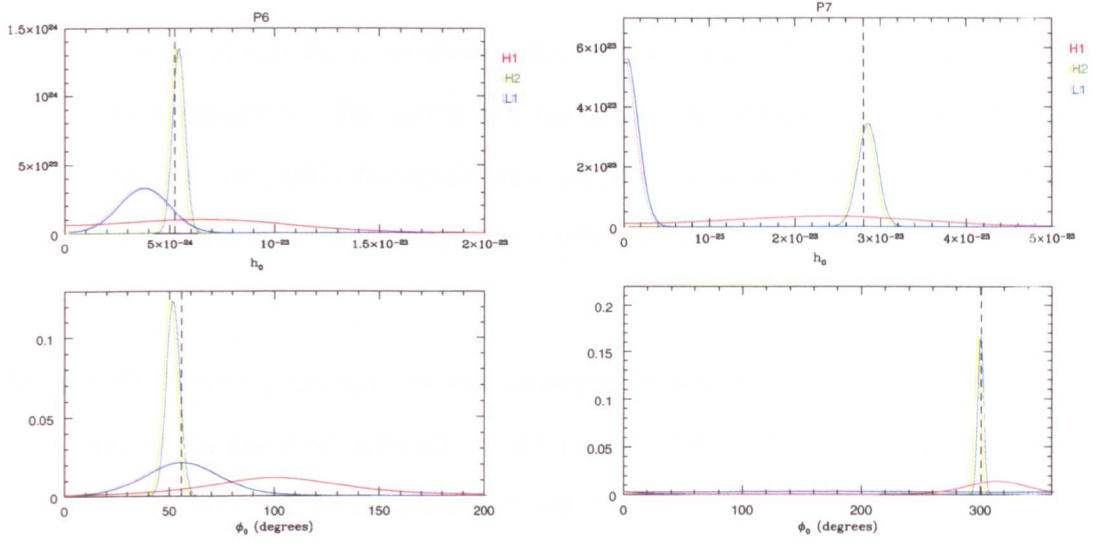


Figure 5.5: Posterior probability distributions for h_0 and ϕ_0 for the injected signals P6 (left) and P7 (right).

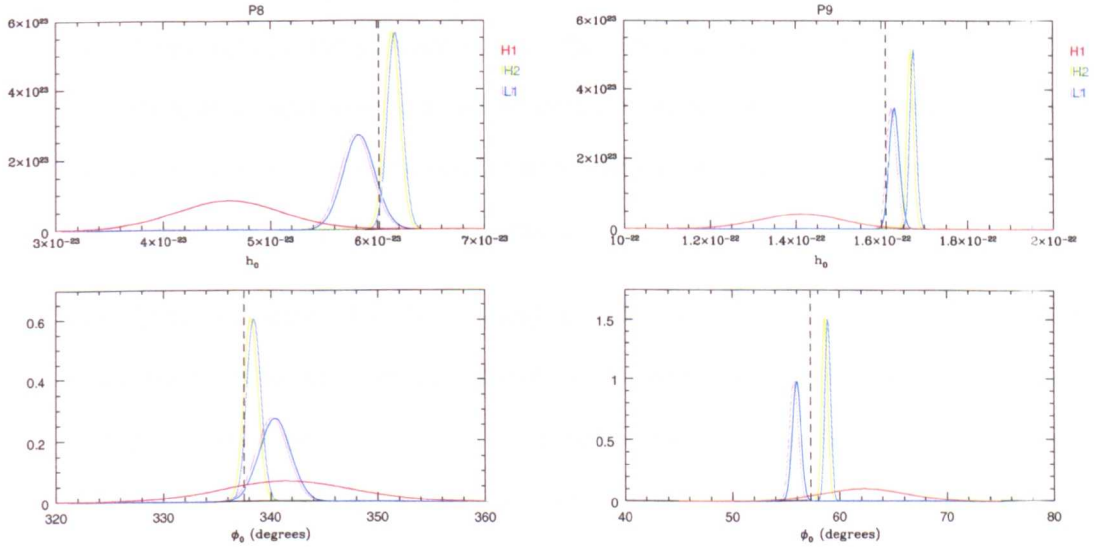


Figure 5.6: Posterior probability distributions for h_0 and ϕ_0 for the injected signals P8 (left) and P9 (right).

signals except P3 and P8 were subsequently switched on again on December 22, 2003 for the last third of the run.

For these signals the posterior distribution functions for h_0 and ϕ_0 are consistent with expectations. The actual injected parameters are typically within the pdfs though in some IFOs the signals are sometimes too weak to produce a sharp peak. But nonetheless in those cases, the injected value is still within the width of the pdfs.

P4 and P9 These signals were the two strongest signals injected in the LIGO IFOs in roughly the first third of the S3 run. P4 was switched off permanently on November 25, 2003 to avoid any unanticipated contamination of the data that may affect other analyses. P9 was switched off on December 1, 2003 for the rest of the run.

Although the pdfs are within $\sim 20\%$ of the injected values, the location and widths of the pdfs are in some cases inconsistent with the injected values. We believe that these problems are due to systematic errors in the phase and amplitude of the calibration information used to inject the signals. These systematic effects are only noticeable for strong signals where they are significant compared to the noise floor. For weak signals, the pdfs are so wide that these systematic shifts are well within the expected uncertainty and no abnormality is noticed. Since these systematic errors come from the signal injection mechanism they should not be an issue when performing the analysis for real signals.

P7 This signal was injected in the first and last third of the run along with the majority of other weak signals. While the pdfs look reasonable for H1 and H2, the signal appears to be completely absent from the L1 data. The problem does not seem to be that the L1 data is simply too insensitive to see the signal. The width of the pdf for h_0 for L1 is similar to the width of the pdf for H2 but is simply shifted near zero as if no signal was present. At the time of writing there are no obvious explanations for this result in the L1 data.

P10 This was the only signal that was injected in both the LIGO IFOs and GEO 600.




Figure 5.7: Posterior probability distributions for h_0 and ϕ_0 for the injected signals P10 using both GEO 600 and LIGO data.

The signal was injected into the LIGO IFOs from December 22, 2003 until the end of the run. It was present during the whole S3 II run for GEO 600.

Although the phases of the signals injected into LIGO agree with each other, they seem to be systematically shifted by about 10° from the injected phase. The amplitudes h_0 recovered from the LIGO IFOs are shifted by what appears to be similar systematic errors as reported for P4 and P9. These systematic effects are noticeable for this signal since this is such an unusually strong signal.

While the signal injected into GEO 600 is apparent, the recovered phase is off by over 90° from the injected phase. The recovered amplitude is about 25% lower than the expected strength. These problems were eventually tracked down to electronic pickup of the signal injection hardware into the interferometer channels. Further work, post-S3, confirmed that when these problems were resolved, the signals injected into GEO 600 were consistent with the signals injected into LIGO [91].

These hardware injections are truly remarkable experiments. The forces acting on the mirrors only displace them by $\sim 10^{-21}$ m and yet we are able (in most cases) to pre-

serve the phase between sites thousands of kilometers apart. While some problems were uncovered extracting the signals from LIGO and GEO 600, these seem to be due to the injection mechanisms as oppose to more fundamental reasons. The post-S3 periodic signals that were injected into GEO 600 were the first truly end-to-end tests that confirmed the phases between the detectors on each continent were consistent.

5.3 Characterisation of the S3 data

The peak sensitivity of GEO 600 during S3 was near 1000 Hz where it was comparable with the LIGO IFOs. Since GEO 600's sensitivity was poor at lower frequencies the preliminary analysis presented here was only carried out on pulsar B1937+21 at 1284 Hz for GEO 600. For the LIGO interferometers the same 28 pulsars (including B1937+21) as S2 were analysed.

The same statistical tests that were applied to the S2 data in Chapter 4 were applied to the S3 data streams for the 28 pulsars. The Kolmogorov-Smirnov tests indicated that for the majority of the pulsars about 5% of the data were rejected for normality at a significance level of 0.05. A few of the pulsars had larger fractions of data rejected up to nearly 50%. Table 5.2 gives the exact percentage of data failing the K-S test for each pulsar. Figure 5.8 graphically shows the distribution of percentage of rejected data.

Similarly to the S2 analyses in Chapter 4, F -tests were applied to the data to assess whether the variance in the first half and second half of 30 minute segments of data were from the same distribution. The results are given in Table 5.2 and Figure 5.9. The percentage of data rejected by the F -tests seems correlated with the amount of data rejected by the K-S test. This was also seen in S2 and is a reasonable outcome since non-stationarity data or a glitch would affect both tests.

For the GEO 600 data the same statistical tests were applied for pulsar B1937+21. A

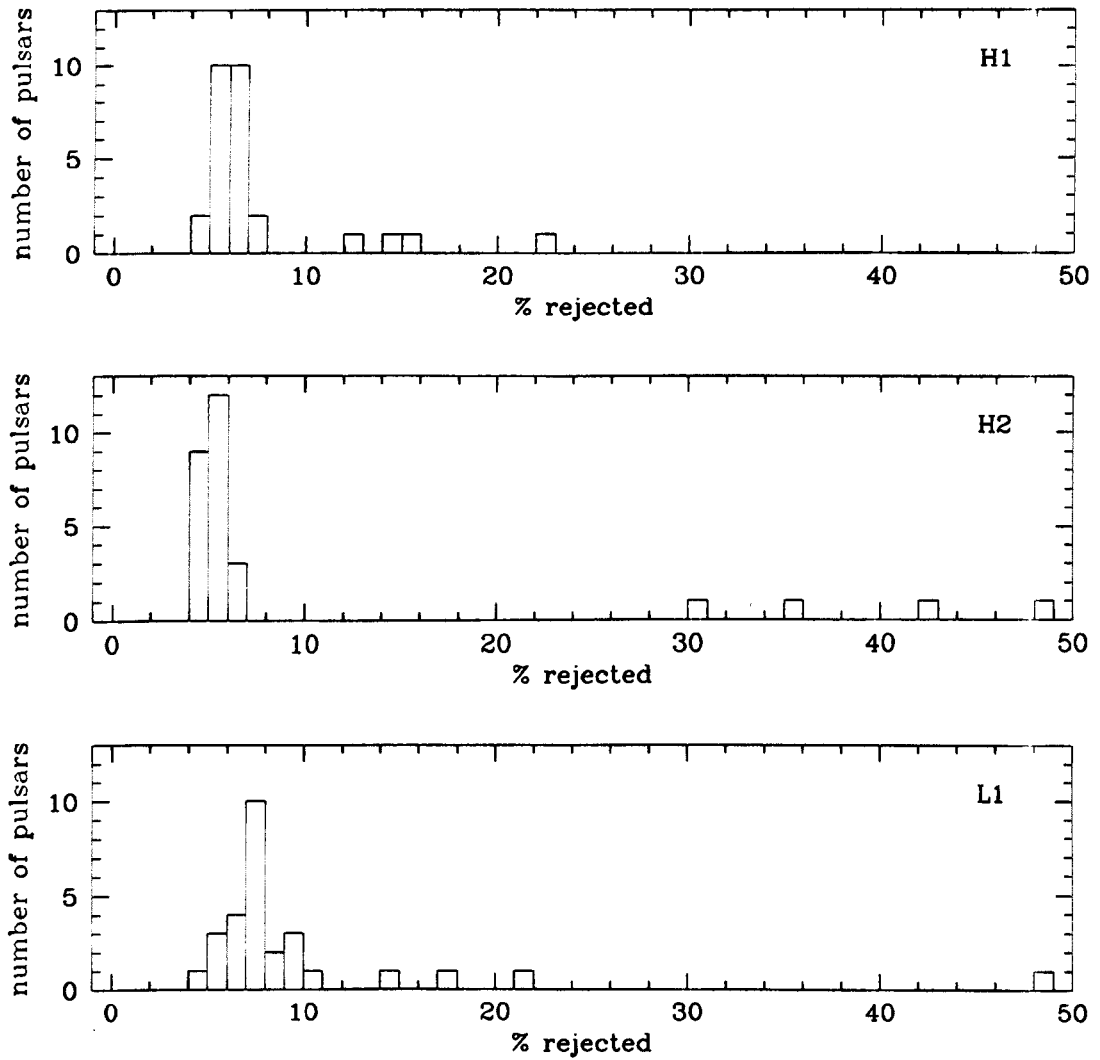


Figure 5.8: Percentage of data rejected at a significance level of 0.05 when a K-S test is applied to S3 data.

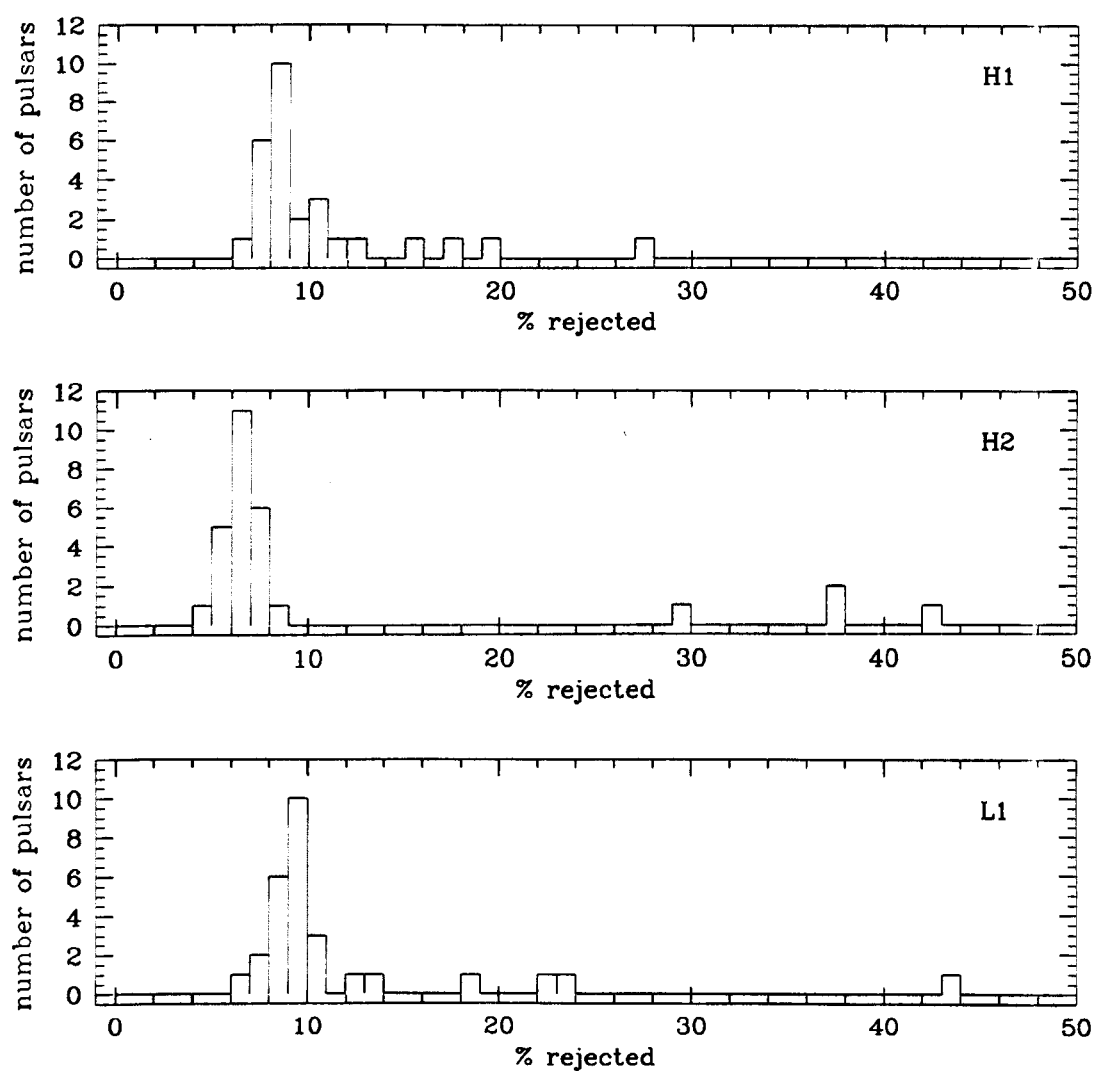


Figure 5.9: Percentage of data rejected at a significance level of 0.05 when an F -test is applied to S3 data.

Pulsar	<i>F</i> -test H1	<i>F</i> -test H2	<i>F</i> -test L1	K-S H1	K-S H2	K-S L1
B0021-72C	8.2%	7.1%	9.9%	4.9%	5.0%	6.5%
B0021-72D	8.1%	7.3%	9.3%	5.7%	5.2%	6.9%
B0021-72F	7.8%	8.7%	8.7%	6.4%	5.1%	7.7%
B0021-72G	7.9%	6.7%	9.3%	5.7%	5.7%	7.1%
B0021-72L	9.2%	6.3%	8.3%	6.0%	5.4%	8.1%
B0021-72M	8.0%	6.7%	9.7%	5.6%	5.2%	5.5%
B0021-72N	7.5%	6.3%	7.9%	6.0%	5.0%	6.7%
J0030+0451	8.5%	6.8%	6.9%	5.8%	4.7%	5.1%
B0531+21	15.9%	29.6%	18.8%	12.8%	30.8%	14.4%
J0711-6830	12.5%	7.3%	10.5%	6.7%	5.9%	7.9%
J1024-0719	8.9%	5.9%	9.9%	5.9%	4.2%	8.7%
B1516+02A	27.2%	37.6%	43.1%	22.5%	35.9%	48.6%
J1629-6902	8.6%	7.3%	10.5%	6.8%	6.5%	7.7%
J1721-2457	7.4%	6.7%	9.7%	5.6%	5.7%	7.5%
J1730-2304	11.1%	6.1%	8.9%	6.2%	6.0%	9.9%
J1744-1134	8.5%	5.7%	7.1%	6.8%	4.7%	6.5%
J1748-2446C	10.4%	7.7%	12.6%	7.0%	5.7%	9.9%
B1820-30A	8.7%	5.9%	8.5%	6.3%	5.4%	7.3%
B1821-24	8.5%	5.6%	8.3%	6.4%	5.1%	7.7%
J1910-5959B	10.4%	6.9%	13.2%	7.0%	4.6%	9.5%
J1910-5959C	9.1%	6.2%	9.3%	5.3%	6.3%	7.7%
J1910-5959D	10.7%	7.9%	10.9%	7.2%	6.8%	10.7%
J1910-5959E	8.4%	6.4%	9.5%	5.2%	4.4%	4.9%
J1913+1011	19.3%	42.3%	22.3%	14.9%	48.7%	21.9%
B1937+21	7.0%	4.8%	9.7%	4.5%	4.9%	5.7%
B1951+32	17.7%	37.8%	23.5%	15.9%	42.2%	18.0%
J2124-3358	8.0%	5.5%	9.3%	5.1%	5.0%	7.7%
J2322+2057	8.3%	6.1%	8.7%	5.7%	5.4%	7.3%

Table 5.2: Percentage of data rejected due to *F*-test and K-S test for each pulsar with S3 data with a significance level of 0.05.

total of 24.4% of the 30 min segments fail the K-S test for normality. The *F*-test rejected 29.6% of the data.

Similarly to S2, even though some of the data were non-normal, we have so far opted to keep it despite the data only having a small amount of useful information. If the non-normal data are excluded from the analysis the results are not significantly affected as shown in the next section.




Figure 5.10: Improvements in S3 upper limits on h_0 for the 28 pulsars compared to the S2 limits.

5.4 Preliminary results for 28 pulsars

The same analysis procedure as in S2 for calculating the posterior distribution function for each parameter was applied. We used the same flat priors and the Student's t likelihood function over 30 minute segments. The total amount of S3 data from each detector was: 62 610 min for H1, 53 970 min for H2, 15 150 min for L1, and 15 960 min for GEO 600 (S3 II only).

The improved sensitivity of the interferometers in S3 translated to a reduction in the upper limits on h_0 by a factors between 1 and 9 as shown in Figure 5.10.

For pulsar B1937+21 the multi-detector analysis included data from GEO 600. Figure 5.11 shows the posterior pdfs for pulsar B1937+21 for data from each of the detectors as well as the joint analysis. For B1937+21 the 95% upper limits set on h_0 were 8.6×10^{-23} for GEO 600, 8.4×10^{-24} for H1, 5.5×10^{-24} for H2, and 2.7×10^{-23} for L1. If the non-normal data rejected by the K-S tests are excluded from the analysis these upper limits on h_0 for pulsar B1937+21 change to 1.0×10^{-22} for GEO 600, 9.0×10^{-24} for H1, 5.7×10^{-24}




Figure 5.11: Posterior probability distributions for pulsar B1937+21 using S3 data from GEO 600 and the three LIGO IFOs.

for H2, and 2.6×10^{-23} for L1. The upper limits for all four detectors were only changed by a few percent when the non-normal data were excluded.

The multi-detector 95% upper limits for each of the pulsars and the corresponding limits on equatorial ellipticity are shown in Table 5.3. The limits on ellipticity were computed using the source parameters listed in Chapter 4 and a moment of inertia of 10^{38} kg m^2 .

Figure 5.12 shows the change in upper limits for all of the pulsars if the non-normal data are excluded from the analysis. The largest difference in upper limits is around 30% and on average the limits are larger when the non-normal data are excluded. Presumably the majority of the rejected data is not Gaussian and thus did not contribute much to the posterior pdf due to its low likelihood.

Pulsar	h_{95}	ϵ_{95}
B0021-72C	3.6×10^{-24}	1.4×10^{-4}
B0021-72D	1.1×10^{-24}	3.4×10^{-5}
B0021-72F	1.6×10^{-24}	1.3×10^{-5}
B0021-72G	7.0×10^{-25}	1.3×10^{-5}
B0021-72L	7.9×10^{-25}	1.7×10^{-5}
B0021-72M	1.2×10^{-24}	1.9×10^{-5}
B0021-72N	1.4×10^{-24}	1.5×10^{-5}
J0711-6830	9.4×10^{-25}	6.9×10^{-6}
B1820-30A	1.4×10^{-24}	7.8×10^{-5}
J1730-2304	6.9×10^{-25}	5.5×10^{-6}
J1721-2457	1.2×10^{-24}	5.6×10^{-6}
J1629-6902	5.5×10^{-25}	6.4×10^{-6}
B1821-24	1.1×10^{-24}	1.4×10^{-5}
B1937+21	5.6×10^{-24}	1.2×10^{-5}
J1910-5959E	8.4×10^{-25}	8.9×10^{-6}
J2124-3358	8.3×10^{-25}	1.2×10^{-6}
J1910-5959C	6.9×10^{-25}	9.8×10^{-6}
J0030+0451	1.4×10^{-24}	1.8×10^{-6}
J1024-0719	1.2×10^{-24}	2.6×10^{-6}
J1910-5959D	5.9×10^{-25}	2.4×10^{-5}
J2322+2057	1.5×10^{-24}	6.4×10^{-6}
B1516+02A	1.7×10^{-24}	9.8×10^{-5}
B1951+32	3.4×10^{-23}	3.2×10^{-2}
J1748-2446C	6.0×10^{-25}	4.6×10^{-4}
J1910-5959B	4.9×10^{-24}	1.7×10^{-5}
J1913+1011	1.0×10^{-23}	1.4×10^{-2}
B0531+21	5.5×10^{-24}	2.9×10^{-3}
J1744-1134	1.5×10^{-24}	2.2×10^{-6}

Table 5.3: Upper limits on h_0 for the 28 selected pulsars using the S3 data. The column h_{95} represents the multi-detector upper limit obtained when the data from all three detectors are included in the likelihood function.




Figure 5.12: Percent change in the S3 upper limits on h_0 if the non-normal data (from the K-S tests) is excluded from the analysis.

5.5 Conclusions

For 11 of the 28 pulsars selected, this preliminary analysis places upper limits on the ellipticity to less than 10^{-5} . For the Crab pulsar we are now within a factor of 4 from the spindown based upper limits (assuming moment of inertia of 10^{38} kg m^2). In addition to these results a few more isolated pulsars have now been discovered and will be analysed with the S3 data in the future. The S3 data will also be used to place upper limits on gravitational waves from several binary pulsars which we have excluded so far.

The results given in this chapter are to be considered preliminary as the calibration information may be improved and the timing solutions for the pulsars have not been verified with radio observations.

Chapter 6

General conclusions

A new generation of gravitational wave detectors have begun collecting data in the past few years. Though their sensitivities are orders of magnitude improved over previous detectors, the expected gravitational waves are very weak. In order to make best use of the data much effort has been devoted to developing optimised data analysis methods.

In this thesis, we have developed a new end-to-end scheme for searching for gravitational waves signals from radio pulsars in data from interferometric detectors. We took a Bayesian approach to the inference problem as it seemed the most natural interpretation of the data.

Furthermore, the developed method was applied to data from the GEO 600 and LIGO interferometers to search for periodic signals from a number of known pulsars. Ultimately, with the most sensitive data from the S3 science run, 10 of the 28 upper limits set on h_0 were below 10^{-24} . The equatorial ellipticities for several pulsars were constrained to less than 10^{-5} thus entering a zone which will soon to be astrophysically interesting.

Further work on known pulsars will attempt to place limits on all known pulsars, including binaries, using the methods described in this thesis. By the next science run, S4, we should be able to place gravitational wave upper limits on the Crab pulsar tighter than

those obtained from spindown arguments. At that point, no doubt, we will be getting closer to traditional astrophysics and closer to making a discovery.

Bibliography

- [1] A. Einstein. The foundation of the general theory of relativity. *Annalen der Physik*, **49**:769, 1916.
- [2] J. Weber. Evidence for Discovery of Gravitational Radiation. *Physical Review Letters*, **22**:1320, 1969.
- [3] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophysical Journal Letters*, **195**:L51, 1975.
- [4] B. S. Sathyaprakash. Gravitational radiation - observing the dark and dense universe. *ArXiv General Relativity and Quantum Cosmology e-prints*, gr-qc/0405136, 2004.
- [5] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. San Francisco, W.H. Freeman and Co., 1973.
- [6] B. F. Schutz. *A first course in general relativity*. Cambridge University Press, 1985.
- [7] B. F. Schutz. Gravitational Radiation. *ArXiv General Relativity and Quantum Cosmology e-prints*, gr-qc/0003069, 2000.
- [8] J. Weber. Gravitational-Wave-Detector Events. *Physical Review Letters*, **20**:1307, 1968.
- [9] P. Astone, M. Bassan, P. Bonifazi, P. Carelli, M. G. Castellano, G. Cavallari, E. Coccia, C. Cosmelli, V. Fafone, S. Frasca, E. Majorana, I. Modena, G. V. Pallottino, G. Pizzella, P. Rapagnani, F. Ricci, and M. Visco. Long-term operation of the Rome

- "Explorer" cryogenic gravitational wave detector. *Physical Review D*, 47:362, January 1993.
- [10] E. Mauceli, Z. K. Geng, W. O. Hamilton, W. W. Johnson, S. Merkowitz, A. Morse, B. Price, and N. Solomonson. The Allegro gravitational wave detector: Data acquisition and analysis. *Physical Review D*, 54:1264, 1996.
- [11] D. G. Blair, I. S. Heng, E. N. Ivanov, and M. E. Tobar. Present Status of the Resonant Mass Gravitational Wave Antenna, NIOBE. In E. Coccia, G. Veneziano, and G. Pizzella, editors, *Second Edoardo Amaldi Conference on Gravitational Wave Experiments*, page 127. Singapore, World Scientific, 1998.
- [12] P. Astone, M. Bassan, P. Bonifazi, P. Carelli, E. Coccia, C. Cosmelli, V. Fafone, S. Frasca, A. Marini, G. Mazzitelli, Y. Minenkov, I. Modena, G. Modestino, A. Moleti, G. V. Pallottino, M. A. Papa, G. Pizzella, P. Rapagnani, F. Ricci, F. Ronga, R. Terenzi, M. Visco, and L. Votano. The gravitational wave detector NAUTILUS operating at $T = 0.1$ K. *Astroparticle Physics*, 7:231, 1997.
- [13] M. Bonaldi, F. Bronzini, E. Cavallini, M. Cerdonio, A. Colombo, P. Falferi, L. Franceschini, R. Macchietto, S. Paoli, G. A. Prodi, L. Taffarello, S. Vitale, and J. P. Zendri. The ultracryogenic gravitational wave antenna AURIGA. *Physica B Condensed Matter*, 194:1, 1994.
- [14] R. Weiss. Gravitational radiation - The status of the experiments and prospects for the future. In *Sources of Gravitational Radiation*, page 7. Cambridge University Press, 1979.
- [15] D. I. Robertson, E. Morrison, J. Hough, S. Killbourn, B. J. Meers, G. P. Newton, N. A. Robertson, K. A. Strain, and H. Ward. The Glasgow 10 m prototype laser interferometric gravitational wave detector. *Review of Scientific Instruments*, 66:4447, 1995.

- [16] H. Billing, K. Maischberger, A. Rudiger, R. Schilling, L. Schnupp, and W. Winkler. An argon laser interferometer for the detection of gravitational radiation. *Journal of Physics E Scientific Instruments*, 12:1043, 1979.
- [17] R. Spero, D. Z. Anderson, Y. Gursel, R. W. P. Drever, G. Gutt, M. Hereld, J. Kaufman, and S. E. Whitcomb. The Caltech Gravitational Wave Detector. In B. Bertotti, F. de Felice, and A. Pascolini, editors, *Proceedings of the 10th International Conference on General Relativity and Gravitation*, page 930. Rome, Consiglio Nazionale delle Ricerche, 1983.
- [18] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gursel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, and K. S. Thorne. LIGO - The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256:325, 1992.
- [19] B. Willke *et al.* (GEO Collaboration). Status of the GEO600 gravitational wave detector. In M. Cruise and P. Saulson, editors, *Proceedings of the SPIE: Gravitational-Wave Detection*, volume 4856, page 238, 2003.
- [20] B. Caron *et al.* (VIRGO Collaboration). The VIRGO Interferometer For Gravitational Wave Detection. *Nuclear Physics B Proceedings Supplements*, 54:167, 1997.
- [21] K. Tsubono. TAMA300 laser interferometer for gravitational wave detection. In M. Kujawinska and M. Takeda, editors, *Proceedings of the SPIE: Interferometry '99: Techniques and Technologies*, volume 3744, page 11, 1999.
- [22] N. A. Robertson. Topical Review: Laser interferometric gravitational wave detectors. *Classical and Quantum Gravity*, 17:19, 2000.
- [23] P. R. Saulson. *Fundamentals of interferometric gravitational wave detectors*. Singapore, World Scientific, 1994.
- [24] C. Cutler and K. S. Thorne. An overview of gravitational wave sources. *ArXiv General Relativity and Quantum Cosmology e-prints*, gr-qc/0204090, 2002.

- [25] A. G. Lyne, M. Burgay, M. Kramer, A. Possenti, R. N. Manchester, F. Camilo, M. A. McLaughlin, D. R. Lorimer, N. D'Amico, B. C. Joshi, J. Reynolds, and P. C. C. Freire. A Double-Pulsar System: A Rare Laboratory for Relativistic Gravity and Plasma Physics. *Science*, **303**:1153, 2004.
- [26] V. Kalogera, C. Kim, D. R. Lorimer, M. Burgay, N. D'Amico, A. Possenti, R. N. Manchester, A. G. Lyne, B. C. Joshi, M. A. McLaughlin, M. Kramer, J. M. Sarkissian, and F. Camilo. The Cosmic Coalescence Rates for Double Neutron Star Binaries. *Astrophysical Journal Letters*, **601**:L179, 2004.
- [27] C. L. Fryer. Stellar Collapse. *International Journal of Modern Physics D*, **12**:1795, 2003.
- [28] K. C. B. New. Gravitational Waves from Gravitational Collapse. *Living Reviews in Relativity*, **6**:2, 2003.
- [29] C. L. Fryer, D. E. Holz, and S. A. Hughes. Gravitational Wave Emission from Core Collapse of Massive Stars. *Astrophysical Journal*, **565**:430, 2002.
- [30] R. W. Klebesadel, I. B. Strong, and R. A. Olson. Observations of Gamma-Ray Bursts of Cosmic Origin. *Astrophysical Journal Letters*, **182**:L85, 1973.
- [31] A. I. MacFadyen and S. E. Woosley. Collapsars: Gamma-Ray Bursts and Explosions in "Failed Supernovae". *Astrophysical Journal*, **524**:262, 1999.
- [32] M. H. P. M. van Putten. Gamma-ray bursts: LIGO/VIRGO sources of gravitational radiation. *Physics Reports*, **345**:1, 2001.
- [33] M. Zimmermann and E. Szedenits. Gravitational waves from rotating and precessing rigid bodies - Simple models and applications to pulsars. *Physical Review D*, **20**:351, 1979.
- [34] D. I. Jones. Gravitational waves from rotating strained neutron stars. *Classical and Quantum Gravity*, **19**:1255, 2002.

- [35] C. Cutler. Gravitational waves from neutron stars with large toroidal B fields. *Physical Review D*, **66**:084025, 2002.
- [36] N. Stergioulas. Rotating Stars in Relativity. *Living Reviews in Relativity*, **6**:3, 2003.
- [37] N. Andersson. A New Class of Unstable Modes of Rotating Relativistic Stars. *Astrophysical Journal*, **502**:708, 1998.
- [38] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson. Gravitational waves from hot young rapidly rotating neutron stars. *Physical Review D*, **58**:084020, 1998.
- [39] G. Ushomirsky, C. Cutler, and L. Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *Monthly Notices of the Royal Astronomical Society*, **319**:902, 2000.
- [40] I. H. Stairs, A. G. Lyne, and S. L. Shemar. Evidence for free precession in a pulsar. *Nature*, **406**:484, 2000.
- [41] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Monthly Notices of the Royal Astronomical Society*, **331**:203, 2002.
- [42] B. Abbott *et al.* (LIGO Scientific Collaboration). Analysis of first LIGO science data for stochastic gravitational waves. *Physical Review D*, **69**:122004, 2004.
- [43] M. S. Turner. Detectability of inflation-produced gravitational waves. *Physical Review D*, **55**:435, 1997.
- [44] B. Allen. The Stochastic Gravity-Wave Background: Sources and Detection. In J.-A. Marck and J.-P. Lasota, editors, *Proceedings of the Les Houches School of Physics: Relativistic Gravitation and Gravitational Radiation*, page 373. Cambridge University Press, 1997.
- [45] A. Vilenkin. Gravitational radiation from cosmic strings. *Physics Letters B*, **107**:47, 1981.

- [46] T. Damour and A. Vilenkin. Gravitational wave bursts from cusps and kinks on cosmic strings. *Physical Review D*, **64**:064008, 2001.
- [47] E. J. Copeland, R. C. Myers, and J. Polchinski. Cosmic F- and D-strings. *Journal of High Energy Physics*, **6**:13, 2004.
- [48] B. Abbott *et al.* (LIGO Scientific Collaboration). 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. *Physical Review D*, **69**(8):082004, April 2004.
- [49] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Physical Review D*, **58**:063001, 1998.
- [50] T. J. Loredo. *From Laplace to Supernova SN 1987A: Bayesian Inference in Astrophysics*. PhD thesis, University of Chicago, 1995.
- [51] L. S. Finn. Issues in Gravitational Wave Data Analysis. In E. Coccia, G. Veneziano, and G. Pizzella, editors, *Second Edoardo Amaldi Conference on Gravitational Wave Experiments*, page 180. Singapore, World Scientific, 1998.
- [52] G. D'Agostini. Overcoming priors anxiety. *ArXiv Physics e-prints*, physics/9906048, 1999.
- [53] G. D'Agostini. Role and meaning of subjective probability: Some comments on common misconceptions. In A. Mohammad-Djafari, editor, *AIP Conference Proceedings 568: Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, page 23, 2001.
- [54] J. H. Taylor and J. M. Weisberg. Further experimental tests of relativistic gravity using the binary pulsar PSR 1913 + 16. *Astrophysical Journal*, **345**:434, 1989.

- [55] T. M. Niebauer, A. Rüdiger, R. Schilling, L. Schnupp, W. Winkler, and K. Danzmann. Pulsar search using data compression with the Garching gravitational wave detector. *Physical Review D*, **47**:3106, 1993.
- [56] D. Sivia. *Data analysis: A Bayesian tutorial*. Oxford University Press, 1996.
- [57] B. Abbott *et al.* (LIGO Scientific Collaboration). Analysis of LIGO Data for GW from Binary Neutron Stars. *Physical Review D*, **69**:102201, 2004.
- [58] B. Abbott *et al.* (LIGO Scientific Collaboration). First upper limits from LIGO on gravitational wave bursts. *Physical Review D*, **69**:102001, 2004.
- [59] B. Abbott *et al.* (LIGO Scientific Collaboration). Detector description and performance for the first coincidence observations between LIGO and GEO. *Nuclear Instruments and Methods in Physics Research A*, **517**:154, 2004.
- [60] J.R. Smith *et al.* (GEO Collaboration). Commissioning, characterization, and operation of the dual-recycled GEO 600. *Classical and Quantum Gravity*, **21**:1737, 2004.
- [61] <http://lappweb.in2p3.fr/virgo/FrameL/>.
- [62] M. Hewitson, H. Grote, G. Heinzel, K. A. Strain, H. Ward, and U. Weiland. Calibration of GEO 600 for the S1 science run. *Classical and Quantum Gravity*, **20**:885, 2003.
- [63] A. Ferrari and R. Ruffini. Theoretical Implications of the Second Time Derivative of the Period of the Pulsar NP 0532. *Astrophysical Journal Letters*, **158**:L71, 1969.
- [64] A.N. Lommen and D.C Backer, private communication.
- [65] <http://www.atnf.csiro.au/research/pulsar/psrcat/>.
- [66] <http://ssd.jpl.nasa.gov/horizons.html>.
- [67] <http://www.lsc group.phys.uwm.edu/daswg/projects/lal.html>.

- [68] <http://pulsar.princeton.edu/tempo/>.
- [69] R. Adhikari *et al.* Calibration of the LIGO detectors for the First LIGO Scientific Run. *LIGO Technical Document T030097-00-D*, 2003.
- [70] M. Bejger and P. Haensel. Accelerated expansion of the Crab Nebula and evaluation of its neutron-star parameters. *Astronomy and Astrophysics*, **405**:747, 2003.
- [71] M. Hereld. *A Search For Gravitational Radiation from PSR 1937+214*. PhD thesis, California Institute of Technology, 1983.
- [72] J. Hough, H. Ward, A. J. Munley, G. P. Newton, B. J. Meers, S. Hoggan, G. A. Kerr, and R. W. P. Drever. Direct observational upper limit to gravitational radiation from millisecond pulsar PSR1937+214. *Nature*, **303**:216, 1983.
- [73] B. Abbott *et al.* (LIGO Scientific Collaboration). Search for binary neutron star coalescence in the Local Group using simultaneous observations by multiple LIGO detectors. *LIGO Technical Document P040024-01-Z*, 2004.
- [74] B. Abbott *et al.* (LIGO Scientific Collaboration). Upper Limits on Gravitational Wave Bursts in LIGO's Second Science Run. *LIGO Technical Document P040040-00-R*, 2004.
- [75] B. Abbott *et al.* (LIGO Scientific Collaboration). A Search for Gravitational Waves Associated with the Gamma Ray Burst GRB030329 Using the LIGO Detectors. *LIGO Technical Document P040007-05-D*, 2004.
- [76] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba. The Hough transform search for continuous gravitational waves. *ArXiv General Relativity and Quantum Cosmology e-prints*, gr-qc/0407001, 2004.
- [77] P. R. Brady and T. Creighton. Searching for periodic sources with LIGO. II. Hierarchical searches. *Physical Review D*, **61**:082001, 2000.
- [78] C. Messenger and A. Vecchio. Searching for gravitational waves from low mass x-ray binaries. *Classical and Quantum Gravity*, **21**:729, 2004.

- [79] M. Kramer, private communication.
- [80] V. M. Kaspi, J. H. Taylor, and M. F. Ryba. High-precision timing of millisecond pulsars. 3: Long-term monitoring of PSRs B1855+09 and B1937+21. *Astrophysical Journal*, 428:713, 1994.
- [81] M. Pitkin and G. Woan. Searching for gravitational waves from the Crab pulsar: the problem of timing noise. *Classical and Quantum Gravity*, 21:843, 2004.
- [82] <http://www.ligo.caltech.edu/pshawhan/lsc/pulreview/s2td/>.
- [83] <http://www.ligo.caltech.edu/docs/G/G020376-00.pdf>.
- [84] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C. The art of scientific computing*. Cambridge University Press, 1992.
- [85] H.W. Lilliefors. On the Kolmogorov-Smirnov test for normality with mean and variance unknowns. *American Statistical Association Journal*, 62:399, 1967.
- [86] G. Gonzalez, M. Landry, and B. O'Reilly. Calibration of the LIGO detectors for S2. *LIGO Technical Document T040060-00*, 2004.
- [87] B. J. Owen and D. I. Jones. When do the pulsar upper limits become interesting? *LIGO Technical Document T040192-00-Z*, 2004.
- [88] T. Suzuki. Search for Continuous Gravitational Wave from Pulsars with Resonant Detector. In E. Coccia, G. Pizzella, and F. Ronga, editors, *First Edoardo Amaldi Conference on Gravitational Wave Experiments*, page 115. Singapore, World Scientific, 1995.
- [89] P. C. Freire, M. Kramer, A. G. Lyne, F. Camilo, R. N. Manchester, and N. D'Amico. Detection of Ionized Gas in the Globular Cluster 47 Tucanae. *Astrophysical Journal Letters*, 557:L105, 2001.

-
- [90] U. Weiland, G. Heinzel, H. Ward, and G. Woan. Hardware injection of simulated continuous gravitational wave signals for GEO 600. *Classical and Quantum Gravity*, 21:861, 2004.
- [91] U. Weiland. *Preparing for gravitational wave astronomy: A verification of the GEO600 detection chain by generation, injection, and extraction of continuous signals*. PhD thesis, University of Hannover, submitted, 2004.

Appendix A

Related publications

The papers in the pages that follow are related to the work presented in this thesis.

Metropolis-Hastings algorithm for extracting periodic gravitational wave signals from laser interferometric detector data

Nelson Christensen,^{1,*} Réjean J. Dupuis,^{2,†} Graham Woan,^{2,‡} and Renate Meyer^{3,§}

¹*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057, USA*

²*Department of Physics and Astronomy, University of Glasgow, G12 8QQ, United Kingdom*

³*Department of Statistics, University of Auckland, Auckland, New Zealand*

(Received 8 January 2004; published 1 July 2004)

The Markov chain Monte Carlo methods offer practical procedures for detecting signals characterized by a large number of parameters and under conditions of low signal-to-noise ratio. We present a Metropolis-Hastings algorithm capable of inferring the spin and orientation parameters of a neutron star from its periodic gravitational wave signature seen by laser interferometric detectors.

DOI: 10.1103/PhysRevD.70.022001

PACS number(s): 04.80.Nn

I. INTRODUCTION

The worldwide network of laser interferometric gravitational wave detectors has begun to acquire scientifically significant data [1–4] and rapidly rotating neutron stars are an important potential source of signals (we will reserve the term “pulsar” to refer to the observed pulsating radio sources). Although a spinning spherically symmetric neutron star will not produce gravitational waves, a number of mechanisms have been proposed that are capable of producing quasi-periodic gravitational waves from biaxial or triaxial neutron stars [5,6]. Any gravitational waves from these neutron stars will likely be seen at Earth as weak continuous wave signals.

The data analysis task of identifying such a signal in the output of a laser interferometer is challenging and difficult, both because of the weakness of the signal and because its time evolution is characterized by a relatively large number of parameters. Radio observations can provide the sky location, rotation frequency and spindown rate of known pulsars, but the problem of looking for unknown (or poorly parameterized) neutron star sources is significantly more challenging. SN1987A is a good example of a poorly parameterized source for which the sky location is approximately known but also for which there is a large uncertainty in the frequency and spindown parameters of the putative neutron star [7].

Much work has already gone into all-sky hierarchical methods for searching for continuous gravitational waves [8,9]. Here we address the specific problem of a “fuzzy” parameter space search, in which a restricted volume of the space needs to be thoroughly investigated. We take a Bayesian approach to this problem and use Markov chain Monte Carlo (MCMC) techniques which have been shown to be especially suited to similar problems involving numerous parameters [10]. In particular, the Metropolis-Hastings (MH)

algorithm [11,12] has been used for estimating cosmological parameters from cosmic microwave background data [13–15], and the applicability of the MH routine has been demonstrated in estimating astrophysical parameters for gravitational wave signals from coalescing compact binary systems [16,17]. MCMC methods have also provided Bayesian inference for noisy and chaotic data [18,19].

Here we demonstrate that a MH algorithm also offers great promise for estimating neutron star parameters from their continuous gravitational wave signals. This work builds on the development (by two of us) of an end-to-end robust Bayesian method of searching for periodic signals in gravitational wave interferometer data [20], summarized in Sec. II. Starting with this Bayesian approach we apply a similar MH routine to that used in [13,17]. The description of the Bayesian MH method is given in Sec. III. In Sec. IV we present the results of this study, using synthesized data, for four and five parameter problems. We believe that this method offers great hope for signal extraction as more parameters are included, and this point is discussed in Sec. V.

II. SIGNAL CHARACTERISTICS

We will initially consider searching for signals from known radio pulsars, and then expand the method to account for an uncertainty in the frequency of the gravitational wave signal. As gravitational waves from pulsars are certainly weak at Earth, long integration periods are required to extract the signal, and we must take account of the antenna patterns of the detectors and the Doppler shift due to the motion of the Earth.

As in the previous study [20,21] we consider the signal expected from a non-precessing triaxial neutron star. The gravitational wave signal from such an object is at twice its rotation frequency, $f_s = 2f_r$, and we characterize the amplitudes of each polarization with overall strain factor, h_0 . The measured gravitational wave signal will also depend on the polarization antenna patterns of the detector $F_{\times,+}$ giving a signal

$$s(t) = \frac{1}{2} F_{+}(t; \psi) h_0 (1 + \cos^2 \iota) \cos \Psi(t) + F_{\times}(t; \psi) h_0 \cos \iota \sin \Psi(t), \quad (1)$$

*Electronic address: nchriste@carleton.edu

†Electronic address: rejean@astro.gla.ac.uk

‡Electronic address: graham@astro.gla.ac.uk

§Electronic address: meyer@stat.auckland.ac.nz

where ψ is the polarization angle of the radiation (which depends on the position angle of the spin axis in the plane of the sky) and ι is the inclination of the pulsar with respect to the line of sight.

Using a simple slowdown model, the phase evolution of the signal can be usefully parameterized as

$$\Psi(t) = \phi_0 + 2\pi \left[f_s(T - T_0) + \frac{1}{2}\dot{f}_s(T - T_0)^2 + \frac{1}{6}\ddot{f}_s(T - T_0)^3 \right], \quad (2)$$

where

$$T = t + \delta t = t + \frac{\mathbf{r} \cdot \mathbf{n}}{c} + \Delta_T. \quad (3)$$

Here, T is the time of arrival of the signal at the solar system barycenter, ϕ_0 is the phase of the signal at a fiducial time T_0 , \mathbf{r} is the position of the detector with regard to the solar system barycenter, \mathbf{n} is a unit vector in the direction of the pulsar, c is the speed of light, and Δ_T contains the relativistic corrections to the arrival time [22].

The signal is *heterodyned* by multiplying the data by $\exp[-i\Psi(t)]$ so that the only time varying quantity remaining is the antenna pattern of the interferometer (which varies over the day). For convenience, the result is low-pass filtered and resampled. We are left with a simple model with four unknown parameters: the overall amplitude of the gravitational wave signal (h_0), its polarization angle (ψ), its phase at time T_0 (ϕ_0), and the angle between the spin axis of the pulsar and the line of sight (ι).

A detailed description of the heterodyning procedure is presented elsewhere [20,21]; here we just provide a summary of this standard technique. The raw signal, $s(t)$, is centered near twice the rotation frequency of the pulsar, but is Doppler modulated due to the motion of the Earth and the orbit of the pulsar if it is in a binary system. The modulation bandwidth is typically 10^4 times less than the detector bandwidth, so one can greatly reduce the effective data rate by extracting this band and shifting it to zero frequency. In its standard form the result is one binned data point, B_k , every minute, containing all the relevant information from the original time series but at only 2×10^{-6} the original data rate. If the phase evolution has been correctly accounted for at this heterodyning stage then the only time-varying component left in the signal will be the effect of the antenna pattern of the interferometer, as its geometry with respect to the neutron star varies with Earth rotation. Any small error, Δf , in the heterodyne frequency will cause the signal to oscillate at Δf , and for the second part of our study we have Δf as our fifth parameter. For both these studies we estimate the noise variance, σ_k^2 , in the bin values, B_k , from the sample variance of the contributing data. It is assumed that the noise is stationary over the 60 s of data contributing to each bin.

III. THE METROPOLIS-HASTINGS ALGORITHM

This section presents a brief review of the Bayesian MH approach to parameter estimation. Comprehensive descriptions of MCMC methods and the MH algorithm can be found elsewhere [10,13,17].

We will denote the output from the above heterodyning procedure as $\{B_k\}$, with a joint probability distribution function (PDF) denoted by $p(\{B_k\}|\mathbf{a})$ conditional on unobserved parameters $\mathbf{a} = (a_1, \dots, a_d)$. The PDF $p(\{B_k\}|\mathbf{a})$ is referred to as the *likelihood* and regarded as a function of the parameters \mathbf{a} . The parameters of interest for our four parameter study are $\mathbf{a} = (h_0, \psi, \phi_0, \iota)$, while for the five parameter study they are $\mathbf{a} = (h_0, \psi, \phi_0, \iota, \Delta f)$.

From Eq. (1), the (now complex) heterodyned signal is

$$y(t_k; \mathbf{a}) = \frac{1}{4} F_+(t_k; \psi) h_0 (1 + \cos^2 \iota) e^{i\phi_0} - \frac{i}{2} F_\times(t_k; \psi) h_0 \cos \iota e^{i\phi_0}, \quad (4)$$

and the binning procedure should, by the central limit theorem, give the noise a near-Gaussian probability density characterized by a variance σ_k^2 for the k th bin. The likelihood that the data in this bin, taken at time t_k , is consistent with the above model is

$$p(B_k|\mathbf{a}) \propto \exp\left(-\frac{|B_k - y(t_k; \mathbf{a})|^2}{2\sigma_k^2}\right), \quad (5)$$

and the joint likelihood that the data in all the bins (taken as independent) are consistent with a particular set of model parameters is

$$p(\{B_k\}|\mathbf{a}) \propto \prod_k \exp\left(-\frac{|B_k - y(t_k; \mathbf{a})|^2}{2\sigma_k^2}\right). \quad (6)$$

Bayesian inference requires the specification of a prior PDF for \mathbf{a} , $p(\mathbf{a})$, that quantifies the researcher's pre-experimental knowledge about \mathbf{a} . The phase and polarization priors are flat in their space, and are set uniform for ϕ_0 over $[0, \pi]$, and for ψ over $[-\pi/4, \pi/4]$. The prior for ι is uniform in $\cos \iota$ over $[-1, 1]$, corresponding to a uniform prior per unit solid angle of pulsar orientation. Finally, in the present study we take a prior for h_0 that is uniform for $0 < h_0 < 1000$ (in our normalized units for which $\sigma_k = 1$), and zero for all other values.

Using Bayes' theorem, the post-experimental knowledge of \mathbf{a} is expressed by the *posterior* PDF of \mathbf{a} :

$$p(\mathbf{a}|\{B_k\}) = \frac{p(\mathbf{a})p(\{B_k\}|\mathbf{a})}{p(\{B_k\})} \propto p(\mathbf{a})p(\{B_k\}|\mathbf{a}), \quad (7)$$

where $p(\{B_k\}) = \int p(\{B_k\}|\mathbf{a})p(\mathbf{a})d\mathbf{a}$ is the marginal PDF of $\{B_k\}$ which can be regarded as a normalizing constant as it is independent of \mathbf{a} . The posterior PDF is thus proportional to the product of prior and likelihood.

The marginal posterior distribution for parameter a_i is the integral of the joint posterior PDF over all other components of \mathbf{a} other than a_i , i.e.,

$$p(a_i|\{B_k\}) = \int \dots \int p(\mathbf{a}|\{B_k\}) \times da_1 \dots da_{i-1} da_{i+1} \dots da_d, \quad (8)$$

and contains all the analysis has to say about the value of a_i alone. However it is often useful to summarize this in a single "point estimate" of a_i using, for example, the posterior mean:

$$\langle a_i \rangle = \int a_i p(a_i|\{B_k\}) da_i. \quad (9)$$

Calculating the normalization constant $p(\{B_k\})$ and calculating each marginal posterior PDF requires difficult d - and $(d-1)$ -dimensional integrations, respectively, that can be evaluated using a sampling approach and MCMC methods [10,13,17]. Rather than sampling directly from $p(\mathbf{a}|\{B_k\})$, a sample from a Markov chain is generated which has $p(\mathbf{a}|\{B_k\})$ as its equilibrium distribution. Thus, after running the Markov chain for a certain "burn-in" period, these (correlated) samples can be regarded as samples from the limiting distribution, provided that the Markov chain has reached convergence. Despite their correlations, the ergodic theorem guarantees that the sample average is still a consistent estimate of the posterior mean Eq. (9) [23].

The specific MCMC technique used for this study was the MH algorithm [11,12]. The MH algorithm generates a sample from the target PDF $p(\mathbf{a}|\{B_k\})$ using a technique that is similar to the well-known simulation technique of *rejection sampling*. A candidate is generated from an auxiliary PDF and then accepted or rejected with some probability. Starting with an arbitrary initial state \mathbf{a}_0 , at time n a new candidate \mathbf{a}' is generated from the candidate generating PDF, $q(\mathbf{a}|\mathbf{a}_n)$, which can depend on the current state \mathbf{a}_n of the Markov chain. This new candidate \mathbf{a}' is accepted with a certain *acceptance probability* $\alpha(\mathbf{a}'|\mathbf{a}_n)$, also depending on the current state \mathbf{a}_n , given by

$$\alpha(\mathbf{a}'|\mathbf{a}_n) = \min \left\{ \frac{p(\mathbf{a}')p(\{B_k\}|\mathbf{a}')q(\mathbf{a}_n|\mathbf{a}')}{p(\mathbf{a}_n)p(\{B_k\}|\mathbf{a}_n)q(\mathbf{a}'|\mathbf{a}_n)}, 1 \right\}. \quad (10)$$

For good efficiency a multivariate normal distribution centered at the current state \mathbf{a}_n is used for $q(\mathbf{a}'|\mathbf{a}_n)$. This then implies that if the posterior probability at \mathbf{a}' is larger than at the current state \mathbf{a}_n , the proposed step to \mathbf{a}' is always accepted. However, if the step is in a direction of lower posterior probability, then this step is accepted only with a certain probability given by the ratio of the posterior PDFs (since our multivariate normal generating function is symmetric in \mathbf{a}' and \mathbf{a}_n and therefore cancels out). If the candidate is accepted, the next state of the Markov chain is $\mathbf{a}_{n+1} = \mathbf{a}'$, otherwise the chain does not move, i.e. $\mathbf{a}_{n+1} = \mathbf{a}_n$.

The steps of the MH algorithm are therefore:

Step 0: Start with an arbitrary value \mathbf{a}_0 ;

Step $n+1$: Generate \mathbf{a}' from $q(\mathbf{a}|\mathbf{a}_n)$ and u from $U(0,1)$;

If $u \leq \alpha(\mathbf{a}'|\mathbf{a}_n)$ set $\mathbf{a}_{n+1} = \mathbf{a}'$ (acceptance),

If $u > \alpha(\mathbf{a}'|\mathbf{a}_n)$ set $\mathbf{a}_{n+1} = \mathbf{a}_n$ (rejection).

$U(0,1)$ is the uniform distribution between 0 and 1. The efficiency of the MH algorithm depends heavily on the choice of the proposal density. The closer the proposal is to the target distribution, the faster convergence will be accomplished. This link between the closeness of the proposal to stationary distribution and speed of convergence has also been substantiated by Holden [24]. In the study presented here we dynamically altered the proposal distribution based on information from the chain's history. The approach, called *pilot adaptation*, is to perform a separate pilot run to gain insight about the target density and then tune the proposal accordingly for the successive runs. Such adaptation can be iterated but allowing it infinitely often will destroy the Markovian property of the chain and thereby often compromise the stationarity of the chain and the consistency of sample path averages ([25]; see [26] for an example).

Based on the central limit theorem, the posterior PDF should be well approximated by a multivariate normal distribution with mean equal to the posterior mode and covariance matrix equal to minus the Hessian evaluated at the posterior mode. Thus, we use a multivariate normal distribution for the proposal density $q(\mathbf{a}|\mathbf{a}_n)$. As the mode is unknown, we try to make use of pilot samples to estimate its covariance matrix. When we initially run the MH algorithm, we sample candidate parameters from a normal distribution with covariance matrix equal to the identity matrix and centered around the current state. After the completion of this pilot run we use the empirical covariance matrix of the sample as covariance matrix of the multivariate normal proposal density, again with mean equal to the current state.

IV. RESULTS

In the first part of our study we reproduced the results presented in [20] where the four unknown parameters were h_0, ι, ψ , and ϕ_0 . The signal $s(t)$ was synthesized assuming a source at right ascension = $4^{\text{h}}41^{\text{m}}54^{\text{s}}$ and declination = $18^{\circ}23'32''$, as would be seen by the LIGO-Livingston interferometer. This was then added to white Gaussian noise, $n(t)$, which is a good approximation to the detector noise in our band. Our normalized data had a noise variance of $\sigma_k^2 = 1$ for each sample, and the amplitude of the signal used in our test runs was varied in the range $h_0 = 0.0$ to 10.0 . We were able to detect signals for $h_0 > 0.1$. The length of the data set corresponded to 14 400 samples or 10 days of data at a rate of one sample per minute (which was the rate used for the LIGO/GEO S1 analysis described in [21]). Although we will work with strains normalized to $\sigma_k = 1$, the results can be cast into a more conventional form by multiplying σ_i and h_0 by $(S_h/60)^{1/2}$, where $(S_h)^{1/2}$ is the strain noise spectral density of the detector at the frequency of interest, in Hz^{-1} .

An example of the MH routine output is shown in Fig. 1. Displayed are the trace plots and the kernel densities (posterior PDFs). For this example the program ran for 10^6 iterations. The first 10^5 iterations were discarded as the burn-in. Short-term correlations in the chain were eliminated by

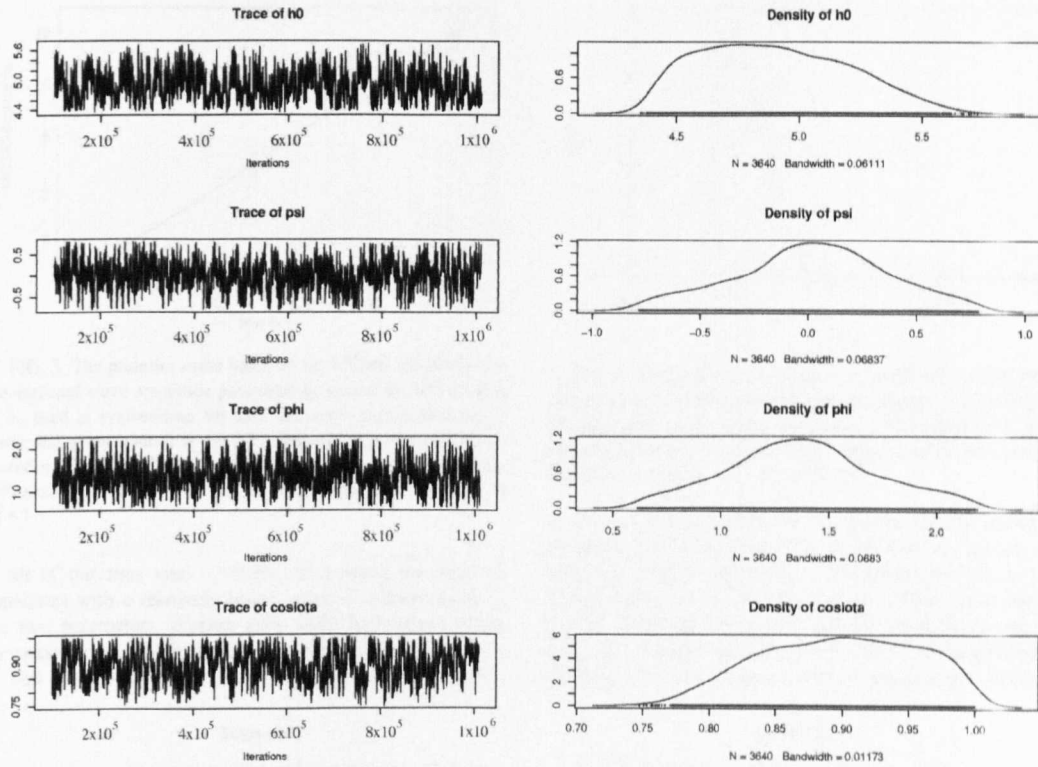


FIG. 1. Trace plot (left) and MCMC-estimated posterior PDFs (right) for the pulsar parameters h_0 , ψ , ϕ_0 and $\cos i$. In this example the true parameters were $h_0=5.0$, $\phi_0=1.0$, $\psi=0.4$, and $i=0.5$ implying $\cos i=0.88$.

“thinning” the remaining terms; we kept every 250th item in the chain. The true parameter values for this run were $h_0=5.0$, $\psi=0.4$, $\phi_0=1.0$ and $i=0.5$ ($\cos i=0.88$). In the example displayed in Fig. 1 the MCMC yielded mean values and 95% posterior probability intervals of $h_0=4.9$ (4.43 to 5.50), $\psi=0.02$ (−0.68 to 0.69), $\phi_0=1.34$ (0.71 to 2.08), and $\cos i=0.90$ (0.79 to 0.99). The 95% posterior probability interval is specified by the 2.5% and 97.5% percentiles of $p(a_i|B_k)$. In Fig. 2 we display the estimated posterior PDF of h_0 on an expanded scale, along with the real and estimated value for h_0 .

It is crucial that our algorithm is sensitive to the true value of the gravitational wave amplitude, h_0 , even under conditions of relatively low signal-to-noise ratio, and Fig. 3 shows injected h_0 values versus their values inferred by the MH routine. The error bars correspond to the 95% posterior probability interval, i.e. the lower and upper bound are specified by the 2.5% and 97.5% percentiles of $p(a_i|B_k)$. The algorithm clearly is successful in finding and estimating h_0 . While the error bars increase as the signal gets larger, the relative error $\Delta h_0/h_0$ does diminish as h_0 increases. The fact that the 95% posterior probability interval grows with h_0 for constant noise level would seem to be counterintuitive. In addition, the widths of the posterior probability distributions for h_0 are larger than would be naively expected from a search for a simple periodic signal. The reason is that these

error bars represent the uncertainty in the parameter rather than just the level of the noise, and this is affected both by the noise level and the posterior covariance between all of the parameters. The MCMC technique also allows one to calculate cross-correlation coefficients from the Markov chains of the parameters, and the value between h_0 and $\cos i$




FIG. 2. An expanded view of the estimated posterior PDF based on the MCMC sample for parameter h_0 . The vertical solid line shows the posterior mean of $h_0=4.9$, while the vertical dotted line marks the true parameter value of $h_0=5.0$.




FIG. 3. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 versus the actual value of h_0 used in synthesizing the data. The error bars correspond to lower and upper bounds at the 2.5% and 97.5% percentiles of the posterior PDF. The solid line has a slope of 1. The calculations were performed over 14 000 data points, each with noise variance of $\sigma_k^2 = 1$.

in all of our runs was ~ -0.95 . As a result the data are consistent with a relatively broad range of combinations of the two parameters, making their *individual* values rather uncertain here—an effect evident from Eq. (1).

The effect of the other unknown parameters (particularly




FIG. 4. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 (dotted line), along with that produced via the method presented in [20] (solid line). In this example the *true* value was $h_0 = 0.5$, while the other *true* parameter values were $\psi = 0.4$, $\phi = 1.0$, and $\epsilon = 0.5$.

ϵ) on the posterior PDF for h_0 can be clearly shown by repeating the analysis for Fig. 3 but with h_0 as the only unknown, namely, all of the other parameters set to their *actual* values in the MCMC routine. Under these circumstances the widths of all 95% posterior probability intervals are 0.116, independent of the value of h_0 . Comprehensive analyses have investigated detection statistics for a periodic

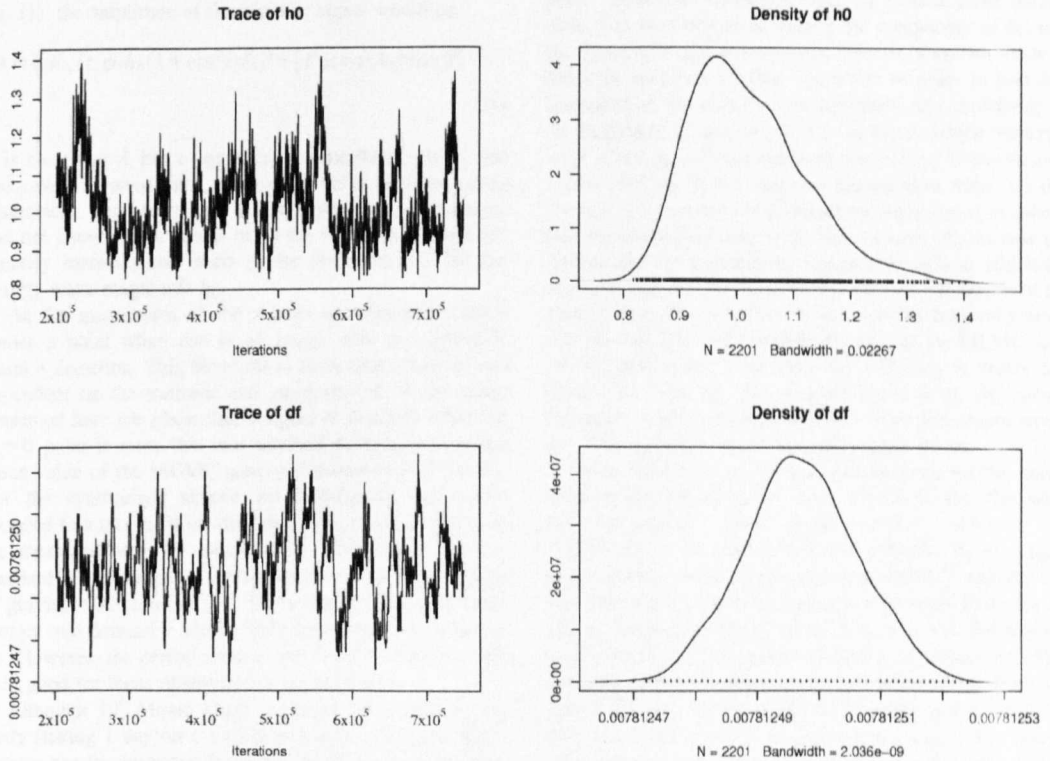


FIG. 5. Trace plot (left) and posterior PDFs (right) for the pulsar parameters h_0 and Δf . In this example from the five parameter problem the true values for these critical parameters were $h_0 = 1.0$ and $\Delta f = 0.0078125$ Hz.




FIG. 6. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 versus the actual value of h_0 used in synthesizing the data. This example is from the five parameter problem. The error bars correspond to the lower and upper bounds being specified by the 2.5% and 97.5% percentiles of the posterior PDF. The solid line has a slope of 1.

signal in a gravity wave detector [27]. However, these statistics are concerned only with the amplitude of the periodic signal, and not with parameter estimation (as described above). If we write Eq. (1) as $s(t) = A \cos(\Psi + \Phi)$ (with A being the periodic signal amplitude and Φ a phase term) then the detection statistic of [27] would apply to finding a signal amplitude A in the presence of the detector noise. In terms of Eq. (1), the amplitude of the periodic signal would be

$$A = \{[F_+(t; \psi)h_0(1 + \cos^2 \iota)/2]^2 + [F_\times(t; \psi)h_0 \cos \iota]^2\}^{1/2}. \quad (11)$$

It is clear that A has a complicated dependence on h_0 and $\cos \iota$. We will never know, *a priori*, the value of all the pulsar parameters. Our study here is about parameter estimation, and not knowing the values of all the pulsar parameters ultimately increases the width in the posterior PDF for the gravity wave magnitude h_0 .

As the magnitudes of the signals are diminished there comes a point when one is no longer able to *confidently* claim a detection. This threshold is somewhat arbitrary, and dependent on the statistics and interpretation. In the study presented here we *claim* that a signal is detected when the $h_0=0$ point is more than two standard deviations from the mean value of the MCMC generated posterior PDF for h_0 . For the synthesized signals we investigated this corresponded to a threshold for detection of $h_0=0.1$; in this case the measured mean of the posterior PDF for h_0 was 2.1 standard deviations away from zero. For an initial detection of gravitational radiation it is likely that the scientific community will demand a significantly larger signal-to-noise ratio. However, the performance of the MCMC routine is still very good for these relatively low signal levels.

Although 10^6 Monte Carlo iterations were used in this study (taking 1 day on a 1 GHz processor) adequate distributions can be generated from 10^5 iterations after the burn-in, so good results can be achieved after just a few hours. In fact the marginalizations discussed above can be tackled




FIG. 7. The posterior mean based on the MCMC sample for the uncertainty in the frequency, Δf , versus the actual value of h_0 used in synthesizing the data. This example is from the five parameter problem. The error bars correspond to the lower and upper bounds being specified by the 2.5% and 97.5% percentiles of the posterior PDF. The horizontal line corresponds to the real value of $\Delta f = 0.007\,812\,500$ Hz.

more quickly using simple summing methods as performed by [20], and the result of a comparison of the two is shown in Fig. 4. The great advantage of the MCMC method for us is its demonstrated ability to deal with problems that have a large number of parameters [10], where other numerical integration techniques (such as employed by [20]) are not feasible. The ultimate goal of our research is to expand this pulsar parameter estimation work to include more parameters. The next step in increasing the complexity of the pulsar signal is to consider potential sources of known location, but with unknown rotation frequency. In order to start this investigation we added a new parameter, the uncertainty in the frequency of the source, Δf . In this example the exact value of the pulsar's gravitational wave signal is uncertain to within 1/60 Hz. In the study we present here there is a difference, Δf , between the gravitational wave signal frequency and the heterodyne frequency. The addition of this new parameter did not significantly increase the rate at which the code ran, but did (by about 20%) increase the length of the burn-in time. If one wanted to increase this frequency range to 5 Hz then this could be done by running the MCMC code on 300 processors, with each run differing in center frequency by 1/60 Hz. The Markov chain using the *correct* frequency would converge, while the other 299 chains would not. This will be a future research project for us.

In our MH code we used a uniform prior for the uncertainty in the frequency, Δf , over $\pm 0.016\,67$ Hz. The injection parameters used were $\psi=0.4$, $\phi=1.0$, $\Delta f=0.007\,812\,5$ Hz, and $\iota=0.5$ ($\cos \iota=0.88$). h_0 was again injected with a number of values between 0.25 and 10.0. In Fig. 5 we show sample trace plots and posterior PDFs for Δf and h_0 when the injected value of h_0 was 1.0. For this example the MCMC algorithm yielded mean values and 95% posterior probability intervals of $h_0=1.02$ (0.86 to 1.26) and $\Delta f=0.007\,812\,497$ Hz (0.007 812 480 Hz to 0.007 812 515 Hz). The frequency PDF is quite narrow, which was responsible for the increase in the burn-in time as the Markov chain must find this narrow region of parameter space. In Fig. 6 we display the estimate for the gravitational wave amplitude

(h_0) predicted by the five parameter MH routine versus the actual h_0 . In Fig. 7 we display the estimate for the difference in frequency Δf predicted by the five parameter MH routine versus the injected h_0 .

V. DISCUSSION

Recent applications of MCMC techniques have provided a Bayesian approach to estimating parameters in a number of physical situations. These include cosmological parameter estimation from cosmic microwave background data [13–15], estimating astrophysical parameters for gravitational wave signals from coalescing compact binary systems [16,17], and parameter estimation of a chaotic system in the presence of noise [18,19]. An all sky survey for periodic gravitational waves from neutron stars must explore a very large parameter space, and this has partially been addressed in [8]. Generically, the signal from a neutron star in a binary

system will be characterized by at least 13 parameters. It is our hope that the MCMC techniques will prove fruitful in dealing with these complex signals.

In this paper we have demonstrated that the success of the MH routine for the five parameter problem: h_0 , ψ , ϕ_0 , t and Δf . Our longer term plans are to account for other parameters, such as spindown rate, pulsar wobble, and possibly location of the signal in the sky. This research is currently in progress.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation grants PHY-0071327 and PHY-0244357, the Royal Society of New Zealand Marsden fund award UOA204, the Natural Sciences and Engineering Research Council of Canada, Universities U.K., and the University of Glasgow.

-
- [1] B. Barish and R. Weiss, *Phys. Today* **52** (10), 44 (1999).
 - [2] B. Willke *et al.*, *Class. Quantum Grav.* **19**, 1377 (2002).
 - [3] B. Caron *et al.*, *Nucl. Phys. B (Proc. Suppl.)* **54**, 167 (1996).
 - [4] K. Tsubono, in *1st Edoardo Amaldi Conference on Gravitational Wave Experiments*, edited by E. Coccia, G. Pizella, and F. Ronga (World Scientific, Singapore, 1995), p. 112.
 - [5] C. Cutler, *Phys. Rev. D* **66**, 084025 (2002).
 - [6] L. Bildsten, *Astrophys. J. Lett.* **501**, L89 (1998).
 - [7] J. Middleditch, J.A. Kristan, W.E. Kunkel, K.M. Hill, R.D. Watson, R. Lucinio, J.N. Imamura, T.Y. Steiman-Cameron, A.S. Shearer, R. Butler, M. Redfern, and A.C. Danks, *New Astron.* **5**, 243 (2000).
 - [8] P. Jaranowski, A. Królak, and B.F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
 - [9] P. Brady and T. Creighton, *Phys. Rev. D* **61**, 082001 (2000).
 - [10] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Chapman and Hall, London, 1996).
 - [11] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
 - [12] W.K. Hastings, *Biometrika* **57**, 97 (1970).
 - [13] N. Christensen, R. Meyer, L. Knox, and B. Luey, *Class. Quantum Grav.* **18**, 2677 (2001).
 - [14] L. Knox, N. Christensen, and C. Skordis, *Astrophys. J. Lett.* **563**, L95 (2001).
 - [15] L. Verde, H.V. Peiris, D.N. Spergel, M. Nolte, C.L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S.S. Meyer, L. Page, G.S. Tucker, E. Wollack, and E.L. Wright, *Astrophys. J., Suppl. Ser.* **148**, 195 (2003).
 - [16] N. Christensen and R. Meyer, *Phys. Rev. D* **64**, 022001 (2001).
 - [17] N. Christensen, R. Meyer, and A. Libson, *Class. Quantum Grav.* **21**, 317 (2004).
 - [18] R. Meyer and N. Christensen, *Phys. Rev. E* **62**, 3535 (2000).
 - [19] R. Meyer and N. Christensen, *Phys. Rev. E* **65**, 016216 (2002).
 - [20] Réjean J. Dupuis and Graham Woan, "A Bayesian Method to Search for Periodic Gravitational Waves."
 - [21] B. Abbott *et al.*, *Phys. Rev. D* **69**, 082004 (2004).
 - [22] J.H. Taylor, *Phys. Rev. D* **66**, 084025 (2002).
 - [23] L. Tierney, *Ann. Stat.* **22**, 1701 (1994).
 - [24] L. Holden, "Adaptive chains," <http://www.statslab.cam.ac.uk/mcmc>
 - [25] W.R. Gilks, G.O. Roberts, and S.K. Sahu, *J. Am. Stat. Assoc.* **93**, 1045 (1998).
 - [26] A.E. Gelfand and S.K. Sahu, *J. Comput. Graph. Stat.* **3**, 261 (1994).
 - [27] B. Allen, M.A. Papa, and B.F. Schutz, *Phys. Rev. D* **66**, 102003 (2002).

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

B. Abbott,¹³ R. Abbott,¹⁶ R. Adhikari,¹⁴ A. Ageev,^{21,28} B. Allen,⁴⁰ R. Amin,³⁵ S. B. Anderson,¹³ W. G. Anderson,³⁰ M. Araya,¹³ H. Armandula,¹³ F. Asiri,^{13,a} P. Aufmuth,³² C. Aubert,¹ S. Babak,⁷ R. Balasubramanian,⁷ S. Ballmer,¹⁴ B. C. Barish,¹³ D. Barker,¹⁵ C. Barker-Patton,¹⁵ M. Barnes,¹³ B. Barr,³⁶ M. A. Barton,¹³ K. Bayer,¹⁴ R. Beausoleil,^{27,b} K. Belczynski,²⁴ R. Bennett,^{36,c} S. J. Berukoff,^{1,d} J. Betzwieser,¹⁴ B. Bhawal,¹³ I. A. Bilenko,²¹ G. Billingsley,¹³ E. Black,¹³ K. Blackburn,¹³ B. Bland-Weaver,¹⁵ B. Bochner,^{14,e} L. Bogue,¹³ R. Bork,¹³ S. Bose,⁴¹ P. R. Brady,⁴⁰ V. B. Braginsky,²¹ J. E. Brau,³⁸ D. A. Brown,⁴⁰ S. Brozek,^{32,f} A. Bullington,²⁷ A. Buonanno,^{6,g} R. Burgess,¹⁴ D. Busby,¹³ W. E. Butler,³⁹ R. L. Byer,²⁷ L. Cadonati,¹⁴ G. Cagnoli,³⁶ J. B. Camp,²² C. A. Cantley,³⁶ L. Cardenas,¹³ K. Carter,¹⁶ M. M. Casey,³⁶ J. Castiglione,³⁵ A. Chandler,¹³ J. Chapsky,^{13,h} P. Charlton,¹³ S. Chatterji,¹⁴ Y. Chen,⁶ V. Chikarmane,¹⁷ D. Chin,³⁷ N. Christensen,⁸ D. Churches,⁷ C. Colacino,^{32,2} R. Coldwell,³⁵ M. Coles,^{16,i} D. Cook,¹⁵ T. Corbitt,¹⁴ D. Coyne,¹³ J. D. E. Creighton,⁴⁰ T. D. Creighton,¹³ D. R. M. Crooks,³⁶ P. Csatorday,¹⁴ B. J. Cusack,³ C. Cutler,¹ E. D'Ambrosio,¹³ K. Danzmann,¹³ R. Davies,⁷ E. Daw,^{17,j} D. DeBra,²⁷ T. Delker,¹³ R. DeSalvo,¹³ S. Dhurandhar,¹² M. Díaz,³⁰ H. Ding,¹³ R. W. P. Drever,⁴ R. J. Dupuis,³⁶ C. Ebeling,⁸ J. Edlund,¹³ P. Ehrens,¹³ E. J. Elliffe,³⁶ T. Etzel,¹³ M. Evans,¹³ T. Evans,¹⁶ C. Fallnich,³² D. Farnham,¹³ M. M. Fejer,²⁷ M. Fine,¹³ L. S. Finn,²⁹ E. Flanagan,⁹ A. Freise,²¹ R. Frey,³⁸ P. Fritschel,¹⁴ V. Frolov,¹⁶ M. Fyffe,¹⁶ K. S. Ganezer,⁵ J. A. Giaime,¹⁷ A. Gillespie,^{13,m} K. Goda,¹⁴ G. González,¹⁷ S. Goßler,³² P. Grandclément,²⁴ A. Grant,³⁶ C. Gray,¹⁵ A. M. Gretarsson,¹⁶ D. Grimmer,¹³ H. Grote,² S. Grunewald,¹ M. Guenther,¹⁵ E. Gustafson,^{27,n} R. Gustafson,³⁷ W. O. Hamilton,¹⁷ M. Hammond,¹⁶ J. Hanson,¹⁶ C. Hardham,²⁷ G. Harry,¹⁴ A. Hartunian,¹³ J. Heefner,¹³ Y. Hefetz,¹⁴ G. Heinzl,² I. S. Heng,³² M. Hennessy,²⁷ N. Hepler,²⁹ A. Heptonstall,³⁶ M. Heurs,³² M. Hewitson,³⁶ N. Hindman,¹⁵ P. Hoang,¹³ J. Hough,³⁶ M. Hrynevych,^{13,o} W. Hua,²⁷ R. Ingley,³⁴ M. Ito,³⁸ Y. Itoh,¹ A. Ivanov,¹³ O. Jennrich,^{36,p} W. W. Johnson,¹⁷ W. Johnston,³⁰ L. Jones,¹³ D. Jungwirth,^{13,q} V. Kalogera,²⁴ E. Katsavounidis,¹⁴ K. Kawabe,^{20,2} S. Kawamura,²³ W. Kells,¹³ J. Kern,¹⁶ A. Khan,¹⁶ S. Killbourn,³⁶ C. J. Killow,³⁶ C. Kim,²⁴ C. King,¹³ P. King,¹³ S. Klimenko,³⁵ P. Kloevekom,² S. Koranda,⁴⁰ K. Kötter,³² J. Kovalik,¹⁶ D. Kozak,¹³ B. Krishnan,¹ M. Landry,¹⁵ J. Langdale,¹⁶ B. Lantz,²⁷ R. Lawrence,¹⁴ A. Lazzarini,¹³ M. Lei,¹³ V. Leonhardt,³² I. Leonor,³⁸ K. Libbrecht,¹³ P. Lindquist,¹³ S. Liu,¹³ J. Logan,^{13,r} M. Lormand,¹⁶ M. Lubinski,¹⁵ H. Lück,^{32,2} T. T. Lyons,^{13,s} B. Machenschalk,¹ M. MacInnis,¹⁴ M. Mageswaran,¹³ K. Mailand,¹³ W. Majid,^{13,h} M. Malec,³² F. Mann,¹³ A. Marin,^{14,s} S. Márka,¹³ E. Maros,¹³ J. Mason,^{13,t} K. Mason,¹⁴ O. Matherny,¹⁵ L. Matone,¹⁵ N. Mavalvala,¹⁴ R. McCarthy,¹⁵ D. E. McClelland,³ M. McHugh,¹⁹ P. McNamara,^{36,u} G. Mendell,¹⁵ S. Meshkov,¹³ C. Messenger,³⁴ V. P. Mitrofanov,²¹ G. Mitselmakher,³⁵ R. Mittleman,¹⁴ O. Miyakawa,¹³ S. Miyoki,^{13,v} S. Mohanty,^{1,w} G. Moreno,¹⁵ K. Mossavi,² B. Mours,^{13,x} G. Mueller,³⁵ S. Mukherjee,^{1,w} J. Myers,¹⁵ S. Nagano,² T. Nash,^{10,y} H. Naundorf,¹ R. Nayak,¹² G. Newton,³⁶ F. Nocera,¹³ P. Nutzman,²⁴ T. Olson,²⁵ B. O'Reilly,¹⁶ D. J. Ottaway,¹⁴ A. Ottewill,^{40,z} D. Ouimette,^{13,q} H. Overmier,¹⁶ B. J. Owen,²⁹ M. A. Papa,¹ C. Parameswariah,¹⁶ V. Parameswariah,¹⁵ M. Pedraza,¹³ S. Penn,¹¹ M. Pitkin,³⁶ M. Plissi,³⁶ M. Pratt,¹⁴ V. Quetschke,³² F. Raab,¹⁵ H. Radkins,¹⁵ R. Rakhola,³⁸ M. Rakhmanov,³⁵ S. R. Rao,¹³ D. Redding,^{13,h} M. W. Regehr,^{13,h} T. Regimbau,¹⁴ K. T. Reilly,¹³ K. Reithmaier,¹³ D. H. Reitze,³⁵ S. Richman,^{14,aa} R. Riesen,¹⁶ K. Riles,³⁷ A. Rizzi,^{16,bb} D. I. Robertson,³⁶ N. A. Robertson,^{36,27} L. Robison,¹³ S. Roddy,¹⁶ J. Rollins,¹⁴ J. D. Romano,^{30,cc} J. Romie,¹³ H. Rong,^{35,m} D. Rose,¹³ E. Rothhoff,²⁹ S. Rowan,³⁶ A. Rüdiger,^{20,2} P. Russell,¹³ K. Ryan,¹⁵ I. Salzman,¹³ G. H. Sanders,¹³ V. Sannibale,¹³ B. Sathyaprakash,⁷ P. R. Saulson,²⁸ R. Savage,¹⁵ A. Sazonov,³⁵ R. Schilling,^{20,2} K. Schlaufman,²⁹ V. Schmidt,^{13,dd} R. Schofield,³⁸ M. Schrempel,^{32,ee} B. F. Schutz,^{1,7} P. Schwinberg,¹⁵ S. M. Scott,³ A. C. Searle,³ B. Sears,¹³ S. Seel,¹³ A. S. Sengupta,¹² C. A. Shapiro,^{29,ff} P. Shawhan,¹³ D. H. Shoemaker,¹⁴ Q. Z. Shu,^{35,gg} A. Sibley,¹⁶ G. Siemens,⁴⁰ L. Sievers,^{13,h} D. Sigg,¹⁵ A. M. Sintes,^{1,33} K. Skeldon,³⁶ J. R. Smith,² M. Smith,¹⁴ M. R. Smith,¹³ P. Sneddon,³⁶ R. Spero,^{13,h} G. Stapfer,¹⁶ K. A. Strain,¹⁶ D. Strom,³⁸ A. Stuver,²⁹ T. Summerscales,²⁹ M. C. Sumner,¹³ P. J. Sutton,^{29,y} J. Sylvestre,¹³ A. Takamori,¹³ D. B. Tanner,³⁵ H. Tariq,¹³ I. Taylor,⁷ R. Taylor,¹³ K. S. Thorne,⁶ M. Tibbits,²⁹ S. Tilav,^{13,hh} M. Tinto,^{4,h} K. V. Tokmakov,²¹ C. Torres,³⁰ C. Torrie,^{13,36} S. Traeger,^{32,ii} G. Traylor,¹⁶ W. Tyler,¹³ D. Ugolini,³¹ M. Vallisneri,^{6,jj} M. van Putten,¹⁴ S. Vass,¹³ A. Vecchio,³⁴ C. Vorvick,¹⁵ S. P. Vyachanin,²¹ L. Wallace,¹³ H. Walter,²⁰ H. Ward,³⁶ B. Ware,^{13,h} K. Watts,¹⁶ D. Webber,¹³ A. Weidner,^{20,2} U. Weiland,³² A. Weinstein,¹³ R. Weiss,¹⁴ H. Welling,³² L. Wen,¹³ S. Wen,¹⁷ J. T. Whelan,¹⁹ S. E. Whitcomb,¹³ B. F. Whiting,³⁵ P. A. Willems,¹³ P. R. Williams,^{1,kk} R. Williams,⁴ B. Willke,^{32,2} A. Wilson,¹³ B. J. Winjum,^{29,d} W. Winkler,^{20,2} S. Wise,³⁵ A. G. Wiseman,⁴⁰ G. Woan,³⁶ R. Wooley,¹⁶ J. Worden,¹⁵ I. Yakushin,¹⁶ H. Yamamoto,¹³ S. Yoshida,²⁶ I. Zawischa,^{32,ll} L. Zhang,¹³ N. Zotov,¹⁸ M. Zucker,¹⁶ and J. Zweizig¹³

(LIGO Scientific Collaboration)^{mm}

¹Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Golm, Germany

²Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany

³Australian National University, Canberra, 0200, Australia

⁴California Institute of Technology, Pasadena, California 91125, USA

⁵California State University Dominguez Hills, Carson, California 90747, USA

⁶Caltech-CaRT, Pasadena, California 91125, USA

- ⁷Cardiff University, Cardiff, CF2 3YB, United Kingdom
⁸Carleton College, Northfield, Minnesota 55057, USA
⁹Cornell University, Ithaca, New York 14853, USA
¹⁰Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA
¹¹Hobart and William Smith Colleges, Geneva, New York 14456, USA
¹²Inter-University Centre for Astronomy and Astrophysics, Pune, 411007, India
¹³LIGO, California Institute of Technology, Pasadena, California 91125, USA
¹⁴LIGO-Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
¹⁵LIGO Hanford Observatory, Richland, Washington 99352, USA
¹⁶LIGO Livingston Observatory, Livingston, Louisiana 70754, USA
¹⁷Louisiana State University, Baton Rouge, Louisiana 70803, USA
¹⁸Louisiana Tech University, Ruston, Louisiana 71272, USA
¹⁹Loyola University, New Orleans, Louisiana 70118, USA
²⁰Max Planck Institut für Quantenoptik, D-85748, Garching, Germany
²¹Moscow State University, Moscow, 119992, Russia
²²NASA/Goddard Space Flight Center, Greenbelt, Maryland 20771, USA
²³National Astronomical Observatory of Japan, Tokyo 181-8588, Japan
²⁴Northwestern University, Evanston, Illinois 60208, USA
²⁵Salish Kootenai College, Pablo, Montana 59855, USA
²⁶Southeastern Louisiana University, Hammond, Louisiana 70402, USA
²⁷Stanford University, Stanford, California 94305, USA
²⁸Syracuse University, Syracuse, New York 13244, USA
²⁹The Pennsylvania State University, University Park, Pennsylvania 16802, USA
³⁰The University of Texas at Brownsville and Texas Southmost College, Brownsville, Texas 78520, USA
³¹Trinity University, San Antonio, Texas 78212, USA
³²Universität Hannover, D-30167 Hannover, Germany
³³Universitat de les Illes Balears, E-07071 Palma de Mallorca, Spain
³⁴University of Birmingham, Birmingham, B15 2TT, United Kingdom
³⁵University of Florida, Gainesville, Florida 32611, USA
³⁶University of Glasgow, Glasgow, G12 8QQ, United Kingdom
³⁷University of Michigan, Ann Arbor, Michigan 48109, USA
³⁸University of Oregon, Eugene, Oregon 97403, USA
³⁹University of Rochester, Rochester, New York 14627, USA
⁴⁰University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA
⁴¹Washington State University, Pullman, Washington 99164, USA

(Received 17 September 2003; published 30 April 2004)

Data collected by the GEO 600 and LIGO interferometric gravitational wave detectors during their first observational science run were searched for continuous gravitational waves from the pulsar J1939+2134 at twice its rotation frequency. Two independent analysis methods were used and are demonstrated in this paper: a frequency domain method and a time domain method. Both achieve consistent null results, placing new upper limits on the strength of the pulsar's gravitational wave emission. A model emission mechanism is used to interpret the limits as a constraint on the pulsar's equatorial ellipticity.

DOI: 10.1103/PhysRevD.69.082004

PACS number(s): 04.80.Nn, 07.05.Kf, 95.55.Ym, 97.60.Gb

I. INTRODUCTION

This work presents methods to search for periodic gravitational waves generated by known pulsars, using data collected by interferometric gravitational wave detectors. To illustrate these methods, upper limits are placed on the strength of waves emitted by pulsar J1939+2134 at its expected 1284 Hz emission frequency during S1 [1]. S1 is the first observational science run of the Laser Interferometer Gravitational Wave Observatory (LIGO) [2,3] and GEO [4,5] detectors and it took place during 17 days between 23 August and 9 September 2002. The sensitivity of the searches presented here surpasses that of previous searches for gravitational waves from this source. However, measurements of

the spin-down rate of the pulsar indicate that a detectable signal is very unlikely given the instrument performance for this data set: for these early observations the detectors were not operating at their eventual design sensitivities. Substantial improvements in detector noise have been achieved since the S1 observations, and further improvements are planned. We expect that the methods presented here will eventually enable the direct detection of periodic gravitational waves.

In Sec. II, we describe the configuration and calibration of the four LIGO and GEO interferometers and derive their expected sensitivities to periodic sources having known locations, frequencies, and spin-down rates. In Sec. III we consider proposed neutron star gravitational wave emission mechanisms and introduce notation for describing the nearly

monochromatic signals emitted by isolated neutron stars. Statistical properties of the data, analysis methods, and results are presented in Sec. IV. These results are then summarized and compared in Sec. V. In Sec. V we also interpret the upper limits on the signal amplitude as a constraint on the ellipticity of the pulsar and consider our results in the context of previous upper limits.

II. DETECTORS

Gravitational waves are a fundamental consequence of Einstein's general theory of relativity [6,7], in which they represent perturbations of the spacetime metric which propagate at the speed of light. Gravitational waves produced by

the acceleration of compact astrophysical objects may be detected by monitoring the motions they induce on freely falling test bodies. The strength of these waves, called the *strain*, can be characterized by the fractional variation in the geodesic separation between these test bodies.

During the past decade, several scientific collaborations have constructed a new type of detector for gravitational waves. These large-scale interferometric detectors include the U.S. Laser Interferometer Gravitational Wave Observatory (LIGO), located in Hanford, WA, and Livingston, LA, built by a Caltech-MIT collaboration [2,3]; the GEO 600 detector near Hannover, Germany, built by a British-German collaboration [4,5]; the VIRGO detector in Pisa, Italy, built by a French-Italian collaboration [8]; and the Japanese TAMA 300 detector in Tokyo [9]. In these detectors, the relative positions of suspended test masses are sensed interferometrically. A gravitational wave produces a time-varying differential displacement $\Delta L(t)$ in an interferometer that is proportional to its arm length L . The amplitude of the gravitational wave is described by the dimensionless strain $h(t) = \Delta L(t)/L$. For realistic periodic astrophysical sources we typically expect strain amplitudes smaller than 10^{-24} .

The following sections introduce the operating configurations of GEO 600 and LIGO detectors during the S1 run. The references provide more detailed descriptions of these detectors.

A. Instrument configurations

The GEO 600 detector comprises a four-beam Michelson delay line system of arm length 600 m. The interferometer is illuminated by frequency-stabilized light from an injection-locked Nd:YAG laser. Before reaching the interferometer, the light is passed through two 8-m triangular mode-cleaning cavities. During S1 approximately 2 W of light was incident on the interferometer. A power recycling mirror of 1% transmission was installed to increase the effective laser power available for the measurement.

LIGO comprises three power-recycled Michelson interferometers with resonant Fabry-Perot cavity arms. A 4-km and a 2-km interferometer are collocated at the Hanford site and are designated H1 and H2, respectively, and a 4-km interferometer at the Livingston site is designated L1. Each interferometer employs a Nd:YAG laser stabilized using a monolithic reference cavity and a 12-m mode-cleaning cavity.

In all four instruments the beam splitters, recycling mirrors, and test masses are hung as pendulums from multilayer seismic isolation filters to isolate them from local forces. The masses and beam paths are housed in high-vacuum enclosures to preclude optical scintillation and acoustic interference.

Sinusoidal calibration forces of known amplitude were applied to the test bodies throughout the observing run. These signals were recovered from the data stream and used to periodically update the scale factors linking the recorded signal amplitude to strain. The principal calibration uncertainties arise from the imprecision in the electromechanical coupling coefficients of the force actuators. These were estimated by comparison with the known laser wavelength by actuating a test mass between interference fringes. In the Hanford interferometers, the calibration was also verified

^aCurrently at Stanford Linear Accelerator Center.

^bPermanent address: HP Laboratories.

^cCurrently at Rutherford Appleton Laboratory.

^dCurrently at University of California, Los Angeles.

^eCurrently at Hofstra University.

^fCurrently at Siemens AG.

^gPermanent address: GReCO, Institut d'Astrophysique de Paris (CNRS).

^hCurrently at NASA Jet Propulsion Laboratory.

ⁱCurrently at National Science Foundation.

^jCurrently at University of Sheffield.

^kCurrently at Ball Aerospace Corporation.

^lCurrently at European Gravitational Observatory.

^mCurrently at Intel Corp.

ⁿCurrently at Lightconnect Inc.

^oCurrently at Keck Observatory.

^pCurrently at ESA Science and Technology Center.

^qCurrently at Raytheon Corporation.

^rCurrently at Mission Research Corporation.

^sCurrently at Harvard University.

^tCurrently at Lockheed-Martin Corporation.

^uCurrently at NASA Goddard Space Flight Center.

^vPermanent address: University of Tokyo, Institute for Cosmic Ray Research.

^wCurrently at The University of Texas at Brownsville and Texas Southmost College.

^xCurrently at Laboratoire d'Annecy-le-Vieux

de Physique des Particules.

^yCurrently at LIGO-California Institute of Technology.

^zPermanent address: University College Dublin.

^{aa}Currently at Research Electro-Optics Inc.

^{bb}Currently at Institute of Advanced Physics, Baton Rouge, LA.

^{cc}Currently at Cardiff University.

^{dd}Currently at European Commission, DG Research, Brussels, Belgium.

^{ee}Currently at Spectra Physics Corporation.

^{ff}Currently at University of Chicago.

^{gg}Currently at LightBit Corporation.

^{hh}Currently at University of Delaware.

ⁱⁱCurrently at Carl Zeiss GmbH.

^{jj}Permanent address: NASA Jet Propulsion Laboratory.

^{kk}Currently at Shanghai Astronomical Observatory.

^{ll}Currently at Laser Zentrum Hannover.

^{mm}<http://www.ligo.org>

against piezoelectric displacement transducers connected to mirror support structures. For the S1 observations, the net amplitude uncertainty near 1.3 kHz was estimated at $\pm 4\%$ for GEO, $\pm 10\%$ for each of the LIGO interferometers. These uncertainties are mostly due to errors in the measurement of the actuator's strengths and in the determination of the time-varying optical gains. The more complex Fabry-Perot optical configuration employed by LIGO contributes some additional calibration uncertainty over that of GEO. Details of the calibration methods can be found in [1] and Refs. [42] and [43] therein.

B. Expected sensitivity

We define the gravitational wave strength h_0 of a continuous signal from a given source as the maximum peak amplitude which could be received by an interferometer if the orientations of the pulsar and detector were both optimal. Thus, h_0 depends on the intrinsic emission strength and source distance, but not on the inclination of the pulsar's spin axis or on the antenna pattern of the detector.

The calibrated interferometer strain output is a time series

$$s(t) = h(t) + n(t), \quad (2.1)$$

where $h(t)$ is the received signal, $n(t)$ is the detector noise, and t is the time in the detector's frame.

The noise $n(t)$ is characterized by its single-sided power spectral density $S_n(f)$. Assuming this noise is Gaussian and taking some fixed observation time¹ T , we can compute the amplitude h_0 of a putative continuous signal which would be detected in, e.g., 90% of experimental trials if truly present, but would arise randomly from the noise background in only 1% of trials (what we call a 1% "false alarm rate" and a 10% "false dismissal rate").

If we fix a false alarm rate, it is clear that the lower the desired false dismissal rate, the higher the signal needs to be. The detection statistic used in Sec. IV C provides the lowest false dismissal rate for a given false alarm rate and signal strength and it is thus optimal in the Neyman-Pearson sense (see, for example, [10]). The amplitude of the average signal that we could detect in Gaussian stationary noise with a false alarm rate of 1% and a false dismissal rate of 10% using the detection statistic described in [11] is given by²

$$\langle h_0 \rangle = 11.4 \sqrt{S_n(f_s)/T}, \quad (2.2)$$

where f_s is the frequency of the signal.³ The upper curves in

¹Here we presume that we know the position, frequency, and spin-down parameters of the source and that T is between a few days and several months.

²The average is over different positions, inclinations, and polarizations of the source.

³This differs from [12] for three reasons: (1) the h_0 used here is twice that defined in [12], (2) we use a different statistic for this detection problem (a chi-square distribution with four degrees of freedom), and (3) we have specified a false dismissal rate of 10% whereas the derivation in [12] has an implicit false dismissal rate of about 50%. If we use this false dismissal rate and the \mathcal{F} statistic, we get $\langle h_0 \rangle = 7.6 \sqrt{S_n(f_s)/T}$.

Fig. 1 show $\langle h_0 \rangle$ for the LIGO and GEO detectors during S1. Observation times for respective interferometers are given in the figure. Because of ground motion, equipment failures, and alignment drifts, the four interferometers were not always fully operational during the S1 run; thus, the observation times vary from detector to detector.

The lower curves in Fig. 1 represent $\langle h_0 \rangle$ corresponding to the design sensitivity of the various detectors. An observation of $T = 1$ yr was assumed.

The solid circles in Fig. 1 show the constraints that measurements of spin-down rates of known pulsars place on the expected gravitational wave signal, under the assumption that the pulsars are rigid rotators with a moment of inertia of 10^{45} g cm² and that all of the observed spin-down rate is due to the emission of gravitational waves.

As shown in Fig. 1, under the above assumptions no detection is expected for any known pulsar at the sensitivity achieved during the S1 run. Furthermore, many known pulsars are rotating too slowly to be detected by the initial ground-based interferometers. However, the number of millisecond pulsars observed in this band continues to increase with new radio surveys, and the known targets plotted here constitute a highly selected sample. Future searches for previously undiscovered rotating neutron stars using the methods presented here will sample a different and potentially much larger subset of the total population.

III. PERIODIC GRAVITATIONAL WAVES

A. Expected emission by neutron stars

The strongest argument that some neutron stars (NSs) are emitting gravitational waves (GWs) with amplitude detectable by Advanced LIGO [13], $h_0 \approx 10^{-27} - 10^{-26}$, is due to Bildsten [14,15]. He noted that the inferred rotation frequencies of low-mass x-ray binaries (LMXBs) are all clustered in the range $f_r \sim 270 - 620$ Hz (an inference strengthened by the recent observations of [16,17]), whereas *a priori* there should be no cutoff in f_r , up to the (estimated) NS breakup frequency of ~ 1.5 kHz. Updating a suggestion by Wagoner [18,19], Bildsten proposed that LMXBs have reached an equilibrium where spin-up due to accretion is balanced by spin-down from GW emission. Since the GW spin-down torque scales like f_r^5 , a wide range of accretion rates then leads to a rather narrow range of equilibrium rotation rates, as observed.

Millisecond pulsars (MSPs) are generally believed to be recycled pulsars: old pulsars that were spun up by accretion during an LMXB phase [20,21]. The rotation rates of MSPs also show a high-frequency cutoff [15]; the fastest (PSR J1939+2134) has $f_r = 642$ Hz. If the GWs that arrest the spin up of accreting NSs continue to be emitted in the MSP phase (e.g., because of some persistent deformation of the NS shape away from axisymmetry), then they could also account for the observed spin down of MSPs. In this case, the GW amplitudes of MSPs would in fact be (very close to) the "spin-down upper limits" shown in Fig. 1. (Note that the MSP spin-down rate is generally attributed entirely to the pulsar magnetic field; indeed, pulsar magnetic fields are typi-




FIG. 1. (Color) Upper curves: characteristic amplitude $\langle h_0 \rangle$ of a known monochromatic signal detectable with a 1% false alarm rate and a 10% false dismissal rate by the GEO and LIGO detectors at S1 sensitivity and with an observation time equal to the up-time of the detectors during S1 (GEO: 401 h, L1: 137 h, H1: 209 h, H2: 214 h). Lower curves: $\langle h_0 \rangle$ for the design sensitivities of the detectors for an assumed 1-yr observation time. Solid circles: upper limit on $\langle h_0 \rangle$ from the measured spin-down rate of known radio pulsars assuming a moment of inertia of 10^{45} g cm^2 . These upper limits were derived under the assumption that all the measured loss of angular momentum of the star is due to the emission of gravitational waves, neglecting the spin-down contribution from electromagnetic and particle emission. The arrow points to the solid circle representing pulsar J1939+2134.

cally inferred this way. However, there appears to be no strong evidence supporting this inference.)

We now turn to the possible physical mechanisms responsible for periodic GWs in this frequency range. The main possibilities that have been considered are (1) NS spin precession, (2) an excited NS oscillation mode (most likely the r -mode), and (3) small distortions of the NS shape away from axisymmetry. At present, the third mechanism (small ellipticity) seems the most plausible source of detectable GWs, and in this paper we set upper limits for this particular mechanism (the three mechanisms predict three different GW frequencies for the same observed rotation frequency). However, we begin by briefly commenting on the other two possibilities.

A NS precesses (or “wobbles”) when its angular momentum \mathbf{J} is not aligned with any principal axis of its inertia tensor. A wobbling NS emits GWs at the inertial-frame precession frequency, which is very nearly the rotation frequency f_r . While large-amplitude wobble could plausibly produce GW amplitudes $h_0 \sim 10^{-27}$ over short time scales, the problem with this mechanism is that dissipation should damp NS wobble quickly [22]; while this dissipation time scale is quite uncertain (it is perhaps of the order of a year for a MSP), it is almost certainly orders of magnitude shorter than the typical lifetimes of MSPs. So unless some mechanism is found that regularly reexcites large-amplitude wobble, it is unlikely that any nearby MSP would be wobbling. Moreover, most MSPs have highly stable pulse shapes and typically appear *not* to be wobbling substantially. In particular, the single-pulse characteristics of PSR J1939+2134 have been observed to be extremely stable with no pulse-to-pulse variation except for occasional giant pulses [23]. It has been verified through radio observations that PSR J1939

+2134 continued to spin according to a simple spin-down model during S1 [24].

r -modes (modes driven by Coriolis forces) have been a source of excitement among GW theorists since 1998, when Andersson [25] and Friedman and Morsink [26] showed that they should be unstable due to gravitational back reaction (the Chandrasekhar-Friedman-Schutz instability). Nonlinear mode-mode coupling is predicted to saturate the growth of r -modes at dimensionless amplitude $\alpha \leq 10^{-3} (f_r/\text{kHz})^{5/2}$ [27]. This implies r -mode radiation from nascent NSs in extragalactic supernovas will not be detectable, but r -mode GWs from old, recycled Galactic NSs could still be detectable by Advanced LIGO. For example, GWs from an excited r mode could balance the accretion torque in accreting NSs, as in the Wagoner-Bildsten mechanism.

We now turn to GWs from small nonaxisymmetries in the NS shape. If h_0 is the amplitude of the signal at the detector from an optimally oriented source, as described above, and if we assume that the emission mechanism is due to deviations of the pulsar’s shape from perfect axial symmetry, then

$$h_0 = \frac{4\pi^2 G_N}{c^4} \frac{I_{zz} f_s^2}{r} \epsilon, \quad (3.1)$$

where r is the distance to the NS, I_{zz} is its principal moment of inertia about the rotation axis, $\epsilon \equiv (I_{xx} - I_{yy})/I_{zz}$ is its ellipticity, and the gravitational wave signal frequency f_s is exactly twice the rotation frequency f_r . Here G_N is Newton’s constant, and c is the speed of light. This is the emission mechanism that we assume produces the gravitational wave signal that we are targeting.

One possible source of ellipticity is tiny “hills” in the NS

crust, which are supported by crustal shear stresses. In this case, the maximum ellipticity is set by the crustal breaking strain $\bar{\sigma}_{\max}$ [28]:

$$\epsilon_{\max} \sim 5 \times 10^{-8} (\bar{\sigma}_{\max}/10^{-3}). \quad (3.2)$$

The coefficient in Eq. (3.2) is low both because the NS crust is rather thin (compared to the NS radius) and because the crust shear modulus μ is small compared to the ambient pressure p : $\mu/p \sim 10^{-3} - 10^{-2}$. (If NSs have solid cores, as well as crusts, then much larger ellipticities would be possible.) For the LMXBs, Ushomirsky, Cutler, and Bildsten [28] showed that lateral temperature variations in the crust of order 5% or lateral composition variations of order 0.5% (in the charge-to-mass ratio) could build up NS ellipticities of order $10^{-8} - 10^{-7}$, but only if the crust breaking strain is large enough to sustain such hills.

Strong internal magnetic fields are another possible source of NS ellipticity. Cutler [29] has argued that if a NS interior magnetic field B has a toroidal topology (as expected if the B field was generated by strong differential rotation immediately after collapse), then dissipation tends to reorient the symmetry axis of the toroidal B field perpendicular to the rotation axis, which is the ideal orientation for maximizing equatorial ellipticity. Toroidal B fields of the order of $10^{12} - 10^{13}$ G would lead to sufficient GW emission to halt the spin-up of LMXBs and account for the observed spin-down of MSPs.

B. Signal received from an isolated pulsar

A gravitational wave signal we detect from an isolated pulsar will be amplitude modulated by the varying sensitivity of the detector as it rotates with the Earth (the detector “antenna pattern”). The detected signal has the form [11]

$$h(t) = F_+(t, \psi) h_0 \frac{1 + \cos^2 \iota}{2} \cos \Phi(t) + F_\times(t, \psi) h_0 \cos \iota \sin \Phi(t), \quad (3.3)$$

where ι is the angle between neutron star’s spin direction \hat{s} and the propagation direction of the waves, \hat{k} , and $\Phi(t)$ is the phase evolution of the signal. $F_{+, \times}$ are the strain antenna patterns of the detector to the plus and cross polarizations and are bounded between -1 and 1 . They depend on the orientation of the detector and source and on the polarization of the waves, described by the polarization angle ψ .⁴

The signal will also be Doppler shifted by the orbital motion and rotation of the Earth. The resulting phase evolution of the received signal can be described by a truncated Taylor series as

⁴Following the conventions of [11], ψ is the angle (clockwise about \hat{k}) from $2 \times \hat{k}$ to $\hat{k} \times \hat{s}$, where \hat{z} is directed to the North Celestial Pole. $\hat{k} \times \hat{s}$ is the x axis of the wave frame—also called the wave’s principal+ polarization direction.




FIG. 2. Histogram of timing residuals between our barycentering routines and TEMPO, derived by comparing the phase evolution of test signals produced by the two software packages. Here 156 locations in the sky were chosen at random and the residuals calculated once an hour for the entire year 2002. The maximum timing error is $< 4 \mu\text{s}$.

$$\Phi(t) = \phi_0 + 2\pi \left[f_s(T - T_0) + \frac{1}{2} \dot{f}_s(T - T_0)^2 + \frac{1}{6} \ddot{f}_s(T - T_0)^3 \right], \quad (3.4)$$

where

$$T = t + \delta t = t - \frac{\mathbf{r}_d \cdot \hat{\mathbf{k}}}{c} + \Delta_{E\odot} - \Delta_{S\odot}. \quad (3.5)$$

Here T is the time of arrival of a signal at the solar system barycenter (SSB), ϕ_0 is the phase of the signal at fiducial time T_0 , \mathbf{r}_d is the position of the detector with respect to the SSB, and $\Delta_{E\odot}$ and $\Delta_{S\odot}$ are the solar system Einstein and Shapiro time delays, respectively [30].

The timing routines used to compute the conversion between terrestrial and SSB time [Eq. (3.5)] were checked by comparison with the widely used radio astronomy timing package TEMPO [31]. This comparison (Fig. 2) confirmed an accuracy of better than $\pm 4 \mu\text{s}$, thus ensuring no more than 0.01 rad phase mismatch between a putative signal and its template. This results in a negligible fractional signal-to-noise ratio loss, of order $\sim 10^{-4}$.

Table I shows the parameters of the pulsar that we have chosen to illustrate our analysis methods [32].

IV. DATA ANALYSES

A. Introduction

Two independent search methods are presented here: (i) a frequency domain method which can be employed for ex-

TABLE I. Parameters for the target pulsar of the analyses presented here, PSR J1939+2134 (also designated PSR B1937+21). Numbers in parentheses indicate uncertainty in the last digit.

Right ascension (J2000)	19 ^h 39 ^m 38 ^s .560 210(2)
Declination (J2000)	+21°34′59″.141 66(6)
RA proper motion	−0.130(8) mas yr ^{−1}
Dec proper motion	−0.464(9) mas yr ^{−1}
Period (1/ f_s)	0.001 557 806 468 819 794(2) s
Period derivative	1.051 193(2) × 10 ^{−19} s s ^{−1}
Epoch of period and position	MJDN 47 500

ploring large parameter space volumes and (ii) a time domain method for targeted searches of systems with an arbitrary but known phase evolution.

Both approaches will be used to cast an upper limit on the amplitude of the periodic gravitational wave signal: a Bayesian approach for the time domain analysis and a frequentist approach for the frequency domain analysis. These approaches provide answers to two different questions and therefore should not be expected to result in the exact same numerical answer [33,34]. The frequentist upper limit refers to the reliability of a procedure for identifying an interval that contains the true value of h_0 . In particular, the frequentist confidence level is the fraction of putative observations in which, in the presence of a signal at the level of the upper limit value identified by the actual measurement, $h_0^{95\%}$, the upper limit identified by the frequentist procedure would have been higher than $h_0^{95\%}$. The Bayesian upper limit, on the other hand, defines an interval in h_0 that, based on the observation made and on prior beliefs, includes the true value with 95% probability. The probability that we associate with the Bayesian upper limit characterizes the uncertainty in h_0 given the observation made. This is distinct from the reliability, evaluated over an ensemble of observations, of a procedure for identifying intervals.

All the software used for the analyses is part of the publicly available LSC Algorithm Library (LAL) [35]. This is a library that comprises roughly 700 functions specific to gravitational wave data analysis.

B. Statistical characterization of the data

As a result of the narrow frequency band in which the target signal has appreciable energy, it is most convenient to characterize the noise in the frequency domain. We divided the data into 60-s blocks and took the Fourier transform of each. The resulting set of Fourier transforms will be referred to as short-time-baseline Fourier transforms (SFTs) and is described in more detail in Sec. IV C 1.

The frequency of the pulsar signal at the beginning of the observation for every detector is reported in Table II. Also reported is the value of the spin-down parameter expressed in units of Hz s^{−1}. We have studied the statistical properties of the data in a narrow frequency band (0.5 Hz) containing the emission frequency. This is the frequency search region, as well as the region used for estimating both the noise background and detection efficiency. Figure 3 summarizes our findings. Two types of distributions are plotted. The first col-

TABLE II. Run parameters for PSR J1939+2134. The different emission frequencies correspond to the different initial epochs at which each of the searches began. Numbers in parentheses indicate the uncertainty in the last digit or digits.

Spin-down parameter f_s	−8.6633(43) × 10 ^{−14} Hz s ^{−1}
f_s at start of GEO observation	1283.856 487 705(5) Hz
f_s at start of L1 observation	1283.856 487 692(5) Hz
f_s at start of H1 observation	1283.856 487 687(5) Hz
f_s at start of H2 observation	1283.856 487 682(5) Hz

umn shows the distributions of bin power; for each SFT (labeled by α) and for every frequency bin (labeled by $1 \leq k \leq M$) in the band 1283.75–1284.25 Hz, we have computed the quantity

$$P_{\alpha k} = \frac{|\bar{x}_{\alpha k}|^2}{\sum_k |\bar{x}_{\alpha k'}|^2 / M}, \quad (4.1)$$

where $\bar{x}_{\alpha k}$ is the SFT datum at frequency index k of the α th SFT and have histogrammed these values. If the data are Gaussian and if the different frequency bins in every SFT are independent realizations of the same random process, then we expect the normalized power variable described above ($P_{\alpha k}$) to follow an exponential distribution with a mean and standard deviation of 1, as shown by the dashed line. The circles are the experimental points. The standard deviation of the measured distribution for GEO data is 0.95. The LIGO Livingston, Hanford 4-km, and Hanford 2-km data are also shown in Fig. 3. The standard deviation of the $P_{\alpha k}$ for all of these is 0.97.

The plots in the second column of Fig. 3 show the distribution of phase differences between adjacent frequency bins. With the same notation as above, we have computed the quantity

$$\Delta\Phi_{\alpha k} = \Phi_{\alpha k} - \Phi_{\alpha k-1}, \quad (4.2)$$

where $\Phi_{\alpha k}$ is the phase of the SFT datum at frequency index k of the α th SFT and the difference is reduced to the range $[-\pi, \pi]$. Therefore, $\Delta\Phi_{\alpha k}$ is the distance in phase between data at adjacent frequency bins. If the data were from a purely random process, we expect this distribution to be uniform between $-\pi$ and π , as observed.

Figure 4 shows the average value of $\sqrt{S_n}$ over a 1-Hz band from 1283.5 to 1284.5 Hz as a function of time in days for the entire S1 run starting from the beginning of S1 (15:00 UTC, 23 August 2002). These plots monitor the stationarity of the noise in the band of interest over the course of the run.

Figure 5 shows $\sqrt{S_n}$ as a function of frequency between 1281 and 1285 Hz. During S1, the received signal is expected to have a frequency of 1283.8 Hz. This frequency is shown as a dashed vertical line. During the S1 observation time, the Doppler modulation changed this signal frequency by no more than 0.03 Hz, two SFT frequency bins. For these plots S_n has been estimated by averaging the power in each frequency bin over the entire S1 run. A broad spectral feature is observed in the GEO data. This feature is 0.5 Hz wide, comparatively broad with respect to the expected Doppler

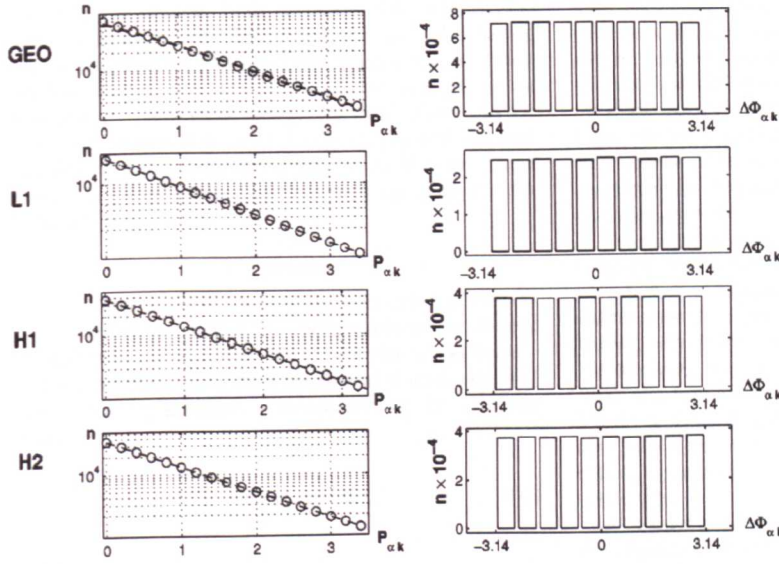


FIG. 3. Histograms of $P_{\alpha k}$ and of $\Delta\Phi_{\alpha k}$ for the four detectors.

shift of the target signal, and represents only a 10% perturbation in the local power spectral density.

C. Frequency domain technique

1. Short-time-baseline Fourier transforms

In principle, the only constraint on the time baseline of the SFTs used in the frequency domain analysis is that the instantaneous frequency of a putative signal not shift during the time baseline by more than half a frequency bin. For frequencies in the kilohertz range this implies a maximum time baseline of the order of 30 min (having assumed an observation time of several months and a source declination

roughly the same as the latitude of the detector). However, in practice, since we are also estimating the noise on the same time baseline, it is advisable for the time baseline to be short enough to follow the nonstationarities of the system. On the other hand, for the frequency domain analysis, the computational time required to carry out a search increases linearly with the number of Fourier transforms. Thus the shorter the time baseline, the higher the computational load. We have chosen for the S1 run a time baseline of 60 s as a compromise between the two opposing needs.

Interruptions in interferometer operation broke each time series into segments separated by gaps representing invalid or contaminated data. Only valid data segments were in-




FIG. 4. The square root of the average value of S_n for all four interferometers over a band of 1 Hz starting at 1283.5 Hz versus time in days starting at the beginning of S1 (23 August 2002, 15:00 UTC).




FIG. 5. $\sqrt{S_n}$ in a band of 4 Hz (starting at 1281 Hz) using the entire S1 data set analyzed from the four interferometers. The noise $\sqrt{S_n}$ is shown in units of $10^{-20} \text{ Hz}^{-1/2}$. The dashed vertical line indicates the expected frequency of the signal received from J1939+2134.

cluded in the analysis. Each valid 60-s data segment was filtered with a fifth-order Butterworth high-pass filter having a knee frequency of 100 Hz. Then a nearly-flat-top Tukey window function was applied to each data segment in the time domain. The window changes the value of less than 1% of the data in each 60-s chunk. Each data segment was then fast Fourier transformed and written to an SFT file. These SFTs were computed once and then used repeatedly for different analyses.

2. \mathcal{F} statistic

The detection statistic that we use is described in [11]. As in [11] we call this statistic \mathcal{F} ,⁵ though differences between our definition and that given in [11] are pointed out below.

The \mathcal{F} statistic derives from the method of maximum likelihood. The log-likelihood function $\ln \Lambda$ is, for Gaussian noise,

$$\ln \Lambda = (s|h) - \frac{1}{2}(h|h), \quad (4.3)$$

where

$$(s|y) = 4\mathcal{R} \int_0^\infty \frac{\tilde{s}(f)\tilde{y}^*(f)}{S_n(f)} df, \quad (4.4)$$

s is the calibrated detector output time series, h is the target signal (commonly referred to as the template), the tilde is the Fourier transform operator, and $S_n(f)$ is the one-sided power spectral density of the noise. The \mathcal{F} statistic is the maximum value of $\ln \Lambda$ with respect to all unknown signals parameters, given our data and a set of known template parameters. In fact, if some or all of the signal's parameters are unknown, it is standard practice to compute the likelihood for different template parameters and look for the highest values. The maximum of the likelihood function is the statistic of choice for matched filtering methods, and it is the optimal detection statistic as defined by the Neyman-Pearson criterion: the lowest false dismissal rate at a fixed false alarm rate (see, for example, Sec. II B).

In our case the known parameters are the position of the source (α, δ angles on the celestial sphere), the emission frequency f_s , and the first-order spin-down parameter value \dot{f}_s . The unknown parameters are the orientation of the pulsar (angle ι), the polarization state of the wave (angle ψ), its initial phase ϕ_0 , and the wave amplitude h_0 .

The core of the calculation of \mathcal{F} consists in computing integrals of the type given in Eq. (4.4), using templates for the two polarizations of the wave. The results are optimally combined as described in [11] except we consider a single-frequency-component signal. Also, we do not treat $S_n(f)$ as constant in time: we reestimate it every 60 s (for every α), based on the average $|\bar{x}_{ak}|^2$ in a 0.5-Hz band around the

signal frequency. Thus, while the method is defined in [11] in the context of stationary Gaussian noise, we adapt it so that it can be used even when the noise is nonstationary. The calculation is easily performed in the frequency domain since the signal energy is concentrated in a narrow frequency band. Using the SFTs described in Sec. IV C 1, some approximations can be made to simplify the calculation and improve computational efficiency while still recovering most ($>98\%$) of the signal power.

The method of computing \mathcal{F} was developed for a specific computational architecture: a cost-effective Beowulf cluster, which is an ensemble of loosely coupled processors with simple network architecture. This becomes crucial when exploring very large parameter-space volumes for unknown sources using long observation periods, because the search depth and breadth are limited by computational resources. The S1 analyses described here were carried out using Condor [36] on the Merlin and Medusa clusters at the AEI and UWM, respectively [37,38]. Each cluster has 300 independent CPUs.

As a point of reference, we note that it takes of order of a few seconds of CPU time on a 1.8-GHz-class CPU to determine the \mathcal{F} statistic for a single template with ~ 16 d of observation time.

3. Setting an upper limit on h_0

The outcome \mathcal{F}^* of a specific targeted search represents the optimal detection statistic for that search. Over an independent ensemble of similar searches in the presence of stationary Gaussian noise, $2\mathcal{F}^*$ is a random variable that follows a χ^2 distribution with four degrees of freedom. If the data also contain a signal, this distribution has a noncentrality parameter λ proportional to the time integral of the squared signal.

Detection of that signal would be signified by a large value \mathcal{F}^* unlikely to have arisen from the noise-only distribution. If instead the value is consistent with pure noise (as we find in this instance), we can place an upper limit on the strength of any signal present, as follows.

Let \mathcal{F}^* be the value of the detection statistic in our actual experiment. Had there existed in the data a real signal with amplitude greater than or equal to $h_0(C)$, then in an ensemble of identical experiments with different realizations of the noise, some fraction C of trials would yield a detection statistic exceeding the value \mathcal{F}^* . We will therefore say that we have placed an upper limit $h_0(C)$ on the strength of the targeted signal, with confidence C . This is a standard frequentist upper limit.

To determine the probability distribution $p(2\mathcal{F}|h_0)$, we produce a set of simulated artificial signals with fixed amplitude h_0 from fictional pulsars at the position of our target source and with the same spin-down parameter value, but with intrinsic emission frequencies that differ from it by a few tenths of a hertz. We inject each of these artificial signals into our data and run a search with a perfectly matched template. For each artificial signal we obtain an independent value of the detection statistic; we then histogram these values. If the SFT data in nearby frequency bins (of order 100 bins) can be considered as different realizations of the same

⁵Note that this detection statistic has nothing to do with the F statistic of the statistical literature, which is ratio of two sample variances, or the F test of the null hypothesis that the two samples are drawn from distributions of the same variance.

TABLE III. Summary of the frequency domain search analyses. T_{obs} indicates the total duration of the analyzed data set. \mathcal{F}^* is the measured value of the detection statistic. $P_0(2\mathcal{F}^*)$ is the probability of getting this value or greater by chance—i.e., in the absence of any signal. h_0^{inject} is the amplitude of the population of fake signals that were injected in the data set such that, when searched for with a perfectly matched template, $C_{\text{meas}}\%$ of the time the resulting value of \mathcal{F} was greater than \mathcal{F}^* . $\langle 1/S_n \rangle$ is the average value of the inverse of the noise in a small frequency band around the target frequency. U_0 is the time integral of the square of the targeted signal with an amplitude of 2×10^{-19} , at the output of the interferometers, for observations times equal to T_{obs} and in the absence of noise. λ_{exp} is the value of the noncentrality parameter that one expects for the distribution of \mathcal{F} from searches with perfectly matched templates on a population of injected signals with amplitude h_0^{inject} and noise with average power $\langle 1/S_n \rangle^{-1}$. $\lambda_{\text{best-fit}}$ is the best-fit noncentrality parameter value derived from the distribution $p(2\mathcal{F}|h_0^{\text{inject}})$ derived from the software signal injections and searches with perfectly matched templates. C_{exp} and $C_{\text{best-fit}}$ are the corresponding confidence values for \mathcal{F}^* .

IFO	T_{obs} [d]	h_0^{inject}	$2\mathcal{F}^*$	$P_0(2\mathcal{F}^*)$	$\langle 1/S_n \rangle^{-1}$ [Hz $^{-1}$]	$U_0/10^{-33}$ [s]	λ_{exp}	$\lambda_{\text{best-fit}}$	C_{exp}	$C_{\text{best-fit}}$	$C_{\text{meas}} \pm \Delta C$
GEO	16.7	1.94×10^{-21}	1.5	0.83	5.3×10^{-38}	1.0	3.6	3.3	95.7%	95.2%	$95.01 \pm 0.23\%$
L1	5.73	2.70×10^{-22}	3.6	0.46	1.4×10^{-40}	0.37	9.6	8.3	96.7%	95.0%	$95.00 \pm 0.23\%$
H1	8.73	5.37×10^{-22}	6.0	0.20	5.4×10^{-40}	0.5	13.3	12.8	96.6%	95.0%	$95.00 \pm 0.23\%$
H2	8.90	3.97×10^{-22}	3.4	0.49	3.8×10^{-40}	0.45	9.3	7.9	96.8%	95.0%	$95.00 \pm 0.23\%$

random process (justified in Sec. IV B), then it is reasonable to assume that the normalized histogram represents the probability density function $p(2\mathcal{F}|h_0)$. One can then compute the confidence

$$C(h_0) = \int_{2\mathcal{F}^*}^{\infty} p(2\mathcal{F}|h_0) d(2\mathcal{F}), \quad (4.5)$$

where $h_0(C)$ is the functional inverse of $C(h_0)$. In practice, the value of the integral in Eq. (4.5) is calculated directly from our simulations as follows: we count how many values of \mathcal{F} are greater or equal to \mathcal{F}^* and divide this number by the total number of \mathcal{F} values. The value derived in this way does not rely on any assumptions about the shape of the probability distribution function (PDF) curve $p(2\mathcal{F}|h_0)$.

There is one more subtlety that must be addressed: all eight signal parameters must be specified for each injected artificial signal. The values of source position and spin-down parameters are known from radio data and are used for these injections. Every injected signal has a different frequency, but all such frequencies lie in bins that are close to the expected frequency of the target signal, 1283.86 Hz. The values of ι and ψ are not known, and no attempt has been made in this analysis to give them informative priors based on radio data. However, the value of the noncentrality parameter that determines the $p(2\mathcal{F}|h_0)$ distribution does depend on these values. This means that, for a given \mathcal{F}^* , a different confidence level can be assigned for the same signal strength, depending on the choice of ι and ψ .

There are two ways to proceed: either inject a population of signals with different values of ι and ψ , distributed according to the priors on these parameters,⁶ or pick a single value for ι and for ψ . In the latter case it is reasonable to choose the most pessimistic orientation and polarization of the pulsar with respect to the detector during the observation time. For fixed signal strength, this choice results in the lowest confidence level and thus, at fixed confidence, in the most conservative upper limit on the signal strength. We have de-

cided to use in our signal injection the worst-case values for ι (which is always $\pi/2$) and ψ —i.e., the values for which the noncentrality parameter is the smallest.

4. Frequency domain S1 analysis for PSR J1939+2134

Table III summarizes the results of the frequency domain analysis. For every interferometer (column 1) the value of the detection statistic for the search for J1939+2134 is reported: $2\mathcal{F}^*$, shown in column 4. Next to it is the corresponding value of the chance probability:

$$P_0(2\mathcal{F}^*) = \int_{2\mathcal{F}^*}^{\infty} p(2\mathcal{F}|h_0=0) d(2\mathcal{F}), \quad (4.6)$$

our estimate of how frequently one would expect to observe the measured value of \mathcal{F}^* or greater in the absence of a signal. As can be seen from $P_0(2\mathcal{F}^*)$, the measured values of $2\mathcal{F}^*$ are not significant; we therefore conclude that there is no evidence of a signal and proceed to set an upper limit.

T_{obs} is the length of the live-observation time. h_0^{inject} is the amplitude of the population of injected signals that yielded a 95% confidence. The upper limit $h_0^{95\%}$ differs from h_0^{inject} only by the calibration uncertainty, as explained in Sec. IV E. Here C_{meas} is the confidence level derived from the injections of fake signals, and ΔC its estimated uncertainty due to the finite sample size of the simulation.

The quantities in the remaining columns can be used to evaluate how far the reported results are from those that one expects. The results shown are remarkably consistent with what one expects based on the noise and on the injected signal: the confidence levels that we determine differ from the expected ones by less than 2%.

Given a perfectly matched template, the expected noncentrality parameter when a signal $h(t)$ is added to white noise with spectral density S_n is

$$\lambda = \frac{2U}{S_n}, \quad (4.7)$$

⁶The time domain analysis assumes uniform priors on $\cos \iota$ and ψ .

where $U = \int_{T_{\text{obs}}} |h(t)|^2 dt$. Here U can also be computed by feeding the analysis pipeline pure signal and by performing the search with a perfectly matched template⁷ having set $S_n(f) = 1$ s. In Table III we report the values of U_0 , for the worst-case $h(t)$ signals for PSR J1939+2134 as “seen” by the interferometers during their respective observation times and with $h_0 = 2 \times 10^{-19}$. The different values of U_0 reflect the different durations of the observations and the different orientations of each detector with respect to the source. The expected value of the noncentrality parameter can be estimated as

$$\lambda_{\text{exp}} = 2U_0 \langle 1/S_n \rangle \left(\frac{h_0^{\text{inject}}}{2 \times 10^{-19}} \right)^2. \quad (4.8)$$

If the noise were stationary, then S_n may be easily determined. Our noise is not completely stationary, so the value determined for the noncentrality parameter λ is sensitive to the details of how S_n is estimated. The value of $\langle 1/S_n \rangle$ used to determine the expected value of λ is computed as

$$\langle 1/S_n \rangle = \frac{\Delta t}{N} \sum_a \frac{1}{\sum_k |\tilde{x}_{ak}|^2 / M}, \quad (4.9)$$

where the frequency index k varies over a band ~ 0.2 Hz around 1283.89 Hz. Here N and Δt are the number of samples and the sampling time of the 60-s time series that are Fourier transformed. We choose an harmonic mean rather than an arithmetic mean because this is the way S_n enters the actual numerical calculation of the \mathcal{F} statistic. This method is advantageous because the estimate it produces is relatively insensitive to very large outliers that would otherwise bias the estimate.

λ_{exp} is the expected value of the noncentrality parameter based on S_n and h_0^{inject} , and $\lambda_{\text{best-fit}}$ is the best-fit value of the noncentrality parameter based on the measured distribution of \mathcal{F} values from the simulation. C_{exp} and $C_{\text{best-fit}}$ are the confidence levels corresponding to these distributions integrated between $2\mathcal{F}^*$ and ∞ .

Figure 6 shows the distributions for $p(2\mathcal{F}|h_0^{\text{inject}})$. The circles result from the simulations described above. The solid lines show the best fit noncentral χ^2 curves. The shaded region is the integral of $p(2\mathcal{F}|h_0^{\text{inject}})$ between $2\mathcal{F}^*$ and ∞ . By definition, this area is 0.95.

D. Time domain search technique

1. Overview

Frequency domain methods offer high search efficiencies when the frequency of the signal and/or the position of the neutron star are unknown and need to be determined along with the other signal parameters. However, in the case of known pulsars, where both the intrinsic rotation frequency of

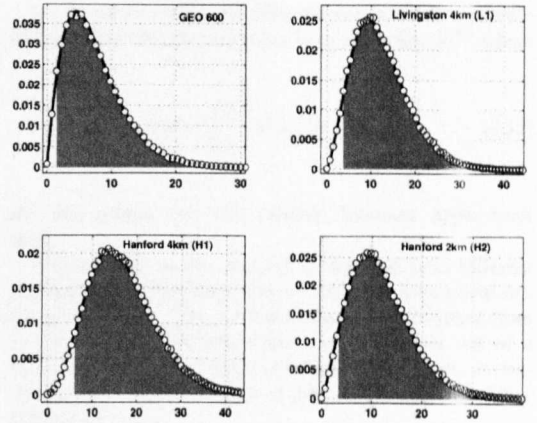


FIG. 6. Measured pdf for $2\mathcal{F}$ for all four interferometer data with injected signals as described in Table III. The circles represent the measured PDF values from the Monte Carlo simulations. The lines represent χ^2 distributions with four degrees of freedom and best-fit noncentrality parameters given in Table III. The filled area represents the integral of the pdfs between $2\mathcal{F}^*$ and $+\infty$.

the neutron star and its position are known to high accuracy, alternative time domain methods become attractive. At some level the two domains are of course equivalent, but issues such as data dropouts and the handling of signals with complicated phase evolutions can be conceptually (and practically) more straightforward in a time series analysis than in an analysis based on Fourier transforms.

The time domain search technique employed here involves multiplying (heterodyning) the quasisinusoidal signal from the pulsar with a unit-amplitude complex function that has a phase evolution equal but of opposite sign to that of the signal. By carefully modeling this expected phase $\Phi(t)$, we can take account of both the intrinsic frequency and spin-down rate of the neutron star and its Doppler shift. In this way the time dependence of the signal is reduced to that of the strain antenna pattern, and we are left with a relatively simple model-fitting problem to infer the unknown pulsar parameters h_0 , ι , ψ , and ϕ_0 defined in Eqs. (3.3) and (3.4).

In the time domain analysis we take a Bayesian approach and therefore express our results in terms of posterior probability distribution functions for the parameters of interest. Such PDFs are conceptually very different from those used to describe the \mathcal{F} statistic used in the frequency domain search and represent the distribution of our *degree of belief* in the values of the unknown parameters, based on the experiments and stated prior PDFs.

The time domain search algorithm comprises stages of heterodyning, noise estimation, and parameter estimation. In outline, the data are first heterodyned at a constant frequency close to the expected frequency of the signal, low-pass filtered to suppress contamination from strong signals elsewhere in the detector band, and rebinned to reduce the sampling frequency from 16384 to 4 Hz. A second (fine) heterodyne is applied to the data to account for the time-varying Doppler shift and spin down of the pulsar and any

⁷This is indeed one of the consistency checks that have been performed to validate the analysis software. We have verified that the two values of U agree within a 1% accuracy.

final instrumental calibration, and the data are rebinned to one sample per minute. We take the data as stationary during this period and make an estimate of the noise variance in each 1-min bin from the variance and covariance of the data contributing to that bin. This variance is used in the likelihood function described below.

The parameter estimation stage, at which we set the Bayesian upper limit on h_0 , proceeds from the joint probability of these 1-min complex samples, $\{B_k\}$. We take these B_k values to have a Gaussian likelihood with respect to our signal model, $y(t_k; \mathbf{a})$, where \mathbf{a} is a vector in our parameter space with components $(h_0, \iota, \psi, \phi_0)$ and t_k is the time stamp of the k th sample. The signal model, the complex heterodyne of Eq. (3.3), is

$$y(t_k; \mathbf{a}) = \frac{1}{4} F_+(t_k; \psi) h_0 (1 + \cos^2 \iota) e^{i\phi_0} - \frac{i}{2} F_\times(t_k; \psi) h_0 \cos \iota e^{i\phi_0}. \quad (4.10)$$

We choose uniform prior probabilities for ϕ_0 over $[0, 2\pi]$ and ψ over $[-\pi/4, \pi/4]$ and a prior for ι that is uniform in $\cos \iota$ over $[-1, 1]$, corresponding to a uniform probability per unit solid angle of pulsar orientation. These uniform priors are uninformative in the sense that they are invariant under changes of origin for the parameters. Although strictly a scale parameter, the prior for h_0 is also chosen as constant for $h_0 \geq 0$ and zero for $h_0 < 0$. This is a highly informative prior, in the sense that it states that the prior probability that h_0 lies between 10^{-24} and 10^{-25} is 10 times less than the prior probability it lies between 10^{-23} and 10^{-24} , but guarantees that our posterior PDF can be normalized.

The joint posterior PDF for these parameters is

$$p(\mathbf{a}|\{B_k\}) \propto p(\mathbf{a}) \exp \left[- \sum_k \frac{\mathcal{R}\{B_k - y(t_k; \mathbf{a})\}^2}{2\sigma_{\mathcal{R}\{B_k\}}^2} \right] \times \exp \left[- \sum_k \frac{\mathcal{I}\{B_k - y(t_k; \mathbf{a})\}^2}{2\sigma_{\mathcal{I}\{B_k\}}^2} \right], \quad (4.11)$$

where $p(\mathbf{a})$ ($\propto \sin \iota$) is the prior on \mathbf{a} , $\sigma_{\mathcal{R}\{B_k\}}^2$ is the variance of the real parts of B_k , and $\sigma_{\mathcal{I}\{B_k\}}^2$ is the variance of the imaginary parts of B_k .

The final stage in the analysis is to integrate this posterior PDF over the ι, ψ , and ϕ_0 parameters to give a marginalized posterior for h_0 of

$$p(h_0|\{B_k\}) \propto \iiint p(\mathbf{a}|\{B_k\}) d\iota d\psi d\phi_0, \quad (4.12)$$

normalized so that $\int_0^\infty p(h_0|\{B_k\}) dh_0 = 1$. This curve represents the distribution of our degree of belief in any particular value of h_0 , given the model of the pulsar signal, our priors for the pulsar parameters, and the data. The width of the curve roughly indicates the range in values consistent with our state of knowledge.

By definition, given our data and priors, there is a probability of 0.95 that the true value of h_0 lies below $h_0^{95\%}$ where

$$0.95 = \int_0^{h_0^{95\%}} p(h_0|\{B_k\}) dh_0, \quad (4.13)$$

and this defines our 95%-credible Bayesian upper limit on h_0 .

An attraction of this analysis is that data from different detectors can be combined directly using the appropriate signal model for each. The combined posterior distribution from all the available interferometers comes naturally out of a Bayesian analysis and, for independent observations, is simply the (normalized) product of the contributing probability distributions—i.e.,

$$p(\mathbf{a}|\text{all data}) \propto p(\mathbf{a}) \times p(\text{GEO}|\mathbf{a}) \times p(\text{H1}|\mathbf{a}) \times p(\text{H2}|\mathbf{a}) \times p(\text{L1}|\mathbf{a}). \quad (4.14)$$

This posterior PDF embodies all we believe we know about the values of the parameters, optimally combining the data from all the interferometers in a coherent way. For interferometers with very different sensitivities, this will closely approximate the result from the most sensitive instrument. Again, we must marginalize over ι, ψ , and ϕ_0 to obtain the posterior PDF for h_0 alone. We note that this is more than simply a combination of the marginalized PDFs from the separate interferometers as the coherence between the instruments is preserved, and it recognizes the different polarization sensitivities of each.

Equipment timing uncertainties due to system response delays of the order of 150 μs , constant during the run but unknown, cautioned against a coherent multi-interferometer analysis with this data set.⁸ In principle, we could assign a suitable prior for the resulting phase offsets and marginalize over them. However, the dominant position of the Livingston 4-km interferometer means that even a fully a coherent analysis would only improve our sensitivity by about 20%, so we have not pursued this. Fully coherent analyses will be possible in future observing runs.

⁸A constant (but unknown) timing offset of 150 μs at 1.3 kHz does not affect the single interferometer (IFO) coherent analysis for a 2-week observation time. For a constant time offset to matter (i.e., reduce the detection statistic by $\sim 20\%$) in the single IFO analysis, the offset must be of order 100 s or larger. This is because the detection statistic is maximized over the unknown phase ϕ_0 of the signal and the received signal is frequency modulated. The effect of a constant time offset δt is small if

$$\delta \ll \frac{10^4}{f_s} \frac{1 \text{ year}}{\min(T_{\text{obs}}, 1 \text{ year})}, \quad (4.15)$$

where f_s is the frequency of the signal and T_{obs} is the observation time (the factor 10^4 is $c/|v|$, with v being the velocity of Earth around the Sun).

2. Time domain S1 analyses for PSR J1939+2134

The time domain search used contiguous data segments 300 s or longer in duration.

The effectiveness of the noise estimation procedure described above was assessed from histograms of $B/\sigma = \mathcal{R}(B_k)/\sigma_{\mathcal{R}(B_k)} + \mathcal{J}(B_k)/\sigma_{\mathcal{J}(B_k)}$. If the estimates are correct and our likelihood function is well modeled by a Gaussian, these histograms (Fig. 7) should also be Gaussian with a variance of 1. Since we divide the noise between the real and imaginary components, we expect the value of χ^2 to be close (within $\sqrt{2N}$) to the number of real and imaginary data, N (twice the number of complex binned data values B_k). A small number of outliers with magnitudes of B_k/σ_k larger than 5 were not included in this or subsequent analyses.

The marginalized posterior PDFs for h_0 are plotted as the solid lines in Fig. 8. These represent the distribution of our degree of belief in the value of h_0 , following S1, for each interferometer. The width of each curve roughly indicates the range in values consistent with our priors and the data from the instruments individually. The formal 95% upper limits from this analysis are the upper bounds to the shaded regions in the plots and are 2.2×10^{-21} for GEO, 1.4×10^{-22} for L1, 3.3×10^{-22} for H1, and 2.4×10^{-22} for H2.

The dotted line in the GEO plot of Fig. 8 shows the (very different) marginalized posterior PDF obtained when a simulated signal is added to the data with an amplitude of 2.2×10^{-21} and with $\phi_0 = 0^\circ$, $\psi = 0^\circ$, and $\iota = 0^\circ$. Here there is a clear nonzero lower limit for the value of h_0 , and a result such as this would have indicated a nominal detection, had we seen it.

E. Estimation of uncertainties

In the frequency domain analysis the uncertainty in the upper limit value, $h_0^{95\%}$, has two contributions. The first stems from the uncertainty in the confidence ($\Delta C \approx 0.23\%$) that results from the finite sample size of the simulations. In order to convert this uncertainty into an uncertainty in $h_0^{95\%}$, we have performed several additional Monte Carlo simulations. For every run we have injected a population of signals with a given strength, h_0^{inject} , near $h_0^{95\%}$, searched for each of them with a perfectly matched template, and derived a value of \mathcal{F} . With these values we were able to estimate the $h_0(C)$ curve near $h_0^{95\%}$ and its slope h_0' and, from this, the uncertainty in the value of h_0^{inject} :

$$\Delta h_0^{95\%} \approx h_0' \Delta C. \quad (4.16)$$

The second contribution to the uncertainty in the value of $h_0^{95\%}$ comes from errors in the calibration of the instruments, which influence the absolute sensitivity scale. In particular, this reflects in an uncertainty in the actual value of the strength of injected signals so that $h_0^{95\%} = h_0^{\text{inject}} \pm \delta h_0^{\text{cal}}$. The sum of this error, estimated in Sec. II A, and the error arising from the finite sample size, Eq. (4.16), is given in the frequentist results in Table IV.

Note that when a pulsar signal is present in the data, errors in the calibration introduce errors in the phase and am-

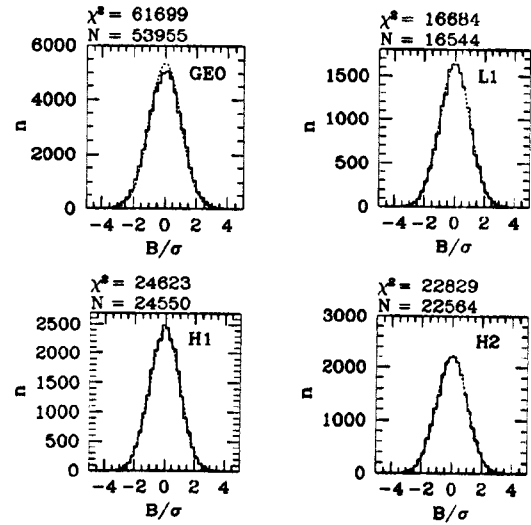


FIG. 7. Histograms of $B/\sigma = \mathcal{R}(B_k)/\sigma_{\mathcal{R}(B_k)} + \mathcal{J}(B_k)/\sigma_{\mathcal{J}(B_k)}$ for each interferometer. The dotted lines represent the expected Gaussian distribution, with $\mu=0$ and $\sigma=1$.

plitude of that signal. The errors in \mathcal{F} due to the signal are quadratic with the errors in the phase and are linear with the errors in the amplitude. However, the estimate of the noise spectral density is also affected by calibration errors and, in particular, by amplitude errors. The net effect on \mathcal{F} is that the resulting error in this quantity (which can be considered a sort of signal-to-noise ratio) is quadratic in calibration errors, thus insensitive, to first order, to calibration errors.

The errors quoted for the Bayesian results in Table IV simply reflect the calibration uncertainties given in Sec. II A. For clarity, no attempt has been made to fold a prior for this calibration factor into the marginal analysis.

V. CONCLUSION

A. Summary of results

Table IV summarizes the 95% upper limit (UL) results that we have presented in the previous sections. We should stress once more that the two analyses address two well-posed but *different* questions, and the common nomenclature is somewhat misleading.

The frequentist upper limit statements made in Sec. IV C refer to the likelihood of measuring a given value of the detection statistic or greater in repeated experiments, assuming a value for h_0 and a least-favorable orientation for the pulsar. The Bayesian limits set in Sec. IV D 1 refer to the cumulative probability of the value of h_0 itself given the data and prior beliefs in the parameter values. The Bayesian upper limits report intervals in which we are 95% certain that the true value resides. We do not expect two such distinct definitions of “upper limit” to yield the same numerical value.

Recall that the frequentist UL is *conservative*: it is calculated for the worst-case values of signal parameters ι and ψ . The Bayesian TDS method marginalizes over these pa-




FIG. 8. For each interferometer, the solid line represents the marginalized posterior PDF for h_0 (PSR J1939+2134) resulting from the S1 data. The 95% upper limits (extent of the shaded region) are 2.2×10^{-21} for GEO, 1.4×10^{-22} for L1, 3.3×10^{-22} for H1, and 2.4×10^{-22} for H2. The dotted line in the GEO plot shows the posterior PDF of h_0 in the presence of a simulated signal injected into the GEO S1 data stream using $h_0 = 2.2 \times 10^{-21}$, $\phi_0 = 0^\circ$, $\psi = 0^\circ$, and $\epsilon = 0^\circ$.

rameters, gathering together the evidence supporting a particular h_0 *irrespective of orientation*. We have also performed an alternative calculation of the frequentist ULs by using a $p(\mathcal{F}|h_0)$ derived from a population of signals with $\cos \iota$ and ψ parameters uniformly distributed, as were the Bayesian priors in the time domain search. As expected, we find that the resulting ULs have somewhat lower values than the conservative ones reported in Table IV: 1.2×10^{-21} , 1.5×10^{-22} , 4.5×10^{-22} , and 2.3×10^{-22} for the GEO, L1, H1, and H2 data sets, respectively.

Note that a conservative UL in one scheme (Bayesian or frequentist) should not be expected to always produce a higher number than an average or optimistic UL in the other scheme. In particular, when \mathcal{F}^* is fairly low (as in the GEO case), it is reasonable for the frequentist conservative UL to actually be lower than the Bayesian UL [39], as we see in the

TABLE IV. Summary of the 95% upper limit values of h_0 for PSR J1939+2134. The frequency domain search (FDS) quotes a conservative frequentist upper limit and the time domain search (TDS) a Bayesian upper limit after marginalizing over the unknown ι , ψ , and ϕ_0 parameters.

IFO	Frequentist FDS	Bayesian TDS
GEO	$(1.9 \pm 0.1) \times 10^{-21}$	$(2.2 \pm 0.1) \times 10^{-21}$
L1	$(2.7 \pm 0.3) \times 10^{-22}$	$(1.4 \pm 0.1) \times 10^{-22}$
H1	$(5.4 \pm 0.6) \times 10^{-22}$	$(3.3 \pm 0.3) \times 10^{-22}$
H2	$(4.0 \pm 0.5) \times 10^{-22}$	$(2.4 \pm 0.2) \times 10^{-22}$

first line of Table IV. Conversely, the large value of \mathcal{F}^* for H1 translates into a relatively large ratio of the frequentist “average” UL to the Bayesian one.

B. Discussion of previous upper limit results

Two prior upper limits have been published on the strain of a signal from our specific pulsar J1939+2134. A limit of $h < 3.1 \times 10^{-17}$ and 1.5×10^{-17} for the first and second harmonics of the rotation frequency of the pulsar, respectively, was set in [40] using 4 d of data from the Caltech 40-m interferometer. A tighter limit $h < 10^{-20}$ was determined using a divided-bar gravitational wave detector at Glasgow University for the second harmonic alone [41].

More sensitive untargated UL results on the strain of periodic GW signals at other frequencies come from acoustic bar detector experiments [42,43,44]. As a consequence of the narrow sensitivity bands of these detectors (less than 1 Hz around each mode) and the fact that their frequencies do not correspond to those of any known pulsars,⁹ studies with bar antennas have not investigated possible emission from any known pulsars.

In [42] a UL of 2.9×10^{-24} was reported for periodic signals from the Galactic center, with $921.32 < f_s < 921.38$ Hz and no appreciable spin down over ~ 95.7 days of observation. These data were collected by the EXPLORER detector in 1991. This UL result was *not* obtained by a coherent search over the entire observation time, due to insufficient timing accuracy.

In [43] a fully coherent 2-day-long all-sky search was performed again on 1991 EXPLORER data in a f_s search band of about 1 Hz centered at 922 Hz and including one spin-down parameter. It resulted in an UL of 2.8×10^{-23} at the 99% confidence level. This search was based on the same detection statistic used in our frequency domain analysis.

Another parameter space search is presented in [44]. Data taken from the ALLEGRO detector during the first three months of 1994 were searched for periodic gravitational wave signals from the Galactic center and from the globular cluster 47Tuc, with no resolvable spin down and with f_s in the two sensitive bands of their antenna, 896.30–897.30 Hz and 919.76–920.76 Hz, with a 10- μ Hz resolution. The resulting UL at 8×10^{-24} is reported.

There exist several results from searches using early broadband interferometric detectors [40,41,46–49]. As a result of the poor sensitivities of these early detector prototypes, none of these upper limits is competitive with the strain sensitivity achieved here. However, many of the new issues and complications associated with broadband search instruments were first confronted in these early papers, laying the foundations for future analyses.

Data from the first science run of the TAMA detector were searched for continuous waves from SN1987A in a 0.05-Hz

⁹With the exception of the Australian detector NIOBE and of the Japanese torsional antenna built specifically to detect periodic signals from the Crab pulsar [45].

band at ~ 934.9 Hz. The reported 99% confidence upper limit was $\sim 5 \times 10^{-23}$ [50].

Improved noise performance and longer observation times achieved with interferometric detectors since S1 has made their sensitivities comparable to or better than the narrow band peak sensitivity of the acoustic bars cited above, over much broader bandwidths. Combined with the advances in analysis methods presented in this paper, we anticipate significant advances in search depth and breadth in the next set of observations.

C. Upper limit on the ellipticity of the pulsar

An UL on h_0 for J1939+2134 can be interpreted as an UL on the neutron star's equatorial ellipticity. Taking the distance to J1939+2134 to be 3.6 kpc, Eq. (3.1) gives an UL on the ellipticity corresponding to $h_0^{95\%} = 1.4 \times 10^{-22}$ of

$$\epsilon^{95\%} = 2.9 \times 10^{-4} \left(\frac{10^{45} \text{ g cm}^2}{I_{zz}} \right). \quad (5.1)$$

Of course, the UL on the ellipticity of J1939+2134 derived from S1 data is about five orders of magnitude higher than the UL obtained from the pulsar measured spin-down rate: $\epsilon \leq 3.80 \times 10^{-9} (10^{45} \text{ g cm}^2 / I_{zz})^{1/2}$. However, an ellipticity of $\sim 10^{-4}$ could in principle be generated by an

interior magnetic field of strength $\sim 10^{16}$ G or it could probably be sustained in a NS with a solid core. Therefore, the above exercise suggests that with improved detector sensitivities, even a null result from a search for *unknown* pulsars will place interesting constraints on the ellipticities of rapidly rotating neutron stars that might exist in our galactic neighborhood.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the U.S. National Science Foundation for the construction and operation of the LIGO Laboratory and the Particle Physics and Astronomy Research Council of the United Kingdom, the Max-Planck-Society, and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Natural Sciences and Engineering Research Council of Canada, the Council of Scientific and Industrial Research of India, the Department of Science and Technology of India, the Spanish Ministerio de Ciencia y Tecnología, the John Simon Guggenheim Foundation, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P. Sloan Foundation.

-
- [1] LIGO Scientific Collaboration, B. Abbott *et al.*, "Detector description and performance for the First Coincidence Observations between LIGO and GEO," Nucl. Instrum. Methods Phys. Res. A **517**, 154 (2004).
 - [2] A. Abramovici *et al.*, Science **256**, 325 (1992).
 - [3] B. Barish and R. Weiss, Phys. Today **52**(10), 44 (1999).
 - [4] B. Willke *et al.*, Class. Quantum Grav. **19**, 1377 (2002).
 - [5] S. Goßler *et al.*, Class. Quantum Grav. **19**, 1835 (2002).
 - [6] A. Einstein, Preuss. Akad. Wiss. Berlin, Sitzber 844 (1915).
 - [7] A. Einstein, Preuss. Akad. Wiss. Berlin, Sitzber 688 (1916).
 - [8] B. Caron *et al.*, Nucl. Phys. B (Proc. Suppl.) **54**, 167 (1997).
 - [9] K. Tsubono, in *Proceedings of the 1st Edoardo Amaldi Conference Gravitational Wave Experiments*, edited by E. Coccia, G. Pizella, and F. Ronga (World Scientific, Singapore, 1995), p. 112.
 - [10] B. Allen, M. A. Papa, and B. F. Schutz, Phys. Rev. D **66**, 102003 (2002).
 - [11] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
 - [12] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
 - [13] P. Fritschel, Proc. SPIE **4856**, 39 (2002).
 - [14] L. Bildsten, Astrophys. J. Lett. **501**, L89 (1998).
 - [15] L. Bildsten, in *Radio Pulsars*, edited by M. Bailes, D. J. Nice, and S. E. Thorsett, ASP Conference Series (The University of Chicago Press, Chicago, 2003), astro-ph/0212004.
 - [16] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. G. Galloway, R. Wijnands, M. van der Klis, and C. Markwardt, Nature (London) **424**, 42 (2003).
 - [17] R. Wijnands, M. van der Klis, J. Homan, D. Chakrabarty, C. B. Markwardt, and E. H. Morgan, Nature (London) **424**, 44 (2003).
 - [18] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).
 - [19] R. V. Wagoner, Nature (London) **424**, 27 (2003).
 - [20] F. Verbunt, Annu. Rev. Astron. Astrophys. **31**, 93 (1993).
 - [21] M. van der Klis, Annu. Rev. Astron. Astrophys. **38**, 717 (2000).
 - [22] D. I. Jones and N. Andersson, Mon. Not. R. Astron. Soc. **331**, 203 (2002).
 - [23] R. Jenet and S. B. Anderson, Astrophys. J. **546**, 394 (2001).
 - [24] A. N. Lommen and D. C. Backer (personal communication).
 - [25] N. Andersson, Astrophys. J. **502**, 708 (1998).
 - [26] J. L. Friedman and S. M. Morsink, Astrophys. J. **502**, 714 (1998).
 - [27] P. Arras, E. E. Flanagan, S. M. Morsink, A. K. Schenk, S. A. Teukolsky, and I. Wasserman, Astrophys. J. **591**, 1129 (2003).
 - [28] G. Ushomirsky, C. Cutler, and L. Bildsten, Mon. Not. R. Astron. Soc. **319**, 902 (2000).
 - [29] C. Cutler, Phys. Rev. Lett. **66**, 084025 (2002).
 - [30] J. H. Taylor, Philos. Trans. R. Soc. London **A341**, 117 (1992).
 - [31] <http://pulsar.princeton.edu/tempo/index.html>
 - [32] V. M. Kaspi, J. H. Taylor, and M. F. Ryba, Astrophys. J. **428**, 713 (1994).
 - [33] A. O'Hagan, *Kendall's Advanced Theory of Statistic* (Halsted Press, New York, 1994), Vol. 2B.
 - [34] L. S. Finn, in *Proceedings of the Second Edoardo Amaldi Conference on Gravitational Waves*, edited by E. Coccia, G. Veneziano, and G. Pizzella (World Scientific, Singapore, 1998), p. 180.
 - [35] <http://www.lsc-group.phys.uwm.edu/lal/>.

- [36] <http://www.cs.wisc.edu/condor/> "Condor is a specialized workload management system for compute-intensive jobs. Like other full-featured batch systems, Condor provides a job queueing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their serial or parallel jobs to Condor, Condor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion."
- [37] <http://www.lsc-group.phys.uwm.edu/beowulf/medusa/index.html>
- [38] <http://pandora.aei.mpg.de/merlin>
- [39] Y. Itoh, LIGO technical document No. LIGO-T030248-00-Z.
- [40] M. Hereld, Ph.D. thesis, California Institute of Technology, 1983.
- [41] J. Hough *et al.*, *Nature (London)* **303**, 216 (1983).
- [42] P. Astone *et al.*, *Phys. Rev. D* **65**, 022001 (2002).
- [43] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Krolak, *Phys. Rev. D* **65**, 042003 (2002).
- [44] E. Mauceli, M. P. McHugh, W. O. Hamilton, W. W. Johnson, and A. Morse, gr-qc/0007023.
- [45] T. Suzuki, in *Gravitational Wave Experiments*, edited by E. Coccia, G. Pizzella, and F. Ronga (World Scientific, Singapore, 1995), p. 115.
- [46] J. Livas, in *Gravitational Wave Data Analysis*, edited by Schutz (NATO ASI Series C) (Plenum, New York, 1988), Vol. 253, p. 217; Ph.D. thesis, Massachusetts Institute of Technology, 1987.
- [47] T. M. Niebauer *et al.*, *Phys. Rev. D* **47**, 3106 (1993).
- [48] G. S. Jones, Ph.D. thesis, University of Wales (Cardiff University), 1995.
- [49] M. Zucker, Ph.D. thesis, California Institute of Technology, 1988.
- [50] K. Sorda, M. Ando, N. Kanda, H. Tagoshi, D. Tatsumi, K. Tsubono, and the TAMA Collaboration, *Class. Quantum Grav.* **20**, S645 (2003).

Limits on gravitational wave emission from selected pulsars using LIGO data

- B. Abbott,¹² R. Abbott,¹⁵ R. Adhikari,¹³ A. Ageev,^{20,27} B. Allen,³⁹ R. Amin,³⁴ S. B. Anderson,¹² W. G. Anderson,²⁹ M. Araya,¹² H. Armandula,¹² M. Ashley,²⁸ F. Asiri,^{12, a} P. Aufmuth,³¹ C. Aulbert,¹ S. Babak,⁷ R. Balasubramanian,⁷ S. Ballmer,¹³ B. C. Barish,¹² C. Barker,¹⁴ D. Barker,¹⁴ M. Barnes,^{12, b} B. Barr,³⁵ M. A. Barton,¹² K. Bayer,¹³ R. Beausoleil,^{26, c} K. Belczynski,²³ R. Bennett,^{35, d} S. J. Berukoff,^{1, e} J. Betzwieser,¹³ B. Bhawal,¹² I. A. Bilenko,²⁰ G. Billingsley,¹² E. Black,¹² K. Blackburn,¹² L. Blackburn,¹³ B. Bland,¹⁴ B. Bochner,^{13, f} L. Bogue,¹² R. Bork,¹² S. Bose,⁴⁰ P. R. Brady,³⁹ V. B. Braginsky,²⁰ J. E. Brau,³⁷ D. A. Brown,³⁹ A. Bullington,²⁶ A. Bunkowski,^{2, 31} A. Buonanno,^{6, g} R. Burgess,¹³ D. Busby,¹² W. E. Butler,³⁸ R. L. Byer,²⁶ L. Cadonati,¹³ G. Cagnoli,³⁵ J. B. Camp,²¹ C. A. Cantley,³⁵ L. Cardenas,¹² K. Carter,¹⁵ M. M. Casey,³⁵ J. Castiglione,³⁴ A. Chandler,¹² J. Chapsky,^{12, b} P. Charlton,¹² S. Chatterji,¹³ S. Chelkowski,^{2, 31} Y. Chen,⁶ V. Chickarmane,^{16, h} D. Chin,³⁶ N. Christensen,⁸ D. Churches,⁷ T. Cokelaer,⁷ C. Colacino,⁴³ R. Coldwell,³⁴ M. Coles,^{15, i} D. Cook,¹⁴ T. Corbitt,¹² D. Coyne,¹² J. D. E. Creighton,³⁹ T. D. Creighton,¹² D. R. M. Crooks,³⁵ P. Csatorday,¹³ B. J. Cusack,³ C. Cutler,¹ E. D'Ambrosio,¹² K. Danzmann,^{31, 2} E. Daw,^{16, j} D. DeBra,²⁶ T. Delker,^{34, k} V. Dergachev,³⁶ R. DeSalvo,¹² S. Dhurandhar,¹¹ A. Di Credico,²⁷ M. Díaz,²⁹ H. Ding,¹² R. W. P. Drever,⁴ R. J. Dupuis,³⁵ J. A. Edlund,^{12, b} P. Ehrens,¹² E. J. Elliffe,³⁵ T. Etzel,¹² M. Evans,¹² T. Evans,¹⁵ S. Fairhurst,³⁹ C. Fallnich,³¹ D. Farnham,¹² M. M. Fejer,²⁶ T. Findley,²⁵ M. Fine,¹² L. S. Finn,²⁸ K. Y. Franzen,³⁴ A. Freise,^{2, l} R. Frey,³⁷ P. Fritschel,¹³ V. V. Frolov,¹⁵ M. Fyfe,¹⁵ K. S. Ganezer,⁵ J. Garofoli,¹⁴ J. A. Giaime,¹⁶ A. Gillespie,^{12, m} K. Goda,¹³ G. González,¹⁶ S. Goßler,³¹ P. Grandclément,^{23, n} A. Grant,³⁵ C. Gray,¹⁴ A. M. Gretarsson,¹⁵ D. Grimmer,¹² H. Grote,² S. Grunewald,¹ M. Guenther,¹⁴ E. Gustafson,^{26, o} R. Gustafson,³⁶ W. O. Hamilton,¹⁶ M. Hammond,¹⁵ J. Hanson,¹⁵ C. Hardham,²⁶ J. Harms,¹⁹ G. Harry,¹³ A. Hartunian,¹² J. Heefner,¹² Y. Hefetz,¹³ G. Heinzel,² I. S. Heng,³¹ M. Hennessy,²⁶ N. Hepler,²⁸ A. Heptonstall,³⁵ M. Heurs,³¹ M. Hewitson,² S. Hild,² N. Hindman,¹⁴ P. Hoang,¹² J. Hough,³⁵ M. Hrynevych,^{12, p} W. Hua,²⁶ M. Ito,³⁷ Y. Itoh,¹ A. Ivanov,¹² O. Jennrich,^{35, q} B. Johnson,¹⁴ W. W. Johnson,¹⁶ W. R. Johnston,²⁹ D. I. Jones,²⁸ L. Jones,¹² D. Jungwirth,^{12, r} V. Kalogera,²³ E. Katsavounidis,¹³ K. Kawabe,¹⁴ S. Kawamura,²² W. Kells,¹² J. Kern,^{15, s} A. Khan,¹⁵ S. Killbourn,³⁵ C. J. Killow,³⁵ C. Kim,²³ C. King,¹² P. King,¹² S. Klimenko,¹⁴ S. Koranda,³⁹ K. Kötter,³¹ J. Kovalik,^{15, b} D. Kozak,¹² B. Krishnan,¹ M. Landry,¹⁴ J. Langdale,¹⁵ B. Lantz,²⁶ R. Lawrence,¹³ A. Lazzarini,¹² M. Lei,¹² I. Leonor,³⁷ K. Libbrecht,¹² A. Libson,⁸ P. Lindquist,¹² S. Liu,¹² J. Logan,^{12, t} M. Lormand,¹⁵ M. Lubinski,¹⁴ H. Lück,^{31, 2} T. T. Lyons,^{12, u} B. Machenschalk,¹ M. MacInnis,¹³ M. Mageswaran,¹² K. Mailand,¹² W. Majid,^{12, b} M. Malec,^{2, 31} F. Mann,¹² A. Marin,^{13, u} S. Márka,¹² E. Maros,¹² J. Mason,^{12, v} K. Mason,¹³ O. Matherny,¹⁴ L. Matone,¹⁴ N. Mavalvala,¹³ R. McCarthy,¹⁴ D. E. McClelland,³ M. McHugh,¹⁸ J. W. C. McNabb,²⁸ G. Mendell,¹⁴ R. A. Mercer,³³ S. Meshkov,¹² E. Messaritaki,³⁹ C. Messenger,³³ V. P. Mitrofanov,²⁰ G. Mitselmakher,³⁴ R. Mittleman,¹³ O. Miyakawa,¹² S. Miyoki,^{12, w} S. Mohanty,²⁹ G. Moreno,¹⁴ K. Mossavi,² G. Mueller,³⁴ S. Mukherjee,²⁹ P. Murray,³⁵ J. Myers,¹⁴ S. Nagano,² T. Nash,¹² R. Nayak,¹² G. Newton,³⁵ F. Nocera,¹² J. S. Noel,⁴⁰ P. Nutzman,²³ T. Olson,²⁴ B. O'Reilly,¹⁵ D. J. Ottaway,¹³ A. Ottewill,^{39, x} D. Ouimette,^{12, y} H. Overmier,¹⁵ B. J. Owen,²⁸ Y. Pan,⁶ M. A. Papa,¹ V. Parameshwariah,¹⁴ C. Parameshwariah,¹⁵ M. Pedraza,¹² S. Penn,¹⁰ M. Pitkin,³⁵ M. Plissi,³⁵ R. Prix,¹ V. Quetschke,³⁴ F. Raab,¹⁴ H. Radkins,¹⁴ R. Rahlkola,³⁷ M. Rakhmanov,³⁴ S. R. Rao,¹² K. Rawlins,¹³ S. Ray-Majumder,³⁹ V. Re,³³ D. Redding,^{12, b} M. W. Regehr,^{12, b} T. Regimbau,⁷ S. Reid,³⁵ K. T. Reilly,¹² K. Reithmaier,¹² D. H. Reitze,³⁴ S. Richman,^{13, y} R. Riesen,¹⁵ K. Riles,³⁶ B. Rivera,¹⁴ A. Rizzi,^{15, z} D. I. Robertson,³⁵ N. A. Robertson,^{26, 35} L. Robison,¹² S. Roddy,¹⁵ J. Rollins,¹³ J. D. Romano,⁷ J. Romie,¹² H. Rong,^{34, m} D. Rose,¹² E. Rothhoff,²⁸ S. Rowan,³⁵ A. Rüdiger,² P. Russell,¹² K. Ryan,¹⁴ I. Salzman,¹² V. Sandberg,¹⁴ G. H. Sanders,^{12, aa} V. Sannibale,¹² B. Sathyaprakash,⁷ P. R. Saulson,²⁷ R. Savage,¹⁴ A. Sazonov,³⁴ R. Schilling,² K. Schlaufman,²⁸ V. Schmidt,^{12, bb} R. Schnabel,¹⁹ R. Schofield,³⁷ B. F. Schutz,^{1, 7} P. Schwinberg,¹⁴ S. M. Scott,³ S. E. Seader,⁴⁰ A. C. Searle,³ B. Sears,¹² S. Seel,¹² F. Seifert,¹⁹ A. S. Sengupta,¹¹ C. A. Shapiro,^{28, cc} P. Shawhan,¹² D. H. Shoemaker,¹³ Q. Z. Shu,^{34, dd} A. Sibley,¹⁵ X. Siemens,³⁹ L. Sievers,^{12, b} D. Sigg,¹⁴ A. M. Sintes,^{1, 32} J. R. Smith,² M. Smith,¹³ M. R. Smith,¹² P. H. Sneddon,³⁵ R. Spero,^{12, b} G. Stapfer,¹⁵ D. Steussy,⁸ K. A. Strain,³⁵ D. Strom,³⁷ A. Stuver,²⁸ T. Summerscales,²⁸ M. C. Sumner,¹² P. J. Sutton,¹² J. Sylvestre,^{12, ee} A. Takamori,¹² D. B. Tanner,³⁴ H. Tariq,¹² I. Taylor,⁷ R. Taylor,¹² R. Taylor,³⁵ K. A. Thorne,²⁸ K. S. Thorne,⁶ M. Tibbits,²⁸ S. Tilav,^{12, ff} M. Tinto,^{4, b} K. V. Tokmakov,²⁰ C. Torres,²⁹ C. Torrie,¹² G. Traylor,¹⁵ W. Tyler,¹² D. Ugolini,³⁰ C. Ungarelli,³³ M. Vallisneri,^{6, gg} M. van Putten,¹³ S. Vass,¹² A. Vecchio,³³ J. Veitch,³⁵ C. Vorvick,¹⁴ S. P. Vyachanin,²⁰ L. Wallace,¹² H. Walther,¹⁹ H. Ward,³⁵ B. Ware,^{12, b} K. Watts,¹⁵ D. Webber,¹² A. Weidner,¹⁹ U. Weiland,³¹ A. Weinstein,¹² R. Weiss,¹³ H. Welling,³¹ L. Wen,¹² S. Wen,¹⁶ J. T. Whelan,¹⁸ S. E. Whitcomb,¹² B. F. Whiting,³⁴ S. Wiley,⁵ C. Wilkinson,¹⁴ P. A. Willens,¹² P. R. Williams,^{1, hh} R. Williams,⁴ B. Willke,³¹ A. Wilson,¹² B. J. Winjum,^{28, e} W. Winkler,² S. Wise,³⁴ A. G. Wiseman,³⁹ G. Woan,³⁵ R. Wooley,¹⁵ J. Worden,¹⁴ W. Wu,³⁴ I. Yakushin,¹⁵ H. Yamamoto,¹² S. Yoshida,²⁵

K. D. Zaleski,²⁸ M. Zanolin,¹³ I. Zawischa,^{31, ii} L. Zhang,¹² R. Zhu,¹ N. Zotov,¹⁷ M. Zucker,¹⁵ and J. Zwiizig¹²
(The LIGO Scientific Collaboration, <http://www.ligo.org>)

M. Kramer⁴¹ and A. G. Lyne⁴¹

- ¹ Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Golm, Germany
² Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany
³ Australian National University, Canberra, 0200, Australia
⁴ California Institute of Technology, Pasadena, CA 91125, USA
⁵ California State University Dominguez Hills, Carson, CA 90747, USA
⁶ Caltech-CaRT, Pasadena, CA 91125, USA
⁷ Cardiff University, Cardiff, CF2 3YB, United Kingdom
⁸ Carleton College, Northfield, MN 55057, USA
⁹ Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
¹⁰ Hobart and William Smith Colleges, Geneva, NY 14456, USA
¹¹ Inter-University Centre for Astronomy and Astrophysics, Pune - 411007, India
¹² LIGO - California Institute of Technology, Pasadena, CA 91125, USA
¹³ LIGO - Massachusetts Institute of Technology, Cambridge, MA 02139, USA
¹⁴ LIGO Hanford Observatory, Richland, WA 99352, USA
¹⁵ LIGO Livingston Observatory, Livingston, LA 70754, USA
¹⁶ Louisiana State University, Baton Rouge, LA 70803, USA
¹⁷ Louisiana Tech University, Ruston, LA 71272, USA
¹⁸ Loyola University, New Orleans, LA 70118, USA
¹⁹ Max Planck Institut für Quantenoptik, D-85748, Garching, Germany
²⁰ Moscow State University, Moscow, 119992, Russia
²¹ NASA/Goddard Space Flight Center, Greenbelt, MD 20771, USA
²² National Astronomical Observatory of Japan, Tokyo 181-8588, Japan
²³ Northwestern University, Evanston, IL 60208, USA
²⁴ Salish Kootenai College, Pablo, MT 59855, USA
²⁵ Southeastern Louisiana University, Hammond, LA 70402, USA
²⁶ Stanford University, Stanford, CA 94305, USA
²⁷ Syracuse University, Syracuse, NY 13244, USA
²⁸ The Pennsylvania State University, University Park, PA 16802, USA
²⁹ The University of Texas at Brownsville and Texas Southmost College, Brownsville, TX 78520, USA
³⁰ Trinity University, San Antonio, TX 78212, USA
³¹ Universität Hannover, D-30167 Hannover, Germany
³² Universitat de les Illes Balears, E-07122 Palma de Mallorca, Spain
³³ University of Birmingham, Birmingham, B15 2TT, United Kingdom
³⁴ University of Florida, Gainesville, FL 32611, USA
³⁵ University of Glasgow, Glasgow, G12 8QQ, United Kingdom
³⁶ University of Michigan, Ann Arbor, MI 48109, USA
³⁷ University of Oregon, Eugene, OR 97403, USA
³⁸ University of Rochester, Rochester, NY 14627, USA
³⁹ University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA
⁴⁰ Washington State University, Pullman, WA 99164, USA
⁴¹ University of Manchester, Jodrell Bank Observatory, Macclesfield, Cheshire, SK11 9DL, United Kingdom

(Dated: October 5, 2004)

We place direct upper limits on the strain of the gravitational waves from 28 isolated radio pulsars by a coherent multi-detector analysis of the data collected during the second science run of the LIGO interferometric detectors. These are the first *direct* upper limits for 26 of the 28 pulsars. We use coordinated radio observations for the first time to build radio-guided phase templates for the expected gravitational wave signals. The unprecedented sensitivity of the detectors allow us to set strain upper limits as low as a few times 10^{-24} . These strain limits translate into limits on the equatorial ellipticities of the pulsars, which are smaller than 10^{-5} for the four closest pulsars.

PACS numbers: 04.80.Nn, 95.55.Ym, 97.60.Gb, 07.05.Kf

^aCurrently at Stanford Linear Accelerator Center

^bCurrently at Jet Propulsion Laboratory

^cPermanent Address: HP Laboratories

^dCurrently at Rutherford Appleton Laboratory

^eCurrently at University of California, Los Angeles

^fCurrently at Hofstra University

A worldwide effort is underway to detect gravitational waves (GWs) and thus test a fundamental prediction of General Relativity. In preparation for long-term operations, the LIGO and GEO experiments conducted their first science run (S1) during 17 days in 2002. The detectors and the analyses of the S1 data are described in Refs. [1] and [2]-[5], respectively. LIGO's second science run (S2) was carried out from 14 Feb - 14 April 2003, with dramatically improved sensitivity compared to S1. During S2 the GEO detector was not operating.

A spinning neutron star is expected to emit GWs if it is not perfectly symmetric about its rotation axis. The strain amplitude h_0 of the emitted signal is proportional to the star's deformation as measured by its ellipticity ϵ [6]. Using data from S2, this paper reports *direct* observational limits on the GW emission and corresponding ellipticities from the 28 most rapidly rotating isolated pulsars for which radio data is complete enough to guide the phase of our filters with sufficient precision. These are the first such limits for 26 of the pulsars. We concentrate on isolated pulsars with known phase evolutions and sky positions to ensure that our targeted search requires relatively few unknown parameters.

The limits reported here are still well above the indirect limits inferred from observed pulsar spindown, where available (Fig. 1). However, fourteen of our pulsars are in globular clusters, where local gravitational accelerations produce Doppler effects that mask the true pulsar spindown, sometimes even producing apparent spinup. For

these pulsars our observations therefore place the first limits that are inherently independent of cluster dynamics, albeit at levels well above what one would expect if all globular cluster pulsars are similar to field pulsars.

Our most stringent ellipticity upper limit is 4.5×10^{-6} . While still above the maximum expected from conventional models of nuclear matter, distortions of this size would be permitted within at least one exotic theory of neutron star structure [7].

Detectors.—The LIGO observatory is composed of three detectors. Each detector is a power-recycled Michelson interferometer, with Fabry-Perot cavities in the long arms. A passing GW produces a time-varying differential strain in these arms, and the resulting differential displacement of the cavity test mass mirrors is sensed interferometrically. Two detectors, the 4 km arm H1 and the 2 km arm H2 detectors, are collocated in Hanford WA. The 4 km arm L1 detector is situated in Livingston Parish LA. Improvements in noise performance between S1 and S2 were approximately an order of magnitude over a broad frequency range. Modifications that were made between S1 and S2 to aid in noise reduction and improve stability include i) increased laser power to reduce high-frequency noise, ii) better angular control of the mirrors of the interferometer and iii) the use of lower noise digital test mass suspension controllers in all detectors.

During S2, the LIGO detectors' noise performance in the band 40-2000 Hz was better than any previous interferometer. The best strain sensitivity, achieved by L1, was $\sim 3 \times 10^{-22} \text{ Hz}^{-1/2}$ near 200 Hz (Fig. 1). The relative timing stability between the interferometers was also significantly improved. Monitored with GPS-synchronized clocks to be better than $10 \mu\text{s}$ over S2, it allowed the coherent combination of the strain data of all three detectors to form joint upper limits.

Analysis method.—In [2] a search for gravitational waves from the millisecond pulsar PSR J1939+2134 using S1 data was presented. In that work, two different data analysis methods were used, one in the time domain and the other in the frequency domain. Here we extend the former method [8] and apply it to 28 isolated pulsars.

Following [2] we model the sources as non-processing triaxial neutron stars showing the same rotational phase evolution as is present in the radio signal and perform a complex heterodyne of the strain data from each detector at the instantaneous frequency of the expected gravitational wave signal, which is twice the observed radio rotation frequency. These data are then downsampled to 1/60 Hz and will be referred to as B_k . Any gravitational signal in the data would show a residual time evolution reflecting the antenna pattern of the detector, varying over the day as the source moved through the pattern, but with a functional form that depended on several other source-observer parameters: the antenna responses to plus and cross polarisations, the amplitude of the gravitational wave h_0 , the angle between the line-of-sight to the pulsar and its spin axis ι , the polarisation angle of the gravitational radiation ψ (all defined in [6])

*Permanent Address: GReCO, Institut d'Astrophysique de Paris (CNRS)

^bCurrently at Keck Graduate Institute

^cCurrently at National Science Foundation

^dCurrently at University of Sheffield

^eCurrently at Ball Aerospace Corporation

^fCurrently at European Gravitational Observatory

^gCurrently at Intel Corp.

^hCurrently at University of Tours, France

ⁱCurrently at Lightconect Inc.

^jCurrently at W.M. Keck Observatory

^kCurrently at ESA Science and Technology Center

^lCurrently at Raytheon Corporation

^mCurrently at New Mexico Institute of Mining and Technology / Magdalena Ridge Observatory Interferometer

ⁿCurrently at Mission Research Corporation

^oCurrently at Harvard University

^pCurrently at Lockheed-Martin Corporation

^qPermanent Address: University of Tokyo, Institute for Cosmic Ray Research

^rPermanent Address: University College Dublin

^sCurrently at Research Electro-Optics Inc.

^tCurrently at Institute of Advanced Physics, Baton Rouge, LA

^uCurrently at Thirty Meter Telescope Project at Caltech

^vCurrently at European Commission, DG Research, Brussels, Belgium

^wCurrently at University of Chicago

^xCurrently at LightBit Corporation

^yPermanent Address: IBM Canada Ltd.

^zCurrently at University of Delaware

^{aa}Permanent Address: Jet Propulsion Laboratory

^{ab}Currently at Shanghai Astronomical Observatory

^{ac}Currently at Laser Zentrum Hannover




FIG. 1: Upper curves: characteristic amplitudes detectable from a known source with a 1% false alarm rate and 10% false dismissal rate, as given by Eq. (2.2) in [2], using S2 sensitivities and observation times. Lower curve: LIGO design sensitivity for 1 yr of data. Stars: upper limits found in this paper for 28 known pulsars. Circles: spindown upper limits for the pulsars with negative spindown values if *all* the measured loss of angular momentum were due to gravitational waves and assuming a moment of inertia of $10^{45} \text{ g} \cdot \text{cm}^2$.

and the phase ϕ_0 of the gravitational wave signal at some fiducial time t_0 . Let \mathbf{a} be a vector in parameter space with components $(h_0, \iota, \psi, \phi_0)$.

The analysis proceeds by determining the posterior probability distribution function (pdf) of \mathbf{a} given the data B_k and the signal model:

$$p(\mathbf{a}|\{B_k\}) \propto p(\mathbf{a})p(\{B_k\}|\mathbf{a}), \quad (1)$$

where $p(\{B_k\}|\mathbf{a})$ is the likelihood and $p(\mathbf{a})$ the prior pdf we assign to the model parameters. We have used a uniform prior for $\cos \iota$, ϕ_0 , ψ and h_0 ($h_0 > 0$), in common with [2]. A uniform prior for h_0 has been chosen for its simplicity and so that our results can readily be compared with other observations. This prior favors high values of h_0 (which comprise the majority of the parameter space) and therefore generates a somewhat conservative upper limit for its value. Indeed the reader might prefer to regard our resulting posterior pdfs for h_0 as marginalised likelihoods rather than probabilities for h_0 — these would be functionally identical using our priors.

As in [2] we use a Gaussian joint likelihood for $p(\{B_k\}|\mathbf{a})$. In [2] the S1 noise floor was estimated over a 60s period from a 4 Hz band about the expected signal frequency. This gave a reliable point estimate for the noise level but was sensitive to spectral contamination within the band, as demonstrated in the analysis of the GEO S1 data. In this paper we exploit the improved stationarity of the instruments and take the noise floor to be constant over periods of 30 min. In addition we restrict the bandwidth to 1/60 Hz, which makes it possible to search for signals from pulsars at frequencies close




FIG. 2: Parameters of the artificial pulsar P1, recovered from 12h of strain data from the Hanford and Livingston interferometers. The results are displayed as marginal pdfs for each of the four signal parameters. The vertical dotted lines show the values used to generate the signal, the colored lines show the results from the individual detectors (H1 green, H2 blue, L1 red), and the black lines show the joint result from combining coherently data from the three.

to strong spectral disturbances. However, the noise level now determined is less certain as the estimate relies on fewer data. We take account of this increased uncertainty by explicitly marginalising with a Jeffreys prior over the constant but unknown noise level for each 30 min period of data [9]. The likelihood for this analysis is then the combined likelihood for all the 30 min stretches of data, labeled by j in Eq. (2), taken as independent:

$$p(\{B_k\}|\mathbf{a}) \propto \prod_j p(\{B_k\}_j|\mathbf{a}), \quad (2)$$

$$p(\{B_k\}_j|\mathbf{a}) \propto \left(\sum_{k=k_{1(j)}}^{k_{2(j)}} |B_k - y_k|^2 \right)^{-m} \quad (3)$$

where y_k is the signal model given by Eq. (4.10) in [2] and $m = k_{2(j)} - k_{1(j)} + 1 = 30$ is the number of B_k data points in a 30 min segment.

In principle the period over which the data are assumed stationary need not be fixed, and can be adjusted dynamically to reflect instrumental performance over the run. We have limited our analysis to continuous 30 min stretches of data, which included more than 88% of the S2 science data set. Inclusion of shorter data sections would at best have resulted in a $\sim 6\%$ improvement on the strain upper limits reported here (Eq. (2.2) of Ref. [2]).

Validation by hardware injections.—The analysis software was validated by checking its performance on fake pulsar signals injected in artificial and real detector noise both in software ([2]) and in hardware. In particular,

two artificial signals (P1, P2) were injected into all three detectors by modulating the mirror positions via the actuation control signals with the strain signal we should expect from a hypothetical pulsar. These injections were designed to give an end-to-end validation of the search pipeline starting from as far up the observing chain as possible.

The pulsar signals were injected for a 12 h period at frequencies of 1 279.123 Hz (P1) and 1 288.901 Hz (P2) with frequency derivatives of zero and $-10^{-8} \text{ Hz s}^{-1}$ respectively, and strain amplitudes of 2×10^{-21} . The signals were modulated and Doppler shifted to simulate sources at fixed positions on the sky with $\psi = 0$, $\cos \iota = 0$ and $\phi_0 = 0$. To illustrate, posterior pdfs for the recovered P1 signal are shown in Fig. 2. The results derived from the different detectors are in broad statistical agreement, confirming that the relative calibrations are consistent and that the assessments of uncertainty (expressed in the posterior widths) are reasonable. Results for P2 were very similar to these.

The phase stability of the detectors in S2 allowed us to implement a *joint* coherent analysis based on data from all three participating instruments. This technique was noted in [2], but could not be performed on the S1 data because of timing uncertainties that existed when those observations were performed. The solid lines in Fig. 2 show marginalisations of the joint posterior from H1, H2 and L1, i.e.,

$$p(\mathbf{a}|\text{H1, H2, L1}) \propto p(\mathbf{a})p(\text{H1}|\mathbf{a})p(\text{H2}|\mathbf{a})p(\text{L1}|\mathbf{a}). \quad (4)$$

With three detectors of roughly similar sensitivities and operational periods these coherent results should be approximately $\sqrt{3}$ times tighter than the individual results. The posteriors for ϕ_0 clearly highlight the relative coherence between the instruments and verify that similar joint methods can be used to set upper limits on our target pulsars.

Results.—From the ATNF pulsar catalogue (www.atnf.csiro.au/research/pulsar/psrcat/) we selected 28 isolated pulsars with rotational frequencies greater than 20 Hz and for which good timing data were available (Table I). For 18 of these, we obtained updated timing solutions from regular timing observations made at the Jodrell Bank Observatory using the Lovell and the Parkes telescopes, adjusted for a reference epoch centred on the period of the S2 run (starred pulsars in Table I). Details of the techniques that were used to do this can be found in [10]. We also checked that none of these pulsars exhibited a glitch during this period.

The list includes globular cluster pulsars (including isolated pulsars in 47 Tuc and NGC6752), the S1 target millisecond pulsar (PSR J1939+2134) and the Crab pulsar (PSR B0531+21). Although Table I only shows approximate pulsar frequencies and frequency derivatives, further phase corrections were made for pulsars with measured second derivatives of frequency. Timing solutions for the Crab were taken from the Jodrell Bank online ephemeris [11], and adjustments were made to its phase over the period of S2 using the method of [12].

pulsar	spin f (Hz)	spindown \dot{f} (Hz s^{-1})	$h_0^{95\%}$ $/10^{-24}$	ϵ $/10^{-5}$
B0021-72C*	173.71	$+1.50 \times 10^{-15}$	4.3	16
B0021-72D*	186.65	$+1.19 \times 10^{-16}$	4.1	14
B0021-72F*	381.16	-9.37×10^{-15}	7.2	5.7
B0021-72G*	247.50	$+2.58 \times 10^{-15}$	4.1	7.5
B0021-72L*	230.09	$+6.46 \times 10^{-15}$	2.9	6.1
B0021-72M*	271.99	$+2.84 \times 10^{-15}$	3.3	5.0
B0021-72N*	327.44	$+2.34 \times 10^{-15}$	4.0	4.3
J0030+0451	205.53	-4.20×10^{-16}	3.8	0.48
B0531+21*	29.81	-3.74×10^{-10}	41	2 100
J0711-6830	182.12	-4.94×10^{-16}	2.4	1.8
J1024-0719*	193.72	-6.95×10^{-16}	3.9	0.86
B1516+02A	180.06	-1.34×10^{-15}	3.6	21
J1629-6902	166.65	-2.78×10^{-16}	2.3	2.7
J1721-2457	285.99	-4.80×10^{-16}	4.0	1.8
J1730-2304*	123.11	-3.06×10^{-16}	3.1	2.5
J1744-1134*	245.43	-5.40×10^{-16}	5.9	0.83
J1748-2446C	118.54	$+8.52 \times 10^{-15}$	3.1	24
B1820-30A*	183.82	-1.14×10^{-13}	4.2	24
B1821-24*	327.41	-1.74×10^{-13}	5.6	7.1
J1910-5959B	119.65	$+1.14 \times 10^{-14}$	2.4	8.5
J1910-5959C	189.49	-7.90×10^{-17}	3.3	1.7
J1910-5959D	110.68	-1.18×10^{-14}	1.7	7.2
J1910-5959E	218.73	$+2.09 \times 10^{-14}$	7.5	7.9
J1913+1011*	27.85	-2.61×10^{-12}	51	6 900
J1939+2134*	641.93	-4.33×10^{-14}	13	2.7
B1951+32*	25.30	-3.74×10^{-12}	48	4 400
J2124-3358*	202.79	-8.45×10^{-16}	3.1	0.45
J2322+2057*	207.97	-4.20×10^{-16}	4.1	1.8

TABLE I: The 28 pulsars targeted in the S2 run, with approximate spin parameters. Pulsars for which radio timing data were taken over the S2 period are starred (*). The right-hand two columns show the 95% upper limit on h_0 , based on a coherent analysis using all the S2 data, and corresponding ellipticity values (ϵ , see text). These upper limit values do not include the uncertainties due to calibration and to pulsar timing accuracy, which are discussed in the text, nor uncertainties in r .

The analysis used 910 hours of data from H1, 691 hours from H2, and 342 hours from L1. There was no evidence of strong spectral contamination in any of the bands investigated, such as might be caused by an instrumental feature or a potentially detectable pulsar signal. A strong gravitational signal would generate a parameter pdf prominently peaked off zero with respect to its width, as for the hardware injections. Such a pdf would trigger a more detailed investigation of the pulsar in question. No such triggers occurred in the analysis of these data, and we therefore simply present upper limits.

The upper limits are presented as the value of h_0 bounding 95% of the cumulative probability of the marginalised strain pdf from $h_0 = 0$. The joint upper

limit $h_0^{95\%}$ therefore satisfies

$$0.95 = \int_{h_0=0}^{h_0^{95\%}} dh_0 \iiint p(a|H1, H2, L1) d\iota d\psi d\phi_0, \quad (5)$$

consistent with [2]. The uncertainty in the noise floor estimate is already included, as outlined above.

The remaining uncertainties in the upper limit values of Table I stem from the calibration of the instrument and from the accuracy of the pulsar timing models. For L1 and H2, the amplitude calibration uncertainties are conservatively estimated to be 10% and 8%, respectively. For H1, the maximum calibration uncertainty is 18%, with typical values at the 6% level. Phase calibration uncertainties are negligible in comparison: less than 10° in all detectors. Biases due to pulsar timing errors are estimated to be 3% or less for J0030+0451, and 1% or less for the remaining pulsars (see [2] for a discussion of the effect of these uncertainties).

Discussion.—The improved sensitivity of the LIGO interferometers is clear from the strain upper limit for PSR J1939+2134, which is more than a factor of ten lower than was achieved with the S1 data [2]. In this analysis the lowest limit is achieved for PSR J1910–5959D at the level of 1.7×10^{-24} , largely reflecting the lower noise floor around 200 Hz.

Table I also gives approximate limits to the ellipticities [6] of these pulsars from the simple quadrupole model

$$\epsilon \simeq 0.237 \frac{h_0}{10^{-24}} \frac{r}{1 \text{ kpc}} \frac{1 \text{ Hz}^2}{f^2} \frac{10^{45} \text{ g cm}^2}{I_{zz}} \quad (6)$$

where r is the pulsar's distance, which we take as the dispersion measure distance using the model of Taylor and Cordes [13], and I_{zz} its principal moment of inertia about the rotation axis, which we take as 10^{45} g cm^2 .

As expected, none of these upper limits improves on those inferred from simple arguments based on the gravi-

tational luminosities achievable from the observed loss of pulsar rotational kinetic energy. However, as discussed in the introduction, for pulsars in globular clusters such arguments are complicated by cluster dynamics, which the direct limits presented here avoid.

The result for the Crab pulsar (PSR B0531+21) is within a factor of about 30 of the spindown limit and over an order of magnitude better than the previous direct upper limit of [14]. The equatorial ellipticities of the four closest pulsars (J0030+0451, J2124+3358, J1024–0719, and J1744–1134) are constrained to less than 10^{-6} .

Once the detectors operate at design sensitivity for a year, the observational upper limits will improve by more than an order of magnitude. The present analysis will also be extended to include pulsars in binary systems, significantly increasing the population of objects under inspection.

Acknowledgments.—The authors gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory and the Particle Physics and Astronomy Research Council of the United Kingdom, the Max-Planck-Society and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Natural Sciences and Engineering Research Council of Canada, the Council of Scientific and Industrial Research of India, the Department of Science and Technology of India, the Spanish Ministerio de Ciencia y Tecnología, the John Simon Guggenheim Foundation, the Leverhulme Trust, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P. Sloan Foundation. This document has been assigned LIGO Laboratory document number LIGO-P040008-A-Z.

-
- [1] B. Abbott et al. (The LIGO Scientific Collaboration), *Nuclear Inst. and Methods in Physics Research A* **517**/1–3 154 (2004).
 - [2] B. Abbott et al. (The LIGO Scientific Collaboration), *Phys. Rev. D* **69** 082004 (2004).
 - [3] B. Abbott et al. (The LIGO Scientific Collaboration), *Phys. Rev. D* **69** 122001 (2004).
 - [4] B. Abbott et al. (The LIGO Scientific Collaboration), *Phys. Rev. D* **69** 102001 (2004).
 - [5] B. Abbott et al. (The LIGO Scientific Collaboration), *Phys. Rev. D* **69** 122004 (2004).
 - [6] P. Jaranowski, A. Królak and B. Schutz, *Phys. Rev. D* **58** 063001 (1998).
 - [7] B.J. Owen and D.I. Jones (2004), LIGO Technical Report T040192-00-Z, <http://www.ligo.caltech.edu/docs/T/T040192-00.pdf>
 - [8] R. Dupuis and G. Woan, in preparation (2004).
 - [9] G.L. Bretthorst "Bayesian spectrum analysis and parameter estimation" *Lecture Notes in Statistics*, Vol 48 (Springer-Verlag, Berlin) (1988).
 - [10] G. Hobbs, A.G. Lyne, M. Kramer, C.E. Martin, and C. Jordan, *MNRAS*, in press (2004).
 - [11] Jodrell Bank Crab Pulsar Timing: Monthly Ephemeris, <http://www.jb.man.ac.uk/research/pulsar/crab.html>.
 - [12] M. Pitkin and G. Woan, *Classical and Quantum Gravity* **21** S843 (2004).
 - [13] J.H. Taylor and J.M. Cordes, *ApJ* **411** 674 (1993).
 - [14] T. Suzuki, in *First Edoardo Amaldi Conference on Gravitational Wave Experiments*, edited by E. Coccia, G. Pizzella and F. Ronga (World Scientific, Singapore, 1995), p. 115.