

Gravitational waves from
accreting neutron stars and
Cassiopeia A

by
Karl William Wette

A thesis submitted for the degree of
Doctor of Philosophy
of
The Australian National University



November 2009

Declaration

This thesis is an account of research undertaken between March 2005 and October 2009. The research was conducted within the Centre for Gravitational Physics at the Australian National University, and during visits to the LIGO Hanford Observatory, the Max Planck Institute for Gravitational Physics (Albert Einstein Institute), the University of Melbourne, and the Pennsylvania State University.

The research described in this thesis is my own work, except where otherwise indicated. Research that has been conducted in collaboration with colleagues is explicitly acknowledged as such. Appropriate references to the work of others are included throughout.

To the best of my knowledge, the work presented in this thesis is original, and has not been submitted for a degree at any university.

Karl W. Wette
November 2009

Acknowledgments

I am grateful for the support I have received, throughout my Ph.D., from an Australian Postgraduate Award, an ANU Vice-Chancellor's Supplementary Scholarship, and from the Australian Research Council through grants DP0451021, SR0567380, and DP0770426. I am also grateful for the support I have received from the LIGO Hanford Observatory, the Pennsylvania State University, the University of Melbourne, and the Max Planck Institute for Gravitational Physics (Albert Einstein Institute).

It has been a pleasure and a privilege to have had the opportunity to work with so many talented scientists from around the world. Thank you to my supervisory panel, Susan Scott and David McClelland, for encouraging my work, and for shouldering the unenviable task of perpetual grant-writing which ensured I was never without funding for an all-expenses-paid trip to somewhere. Thank you to Greg Mendell for getting me started in gravitational wave data analysis, and for your kind hospitality during my long sojourn in Hanford. Thank you to Ben Owen for guiding me through the Cas A search, for interesting discussions on all aspects of this work, and for your continued support and encouragement. Thank you to Andrew Melatos for working with me on magnetic mountains, for your continued enthusiasm and encouragement of my work, and patience with my uneven rate of progress. Thank you to Mike Ashley for help wrestling with grid certificates, to Joseph Betzwieser for answering questions about upper limits, and to Matthias Vigelius for answering questions about ZEUS-MP. Thank you to Maria Alessandra Papa and Bruce Allen for taking an interest in my work, and for inviting me to the Albert Einstein Institute which greatly assisted the completion of much of this work. Thank you to Chris Messenger for many helpful discussions and guidance on numerous finicky data analysis issues. Thank you to Mike Landry, Keith Riles, and Graham Woan for your leadership of the LSC Continuous Wave Group, and for organising innumerable telecons. Thank you to Teviet Creighton, Fred Raab, and Peter Shawhan for carefully reviewing the Cas A work.

Thank you to Ra Ina, David McClelland, Andrew Melatos, Greg Mendell,

Ben Owen, Reinhard Prix, Susan Scott, Antony Searle, and Gerhard Wette for timely and helpful comments on thesis drafts. Thank you to my Ph.D. examiners for their careful reading of this thesis, and for their many helpful comments and suggestions. Any remaining errors or omissions are, of course, the responsibility of the author.

Special thanks to all the people I've met over the last five years, who have helped to distract me from the many trials and frustrations of a Ph.D. life. Most especially, thanks to my wonderful family for your support and encouragement during my many years as a student, for the letters and parcels from home to wherever I was in the world, and for counselling me through the dark and difficult times.

In loving memory of my grandparents:

Nana Joan and Poppa Jack,

Oma Erna und Opa Jobst

Abstract

This thesis is concerned with the mysteries of neutron stars and the quest for gravitational waves. Rapidly-rotating neutron stars are anticipated sources of periodic gravitational waves, and are expected to be detectable within the next decade using kilometre-scale laser interferometry.

We first perform ideal-magnetohydrodynamic axisymmetric simulations of a magnetically confined mountain on an accreting neutron star. Two scenarios are considered, in which the mountain sits atop a hard surface or sinks into a soft, fluid base. We quantify the ellipticity of the star, due to a mountain grown on a hard surface, and the reduction in ellipticity due to sinking. The consequences for gravitational waves from low-mass x-ray binaries are discussed.

We next present two approaches to reducing the computational cost of searches for periodic gravitational waves. First, we generalise the PowerFlux semi-coherent search method to estimate the amplitudes and polarisation of the periodic gravitational wave signal. The relative efficiencies of the generalised and standard methods are compared using simulated signals. Second, we present an algorithm which minimises the number of templates required for a fully coherent search, by using lattice sphere covering to optimally place templates in the search parameter space. An implementation of the algorithm is tested using Monte Carlo simulations.

Finally, we present a coherent search for periodic gravitational waves targeting the central compact object in the supernova remnant Cassiopeia A, using data from the fifth science run of the Laser Interferometer Gravitational-Wave Observatory. The search parameter space is determined by the sensitive frequencies of the detectors, by the age of the compact object, and a range of braking indices. No gravitational wave signal is detected. We set an upper limit on the strength of gravitational waves from the compact object in Cassiopeia A, which surpasses the theoretical limit based on energy conservation. Cassiopeia A is thus one of only a few astronomical objects, to date, where gravitational wave observations are beginning to constrain astrophysics.

Contents

Declaration	iii
Acknowledgments	v
Abstract	ix
1 Introduction	1
1.1 Author contributions and publications	3
2 Neutron stars and gravitational waves	7
2.1 The last hundred years	7
2.2 Neutron stars	9
2.3 Gravitational waves	12
2.4 Periodic gravitational waves from spinning neutron stars . . .	17
2.5 Summary	20
3 A method of building stable mountains with sinking	21
3.1 Magnetic burial	21
3.2 Growing a realistically sized mountain by injection	23
3.2.1 Previous work	23
3.2.2 Outline of the method	24
3.2.3 Initial setup	27
3.2.4 Injection procedure	29
3.3 Summary	31
3.A Additional material	31
3.A.1 Matching a fluid base to a Grad-Shafranov mountain .	31
3.A.2 Custom injection procedure	32
4 Sinking of a mountain on an accreting neutron star	35
4.1 Simulations of magnetic mountains with sinking	35
4.1.1 Verification	37
4.1.2 Illustrative example	40

4.2	Comparison of mountains grown on hard and soft bases	41
4.2.1	Sinking scenarios	46
4.2.2	Magnetic line tying	48
4.2.3	Mass quadrupole moment	51
4.3	Comparison with Choudhuri & Konar	53
4.4	Implications for gravitational waves	54
4.5	Summary	54
5	Gravitational wave data analysis	57
5.1	Summary of gravitational waves searches	57
5.1.1	Compact binary coalescences	58
5.1.2	Bursts	59
5.1.3	The stochastic background	60
5.1.4	Periodic gravitational waves	60
5.2	Periodic gravitational wave data analysis	61
5.2.1	The signal model	62
5.2.2	The \mathcal{F} statistic	67
5.2.3	Computational cost	72
5.3	Summary	74
6	Parameter estimation using generalised PowerFlux	75
6.1	Semi-coherent data analysis	75
6.2	A derivation of the PowerFlux method	76
6.3	Generalisations of PowerFlux	77
6.3.1	Estimation of A_+^2 and A_\times^2	78
6.3.2	Estimation of A_+^2 , A_\times^2 , and ψ	79
6.4	Comparison of the methods	80
6.4.1	Parameter distributions	83
6.4.2	Detection efficiencies	84
6.5	Summary	84
7	Template bank generation using optimal lattices	85
7.1	Lattices and sphere coverings	85
7.1.1	Lattices	86
7.1.2	Sphere coverings	87
7.2	The parameter space metric	89
7.2.1	The metric of the \mathcal{F} statistic	91
7.3	A template bank generation algorithm	92
7.3.1	Iteration over the parameter space	93
7.3.2	The lattice generator	97
7.3.3	Number of templates required for coverage	101

7.4	Implementation	102
7.4.1	Testing the implementation	103
7.5	Summary	105
8	A search for gravitational waves from Cassiopeia A	107
8.1	The central compact object	107
8.1.1	Motivation for a gravitational wave search	109
8.2	Indirect upper limits	110
8.2.1	Uncertainty in the upper limits	112
8.3	Choice of analysis method and time span of data set	113
8.4	The frequency and spindown parameter space	113
8.4.1	Application of the template bank algorithm	115
8.5	Sensitivity of the search	119
8.6	Computational cost and expected upper limits	123
8.7	The expected largest $2\mathcal{F}$	123
8.7.1	Number of statistically independent templates	125
8.8	Summary	128
9	Upper limits on gravitational waves from Cassiopeia A	129
9.1	The LIGO S5 science run	129
9.1.1	Selection of data for the search	130
9.2	The search pipeline	134
9.3	Vetoing of spuriously large $2\mathcal{F}$	134
9.3.1	Contamination of the \mathcal{F} statistic by lines	135
9.3.2	Identification of narrow line features	137
9.3.3	Application to the search results	138
9.4	The largest $2\mathcal{F}$	145
9.5	Upper limits	147
9.5.1	Analytic estimation of upper limits	149
9.5.2	Upper limits on h_0 and ϵ	150
9.6	Summary	153
9.A	Additional material	153
9.A.1	Known instrumental lines	153
10	Conclusion	159
	Bibliography	163

Figures

2.1	Periods, period derivatives, and surface magnetic fields of known pulsars	11
2.2	Perturbation of a ring of test masses due to plus-polarised and cross-polarised plane gravitational waves	13
2.3	Schematic of a Michelson interferometer	15
3.1	Diagrams illustrating schematically three mountain growth scenarios	25
3.2	The accretion rate, and the initial mass-flux distribution . . .	29
4.1	Radial component of the magnetic field along $\theta = 0$	38
4.2	Accreted mass, and masses in the stellar atmosphere and fluid base	39
4.3	Kinetic energy normalised by the magnetic energy	40
4.4	Contours of the absolute, normalised divergence of \mathbf{B} of mountain $\mathcal{H}(10^3)$	41
4.5	Hydromagnetic structure of hard-surface mountain $\mathcal{H}(10^2)$: contours of accreted density	42
4.6	Hydromagnetic structure of hard-surface mountain $\mathcal{H}(10^2)$: contours of magnetic flux	43
4.7	Comparison of the hydromagnetic structure of hard- and soft-surface mountains: contours of accreted density	44
4.8	Comparison of the hydromagnetic structure of hard- and soft-surface mountains: contours of magnetic flux	45
4.9	Shaded contours of the absolute, normalised difference in total density between $\mathcal{S}(r_{\min}, 10)$ and $\mathcal{S}(R_{\star}, 10)$	47
4.10	Angle between the magnetic field and the radial unit vector .	49
4.11	Ellipticity as a function of time, for $M_a/M_c = 1, 10, 10^2$, and 10^3	51
4.12	Ellipticity as a function of M_a for \mathcal{H} , $\mathcal{S}(r_{\min})$, and $\mathcal{S}(R_{\star})$. . .	52

5.1	Transformation from the wave frame to the detector frame, via the celestial sphere frame	65
5.2	Optimal signal-to-noise ratio ρ^2 as a function of the observation time, and the resultant probability distributions of $2\mathcal{F}$	71
6.1	Injected versus detected normalised squared amplitudes for linear, circular, and generalised PowerFlux	81
6.2	Detections efficiencies versus the normalised injected amplitude for linear, circular, and generalised PowerFlux, and Stack-Slide	82
7.1	The hexagonal lattice and its covering spheres.	86
7.2	Edges effects at the boundary of a two-dimensional parameter space	95
7.3	Transformation of the the original hexagonal lattice generator to the final generator for a 2-dimensional template bank	98
7.4	Testing the <code>FlatLatticeTiling</code> implementation of the template bank algorithm	104
8.1	Histograms of the probability distributions of the upper limits on the ellipticity ϵ , assuming $f = 100$ Hz, and the strain h_0	112
8.2	Visualisation of the Cas A search parameter space	114
8.3	\dot{f} - \ddot{f} cross-section through the template tiling of the Cas A parameter space, at $f = 200$ Hz	116
8.4	\dot{f} - \ddot{f} cross-section through the template tiling of the Cas A parameter space, at $\ddot{f} = 8.38 \times 10^{-19}$ Hz s ⁻²	117
8.5	Cumulative number of templates required to cover the Cas A search parameter space	119
8.6	False alarm rate p_α versus detection threshold $2\mathcal{F}_\alpha$, and false dismissal rate p_β versus sensitivity Θ	122
8.7	Expected upper limits obtainable by the Cas A search	122
8.8	Distribution of the largest $2\mathcal{F}$, in each search job frequency band, and for the entire search	126
8.9	Fractional difference between the theoretical distribution of the \mathcal{F} statistic and the distribution of the \mathcal{F} statistic as implemented by <code>ComputeFStatistic_v2</code>	127
9.1	Figure of merit $F(t_0)$ for all 12-day intervals containing SFTs before GPS 861000000	132
9.2	Spectrograms of the H1 and L1 SFTs comprising the Cas A search data set	133

9.3	Largest $2\mathcal{F}$ found in each search job frequency band, before post-processing	135
9.4	The action of a spectral running median on a discrete spectrum containing lines	136
9.5	Histograms of the StackSlide power of the Cas A search SFTs	137
9.6	Times during which pulsar hardware injections were injected into L1 and H1	140
9.7	Veto bands triggered by pulsar hardware injections no. 3 and no. 8.	141
9.8	Veto bands triggered by probable instrumental lines in L1 at ~ 139.5 Hz and ~ 209.3 Hz	144
9.9	Largest $2\mathcal{F}$ found in each search job frequency band, after post-processing	146
9.10	The search job band containing the largest non-vetoed value of $2\mathcal{F}$	146
9.11	The largest value of $2\mathcal{F}$ found by the Cas A search, and its expected distribution	147
9.12	Confidence of the upper limits in each upper limit band as determined by software injections	149
9.13	Upper limits at 95% confidence on the strain of gravitational waves from Cas A and the ellipticity of Cas A	151
9.14	Veto bands triggered by instrumental lines at ~ 119.9 Hz, ~ 128.0 Hz, and ~ 179.8 Hz	154
9.15	Veto band triggered by instrumental L1 line at ~ 139.2 Hz . .	155
9.16	Veto band triggered by instrumental L1 line at ~ 144.8 Hz . .	155
9.17	Veto band triggered by instrumental L1 line at ~ 185.6 Hz . .	156
9.18	Veto band triggered by instrumental L1 line at ~ 193.0 Hz . .	156

Tables

3.1	Conversion of physical quantities into dimensionless variables in the simulations	28
4.1	Simulations of magnetic mountains presented in this chapter .	36
4.2	Important physical parameters of accreting neutron stars, and a summary of the results of the simulations	37
7.1	Thicknesses of a selection of lattice sphere coverings, in up to the 6th dimension	88
9.1	Data quality flags applied to the L1 and H1 science data . . .	131
9.2	Vetoed frequency bands identified by the threshold on $\mathcal{P}(f)$.	138
9.3	Doppler parameters of the two pulsar hardware injections strong enough to be seen by the Cas A search	140

Chapter 1

Introduction

This thesis is concerned with two exciting arenas of modern physics: neutron stars, and gravitational waves. Since the beginning of the twentieth century, experiments have been conducted, with increasing precision, to probe the nature of gravity, and to test the predictions of its most successful theory: the general theory of relativity. Gravitational waves, predicted by the theory, are the next frontier of experimental gravity, and a first direct detection is widely anticipated within the next decade. Meanwhile, the existence of neutron stars has been confirmed by the discovery of pulsars, which have been studied in detail over the last half-century. Neutron stars are important gravitational wave sources; in turn, gravitational wave astronomy offers the possibility of gaining further insights into neutron star physics. Chapter 2 introduces the physics of neutron stars and gravitational waves, and discusses one important link between them.

Chapters 3–4 are concerned with the physics of accreting neutron stars. Specifically, we investigate the formation by the accreted matter of a mountain on the stellar surface, held in a stable equilibrium by the star’s powerful magnetic field. The burial of the magnetic field by the accreted mountain can explain why the magnetic field of a neutron star reduces as the star accretes, which is important for models of the evolution of binary pulsar systems. Magnetic mountains are also a plausible means of generating gravitational radiation from low-mass x-ray binaries. The effect of the magnetic mountain sinking into the neutron star crust has not been accounted for in previous work, and is the key advance of the work presented here.

In Chapter 3, we review previous work on this problem, and present a numerical procedure capable of building magnetic mountains with realistic masses. In Chapter 4, we present a detailed comparison of two scenarios, where the magnetic mountain either sits atop a hard surface or sinks into a soft, fluid base. We discuss the evolution of a magnetic mountain during

accretion, and compare the hydromagnetic structures of mountains grown on hard and soft bases. We allow the mountain to sink in two different but theoretically identical scenarios, and confirm that the resultant equilibria are equivalent. We quantify the ellipticity of the neutron star, due to a mountain grown on a hard surface, and the reduction in ellipticity due to sinking. Finally, we compare our simulations to the model of Choudhuri & Konar (2002), and discuss the consequences for gravitational waves from low-mass x-ray binaries.

Chapters 5–9 are concerned with gravitational waves, and the challenge of analysing the output of kilometre-scale laser interferometric detectors in search of their signatures. In Chapter 5, we summarise results from the searches for gravitational waves conducted to date. We then review the analysis of periodic gravitational waves, which are anticipated to be generated by rapidly rotating neutron stars. We present the analytic model of the periodic gravitational wave signal, the coherent matched filtering technique used to search for them, and discuss the computational cost of such searches. We then explore, in Chapters 6–7, two different approaches to lowering the computational cost of periodic gravitational wave searches.

In Chapter 6, we consider semi-coherent search techniques. Compared to fully coherent matched filtering, these techniques are less sensitive, but are also less computationally intensive; as a result, they can achieve greater overall sensitivity by searching longer stretches of data than would be computationally feasible using a coherent method. We consider the PowerFlux semi-coherent method, and present an alternative derivation to that of Dergachev & Riles (2005). We then generalise the PowerFlux method to estimate the amplitudes of the plus and cross polarisations, and the polarisation angle of the periodic gravitational wave signal. Using simulated signals injected into Gaussian noise, we compare the parameter estimation and detection efficiencies of the generalised and standard PowerFlux methods.

In Chapter 7, we present an algorithm which generates a bank of templates for a coherent search over a given template parameter space. The algorithm is designed to minimise the number of templates required to cover the parameter space, thus minimising the computational cost of the search, while ensuring that any potential signal will still be closely matched by some template in the bank. The algorithm uses sphere coverings on optimally thin lattices to position the points in the parameter space, and a metric on the parameter space to ensure the correct spacing. Particular care is taken to generate extra templates along the edges of the parameter space to ensure that they are completely covered. The chapter introduces the necessary background material, presents the algorithm, discusses how to estimate the number of templates the algorithm requires for coverage, and tests the per-

formance of an implementation of the algorithm.

Finally, in Chapters 8–9, we present a search for periodic gravitational waves targeted at the central compact object in the supernova remnant Cassiopeia A. The compact object is likely the youngest known neutron star, and has been widely studied by astronomers since its discovery a decade ago. No pulsations are observed from the compact object, and it therefore has no known spin frequency. The search uses data from the fifth science run (S5) of the Laser Interferometer Gravitational-Wave Observatory (LIGO), and is the first gravitational wave search to target a known non-pulsing neutron star. An indirect upper limit on the strength of gravitational waves from the compact object can be beaten, over a range of frequencies, using a coherent search of 12 days of LIGO S5 data. Cassiopeia A is therefore one of the few periodic gravitational wave sources which could conceivably be seen by LIGO at its present sensitivity.

In Chapter 8, we review electromagnetic observations of Cassiopeia A, and the motivation for a gravitational wave search. We derive the indirect upper limit on the compact object, and present details of the proposed search, including the choice of analysis method and the time span of the data set, the parameter space of frequencies and frequency derivatives to be searched, and its estimated sensitivity. We confirm that the search will beat the indirect upper limit, and is computationally feasible. Chapter 9 presents the implementation of the search, which includes the selection of data from the LIGO S5 run, the execution of the search pipeline, and post-processing procedures to remove false candidates arising from instrumental noise. After performing these steps, we find that there is no evidence for the detection of a gravitational wave signal from Cassiopeia A. We then determine upper limits on the strength of gravitational waves from Cassiopeia A which, as expected, beat the indirect limit over the range of frequencies searched. Cassiopeia A is one of only a handful of astronomical objects for which this has been achieved.

Chapter 10 summarises the thesis and considers possible directions for further research.

1.1 Author contributions and publications

While this thesis is substantially the work of the author, it also includes work that was done in collaboration with colleagues. This section describes in full the contributions made by the author to the research presented in each chapter, and any publications on which the chapter is based.

Chapter 2 reviews background information relevant to the thesis as a whole.

Chapters 3–4 present work done in collaboration with Andrew Melatos and Matthias Vigelius (University of Melbourne). The two chapters are closely based on the following publication:

Wette *et al.* (2010):

K. Wette, M. Vigelius, and A. Melatos, 2010. Sinking of a magnetically confined mountain on an accreting neutron star. *Monthly Notices of the Royal Astronomical Society* 402, 1099.

Chapter 5 reviews background information relevant to Chapters 6–9.

Chapter 6 presents work done in collaboration with Gregory Mendell (LIGO Hanford Observatory). The chapter is closely based on the following publication:

Mendell & Wette (2008):

G. Mendell and K. Wette, 2008. Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves. *Classical and Quantum Gravity* 25(11), 114044.

Chapter 7 presents the work of the author, and has not been published.

Chapters 8–9 present work conducted within the LIGO Scientific Collaboration (LSC). The LSC works alongside the LIGO Laboratory to undertake the science of LIGO. This includes operating the Observatory’s twin gravitational wave detectors during data acquisition, monitoring the detectors’ performance to ensure the acquired data is of science quality, the correct calibration of the raw detector output, the generation of derived data products from the raw output, authoring the scientific software used in the analysis of the data, and the operating and maintaining of large-scale computer clusters on which the analyses are performed.

Author contributions The author participated in the acquisition of LIGO data by serving as a science monitor at the LIGO Hanford Observatory. The author held primary responsibility for conducting the search for periodic gravitational waves from the supernova remnant Cassiopeia A presented in Chapters 8–9. This included deciding on the various parameters of the search, the design and authoring of the template bank generation algorithm described in Chapter 7, authoring scripts to manage the search pipeline, submitting and managing search jobs on computer clusters, designing and executing the post-processing of the results, and determining the upper

limits. The author greatly benefitted from discussions, at regular teleconferences and at face-to-face meetings, with colleagues in the Continuous Wave (CW) Working Group of the LSC. The author was also responsible for presenting the search to the CW Review Committee, which approved the presentation of preliminary results at the 8th Edoardo Amaldi Conference on Gravitational Waves. The same results are reproduced in this thesis.

Publications Chapters 8–9 expand upon the following publications:

Wette *et al.* (2008):

K. Wette, B. J. Owen, B. Allen, M. Ashley, J. Betzwieser, N. Christensen, T. D. Creighton, V. Dergachev, I. Gholami, E. Goetz, R. Gustafson, D. Hammer, D. I. Jones, B. Krishnan, M. Landry, B. Machenschalk, D. E. McClelland, G. Mendell, C. J. Messenger, M. A. Papa, P. Patel, M. Pitkin, H. J. Pletsch, R. Prix, K. Riles, L. S. de la Jordana, S. M. Scott, A. M. Sintes, M. Trias, J. T. Whelan, and G. Woan, 2008. Searching for gravitational waves from Cassiopeia A with LIGO. *Classical and Quantum Gravity* 25(23), 235011.

Abadie *et al.* (2010):

J. Abadie *et al.* (LIGO Scientific Collaboration), 2010. First search for gravitational waves from the youngest known neutron star. *The Astrophysical Journal* in press.

Disclaimer The search for Cassiopeia A is, at the time of writing, under internal review within the LSC; the material presented in Chapters 8–9 is therefore subject to change. The views and opinions expressed in this thesis regarding the Cassiopeia A search are those of the author, and do not necessarily reflect the views and opinions of the LSC.

Acknowledgements The author gratefully acknowledges the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, and the support of the National Science Foundation, the Australian Research Council, and many other funding agencies for the research by members and member institutions of the LSC.

Chapter 2

Neutron stars and gravitational waves

This chapter introduces the two physical phenomena which concern this thesis: neutron stars, and gravitational waves. We begin with a brief history of their discovery and study since the beginning of the 20th century (section 2.1). We then review the formation, structure, and taxonomy of neutron stars (section 2.2), gravitational waves, and their detection using large-scale laser interferometry (section 2.3). In section 2.4, we discuss one important link between the two phenomena: periodic gravitational waves are expected to be generated by spinning neutron stars.

2.1 The last hundred years

In the early years following the debut of general relativity (Einstein 1916), only three tests of the theory were experimentally accessible (Bertotti *et al.* 1962, Will 1993, 2006, and references therein). While the perihelion advance of Mercury was immediately confirmed, and to 1% accuracy, attempts to conduct the remaining two tests were hampered by experimental limitations; measurements of the deflection of light around the Sun achieved 50% accuracies at best, while the gravitational red-shift of light was never reliably detected during this period. It was not until the latter half of the century that gravitational experiments began to be conducted to high precision.

Meanwhile, the neutron was discovered in 1932 (Shapiro & Teukolsky 1983, and references therein); soon after, it was suggested that cold dense stars, composed of the new particle, might be formed in supernovae (Baade & Zwicky 1934). Initial work on the internal structure of neutron stars (beginning with Oppenheimer & Volkoff 1939) was motivated by the possibility

that neutron cores were the source of stellar energy in massive stars (Landau 1967); when this idea was superseded by thermonuclear fusion, interest in neutron stars declined. General relativity was also neglected: in addition to the three classical tests, it could explain the observed expansion of the Universe, but predicted an age, based on initial observations, that was younger than Earth. In the absence of definitive experimental corroboration of general relativity, a number of alternative theories of gravity emerged, which generally eschewed curved spacetime, as well as competing cosmologies such as the steady-state Universe.

Beginning in the 1960s, a number of important discoveries revived interest in both general relativity and neutron stars (Shapiro & Teukolsky 1983, Will 1993, 2006). Non-solar x-ray sources, first observed in 1962, were thought to be young neutron stars. Quasars were discovered in 1963; to understand their vast expenditure of energy and compact size, general relativity was applied to problems in astrophysics for the first time. This motivated a systematic testing of the theory, beginning with its weak field effects, and which made use of advances in quantum technologies such as lasers. This experimental effort, which saw its heyday between 1960 and 1980, included laboratory measurements of the gravitational red-shift, the use of radio interferometry to measure the light deflection from quasars, and lunar laser ranging experiments to constrain the Nordtvedt effect.

While initial hypotheses linking quasars to neutron stars were discounted, pulsars, discovered in 1967 (Hewish *et al.* 1968), were soon identified as rotating neutron stars (Shapiro & Teukolsky 1983, Will 1993, 2006). Within a year, the Crab and Vela pulsars were found in supernova remnants, confirming the hypothesised formation of neutron stars in supernovae. X-ray pulsars, discovered a few years later, were found to be neutron stars accreting matter from a binary companion star. Then, in 1974, came the first discovery (Hulse & Taylor 1975) of two neutron stars, one of them a pulsar, orbiting each other – a binary pulsar. The relativistic gravitational field of the binary system, combined with the precision measurements obtainable from radio observations of the pulsar, allowed for accurate measurements of astrophysical parameters, such as the neutron star masses, and permitted general relativistic effects to be tested for the first time. In particular, from the gradual shortening of the orbital period, it was deduced that the system was losing orbital energy in the form of gravitational radiation, at the level predicted by general relativity (Taylor & Weisberg 1982, 1989). Since their initial discovery, the population of known pulsars now numbers in the thousands; it includes rotation- and accretion-powered pulsars, high- and low-mass x-ray binaries, and double pulsar systems (Lorimer 2008).

With strong, but indirect, evidence for the existence of gravitational

waves, a new experimental effort was begun to attempt a direct detection using Earth-based instruments (Saulson 1994, and references therein). The first generation of gravitational wave detectors were resonant mass (or bar) detectors at room temperatures; early claims of detections using these instruments were never independently confirmed, and remain controversial. The next generation of cryogenically-cooled bar detectors started operation in the 1980s, and grew to a network of up to five detectors capable of long observation times; these detectors targeted high-frequency gravitational wave bursts (Astone *et al.* 2003, 2007). In addition, the bar detector in Western Australia was used to search for continuous gravitational waves from a millisecond pulsar whose gravitational wave frequency fortuitously coincided with the resonant frequency of the bar (Dhurandhar *et al.* 1996). The idea of using kilometre-scale interferometers as gravitational wave detectors began to be seriously considered in the 1980s, and construction on a number of projects began in the 1990s. Four kilometre-scale gravitational wave interferometers are currently in operation (Acernese *et al.* 2008, Abbott *et al.* 2009e), and are sensitive to spacetime perturbations on the order of 10^{-18} m.

2.2 Neutron stars

In a normal star, the inward pressure of the star's self-gravity is balanced by the outward radiation pressure sustained by nuclear fusion, which takes place initially within the stellar core, and subsequently within a thin shell that expands outward as the star burns through its nuclear fuel (Shapiro & Teukolsky 1983, and references therein). The stellar core collapses under its self-gravity and heats up, providing sufficient energy to catalyse further nuclear reactions, the precise sequence of which is sensitively dependent on the initial mass of the star. Inevitably, however, the weakening nuclear furnace is unable to support the star's own weight; the star ceases to burn, and collapses.

The circumstances under which a star collapses to a particular type of compact object are not known for certain. The star's pre-collapse mass and metallicity are believed to be important parameters (e.g. Shapiro & Teukolsky 1983, Heger *et al.* 2003, Janka 2004, Zhang *et al.* 2008), and recent insights have come from numerical simulations of supernova explosions (e.g. Buras *et al.* 2003, Heger *et al.* 2005, Ott *et al.* 2006, Burrows *et al.* 2007, Dimmelmeier *et al.* 2008, Marek & Janka 2009). If the mass of the star is a few times the mass of the Sun (denoted M_{\odot}), it will eject its outer layers into a planetary nebula and form a white dwarf star, which supports its weight by the pressure of degenerate electrons (Shapiro & Teukolsky 1983). Neutron

stars are believed to form during the collapse and subsequent supernova of stars with masses in the vicinity of 10 to $20M_{\odot}$; the collapse of more massive stars are expected to form black holes (Shapiro & Teukolsky 1983, Fryer 1999, Heger *et al.* 2003, Kokkotas 2008). Massive stars which undergo rapid mass loss due to a strong stellar wind (known as Wolf-Rayet stars) may, however, be left with insufficient mass to trigger a supernova. A white dwarf may also collapse further to form a neutron star, after the accretion of additional matter from a companion star.

Neutron stars have typical masses of $1.4M_{\odot}$ and radii of ~ 10 km (Shapiro & Teukolsky 1983, Fryer & Kalogera 2001). They are sufficiently dense that electrons are captured by protons to form neutrons, in the process radiating neutrinos which rapidly cools the star. As the name suggests, neutrons become the dominant particle in neutron stars, and their degeneracy pressure is sufficient to halt any further gravitational collapse. Despite this uniformity in composition, the internal structure of neutron stars is far from simple, and indeed is not entirely understood (Chamel & Haensel 2008, and references therein). The interior is first divided into two regions: a central core, where the density exceeds typical values for atomic nuclei of $\sim 10^{14}$ g cm $^{-3}$, and a surrounding outer layer, the crust, which contains only a small fraction of the total mass, and at sub-nuclear densities. The crust is further divided into an inner and outer crust, where the division is at the density of neutron drip ($\sim 10^{11}$ g cm $^{-3}$), where neutrons begin to no longer be bound within nuclei. The high-density physics of the core is a challenging theoretical problem, and remains the least understood. The structure and composition of the crust differs according to, among other factors, how the star was formed, and whether it is undergoing accretion. In the latter case, for example (Brown & Bildsten 1998, Chamel & Haensel 2008), the star acquires a series of outer shells composed of hydrogen burning to helium, which sinks to atop the iron crust and is compressed; at $\sim 10^7$ g cm $^{-3}$, a helium flash is ignited, leading to explosive burning of the outer layers and a burst of x-rays.

Pulsars are rapidly rotating neutron stars with strong ($\sim 10^{12}$) magnetic fields (Shapiro & Teukolsky 1983, Kramer 2005, and references therein). In a simple model, proposed soon after their discovery, the magnetic pole is misaligned to the rotation axis, generating electromagnetic radiation which, on Earth, is most commonly observed at radio frequencies. This radiation is collimated into beams which, if they sweep across the Earth, are observed as radio pulses at the frequency of rotation. Subsequent research has led to a more sophisticated model: the magnetic field of the star induces a quadrupole electric field strong enough to rip charged particles from its own surface. These charges form a co-rotating plasma, the magnetosphere, which is confined by the star's magnetic field, and extends out to the radius of

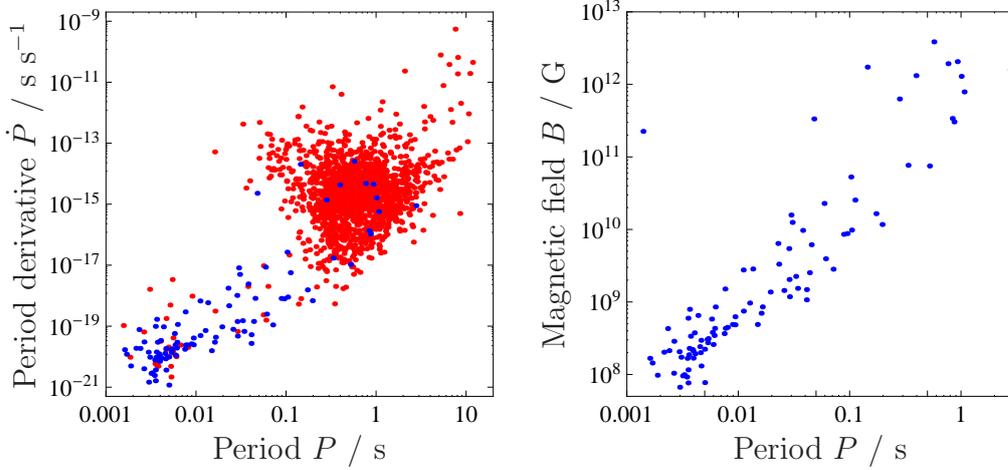


Figure 2.1: Periods P , period derivatives \dot{P} , and surface magnetic fields B of known pulsars, as given by the ATNF (2009) pulsar catalogue (Manchester *et al.* 2005). (left) \dot{P} versus P for isolated (red) and binary (blue) pulsars. (right) B versus P for binary pulsars. B is calculated from P and \dot{P} by assuming rotational energy is lost through magnetic dipole radiation, which gives the relation $B/\text{G} = 3.2 \times 10^{19} (P\dot{P}/\text{s})^{1/2}$ (e.g. Kramer 2005).

the so-called light cylinder, where the co-rotation velocity of the plasma approaches light speed. Charged particles may escape along open magnetic field lines, i.e. those which pass through this cylinder, and it is these particles which generate the radio emission, although the precise mechanism has yet to be determined (e.g. Kramer *et al.* 1997, Kramer 2005, Lorimer 2008). Not all neutron stars are observed as pulsars; for example, the supernova remnant Cassiopeia is believed to contain the youngest known neutron star, which is not observed as a pulsar (see section 8.1). The population of non-pulsing neutron stars is targeted by all-sky searches for periodic gravitational waves (see section 5.1.4).

The periods of all known pulsars span four orders of magnitude, from $P \sim 10^{-3}$ s (ATNF 2009). Isolated pulsars typically spin slower than 0.1 s, whereas pulsars in binary systems typically spin faster than this. Pulsars with periods in the vicinity of 10^{-3} to 10^{-2} s are often specifically referred to as millisecond pulsars. Characteristically, pulsar periods increase over time due to the conversion of rotational kinetic energy into electromagnetic (and, potentially, gravitational) radiation. Spindown rates range from $\dot{P} \sim 10^{-21} \text{ s s}^{-1}$ to 10^{-8} s s^{-1} (ATNF 2009). The most interesting exception to this rule are glitches: occasionally, some pulsars are observed to suddenly spin

up rapidly, then return to their former spindown rate over a short timescale. A logarithmic plot of \dot{P} versus P (Figure 2.1, left) can be used to describe the evolution of pulsars, by analogy with the Hertzsprung–Russell diagram for normal stars (Kramer 2005). A young isolated pulsar begins life around the top-left hand corner of the diagram, i.e. with a short period and high spindown. As the pulsar loses rotational kinetic energy, thereby increasing its period, the rate of energy loss also decreases, and with it the period derivative. Thus, a pulsar will move roughly down and to the right in the \dot{P} – P diagram as it ages. It is believed that pulsars eventually reach a state where their slow rotation can no longer power their radio emission, and they cease to be observed as a pulsar. This transition is often marked by a *pulsar death line* in the \dot{P} – P diagram, although this does not explain the existence of some slowly-spinning pulsars (e.g. Young *et al.* 1999). Binary pulsars, on the other hand, follow a different evolutionary path (see below).

Observed pulsars are associated with the Galaxy, the Large and Small Magellanic Clouds, and globular clusters (Lorimer 2008, and references therein). While the majority are isolated pulsars, a few percent are found in supernova remnants (confirming the link between supernovae and neutron stars), and in binary systems with white dwarf, main sequence, or neutron star companions. (Two are even orbited by a planet.) The taxonomy of pulsars can be explained by a simple model describing the evolution of a binary star. When the more massive star supernovas, the binary is most likely to be disrupted, due to the likely asymmetric explosion, and separate into a lone star and a young pulsar. If it survives, and the second star has sufficient mass to become a red giant which overflows its Roche lobe (the region within which matter is gravitationally bound to the star), the neutron star will accrete matter from its companion and emit x-rays: the system is now observed as an x-ray binary. Accretion will gradually decrease the period of the pulsar as angular momentum is transferred from its companion; this spinning up process is also known as recycling. During this process, the magnetic field of the neutron star is observed to decrease in proportion to the duration of the accretion process (Choudhuri & Konar 2002, and references therein). This is shown in Figure 2.1 (right): binary pulsars with shorter periods (indicating longer accretion phases) possess much weaker magnetic fields.

The physical mechanism responsible is not known for certain, and a number of models have been proposed (see Vigelius 2008, and references therein). One in particular, that the accreted matter acts to bury the magnetic field within the crust, is the subject of Chapters 3 and 4 of this thesis.

The ultimate fate of an x-ray binary lies with the relative mass of the companion star. A relatively low-mass companion will eventually collapse to a white dwarf, by which time the neutron star has spun up to a millisecond

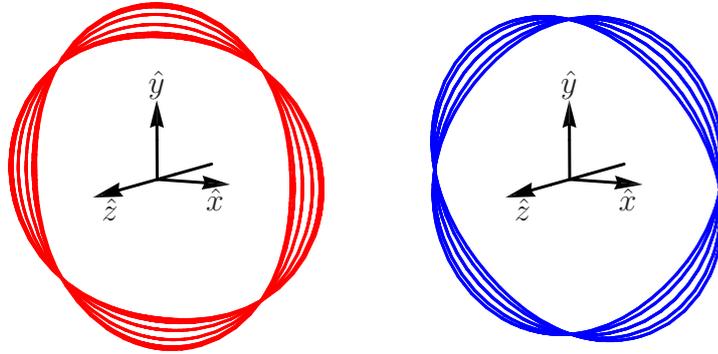


Figure 2.2: Perturbation of a ring of test masses due to either a plus-polarised (left, red) or cross-polarised (right, blue) plane gravitational waves travelling in the \hat{z} direction.

period; this is consistent with the period distribution of binary pulsars seen in Figure 2.1 (left). A relatively high-mass companion will itself supernova and form a second neutron star; if the binary again survives, it will become a rare double pulsar system.

2.3 Gravitational waves

Gravitational waves are wave-like solutions of the linearised field equations of general relativity (Shapiro & Teukolsky 1983, Thorne 1980, 1987, Kokkotas 2008, and references therein). The linearised equations are applicable sufficiently far from the source, so that the waves can be regarded as small perturbations to an otherwise flat spacetime. Gravitational radiation is quadrupolar or higher; monopole and dipole radiation are forbidden by conservation of energy and momentum respectively. The gravitational wave strain is proportional to the second time derivative of the quadrupole moment of its source. Thus, in order to produce gravitational waves, a source must possess a non-zero quadrupole moment; a non-zero second derivative is usually created by the quadrupole moment changing cyclically in time.

In the so-called transverse traceless gauge, spacetime coordinates coincide with world lines of freely-falling test masses, and the spatial gravitational wave strain

$$\mathbf{H} = h_+ \mathbf{H}^+ + h_\times \mathbf{H}^\times, \quad (2.1)$$

is written in terms of two polarisations, *plus* and *cross*, with respective amplitudes h_+ and h_\times , and respective bases \mathbf{H}^+ and \mathbf{H}^\times (see section 5.2.1).

When expressed in the transverse traceless gauge, a gravitational wave does not accelerate masses, but instead changes the measured distance

$$\int_{\text{path}} \sqrt{d\vec{x} \cdot (\mathbf{1} + \mathbf{H}) d\vec{x}} \quad (2.2)$$

between objects along a particular path. Figure 2.2 illustrates the (greatly exaggerated) change in distance, measured radially, from the origin to a ring of test masses, due to the passing of a plane gravitational wave of pure plus and pure cross polarisations. The measured distance along a particular axis is alternately stretched and shrunk at the frequency of the wave; the same perturbation, but with opposite phase, occurs along an orthogonal axis. The perturbations due to the plus and cross polarisations are rotated 45% (about the \hat{z} axis) with respect to each other.

There are many astrophysical sources of gravitational waves (Schutz & Sathyaprakash 2009, and references therein). They may be roughly divided into a number of frequency bands (Hughes 2003); sources in different frequency bands require different experimental techniques to detect them. Gravitational waves with frequencies below 10^{-13} Hz (which corresponds to wavelengths comparable to the size of the universe) are likely to have been generated during inflation, and are expected to leave imprints in the cosmic microwave background detectable by ground- and space-based experiments (Baskaran *et al.* 2010, and references therein). Between 10^{-9} and 10^{-7} Hz, a background of massive binary black holes are expected to be detectable by pulsar timing, which uses time-of-arrival delays in the observed pulses of a suite of millisecond pulsars (Lorimer 2008, and references therein). Space-based interferometric detectors, such as the planned Laser Interferometer Space Antenna (LISA), will be sensitive to gravitational waves from stellar-mass, white dwarf, and black hole binaries, radiating at frequencies between 10^{-5} and 1 Hz (Shaddock 2008, and references therein). Finally, gravitational wave sources with frequencies above 1 Hz include coalescence of compact binary stars, bursts from e.g. core collapse supernovae, rapidly spinning neutron stars, and the stochastic background (see section 5.1). These last sources are targeted by ground-based detectors, such as cryogenic bars and large-scale interferometers. The sensitive range of cryogenic bars is limited to relatively narrow bands surrounding the resonant frequency of the bar, which ranges between 700 and 900 Hz (Astone *et al.* 2003). Ground-based interferometers are generally sensitive to gravitational waves over a broad range of frequencies, typically from 1 Hz to several kHz (e.g. Abbott *et al.* 2009e).

Figure 2.3 shows a schematic diagram of a Michelson interferometer, the base configuration of current large-scale gravitational wave interferometers.

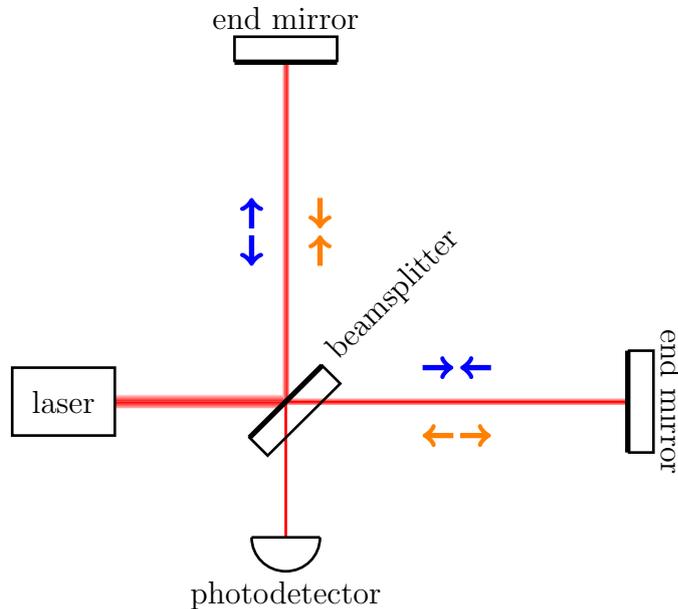


Figure 2.3: Schematic of a Michelson interferometer. The arrows show the expansion and contraction of space caused by a passing gravitational wave during the first half (blue) and second half (orange) of its cycle.

The beamsplitter and end mirrors are freely suspended, so that they approximate freely-falling test masses in the transverse traceless gauge. A laser beam is divided by the beamsplitter; each beam travels down one arm of the interferometer and is reflected back to the beamsplitter. The resulting interference pattern is sensitively dependent on the difference in the optical path lengths of each arm. If the arm lengths differ by an integer-and-a-half multiple of the laser wavelength, the returning beams destructively interfere at the beamsplitter, and no light reaches the photodetector. A gravitational wave, travelling perpendicular to the plane of the interferometer's arms, will tend to increase the measured length of one arm, and decrease the measured length of the other arm, during one half of its cycle, then reverse the perturbation during the second half (see Figure 2.2). The change in the relative lengths of the arms, and thus in the optical path length, produces a cyclical output, of the same frequency of the wave, at the photodetector.

The operation of a real kilometre-arm-length interferometer, such as LIGO, is of course a more complicated affair (Saulson 1994, Abbott *et al.* 2009e). For example, the single end mirror is commonly replaced by a Fabry-Perot cavity, where photons are bounced many times between the two cavity mirrors before being returned to the beamsplitter; this increases the effective arm-length

of the interferometer, which maximises the optical path length difference. The mirrors, beamsplitter, and other optical components hang from steel wire pendula, which mitigate seismic vibrations by ~ 4 orders of magnitude, and which themselves hang from complicated isolation stacks which reduce vibrations by a further ~ 8 orders of magnitude. All optical components are housed in an ultra-high vacuum system, the creation of which required, for example, the electrical heating to 160° C of kilometres of stainless steel tubing. Locking, the process of bringing the many optical components into the correct alignment, requires precise sensing and actuation control. The raw output of the interferometer must also be calibrated, which requires, among other things, the ability to physically manipulate the mirrors at the $\sim 10^{-18}$ m level, using electromagnetic actuators, to simulate a test gravitational wave signal. Sources of noise affecting the operation and overall sensitivity of the instrument include: the quantum statistics of the laser light, thermal noise, the temperature-dependent refractive index of the mirrors (which distorts the laser beam when the mirrors are heated – by the laser itself), electronics, wind, local traffic, tides, passing aircraft, remote earthquakes, and so on. In addition, the analysis of the calibrated data in search of faint gravitational wave signals is itself a significant challenge (see Chapter 5).

Three kilometre-scale gravitational wave observatories are currently in operation. The Laser Interferometer Gravitational-Wave Observatory (LIGO) has constructed two observatories in the United States: in Hanford, Washington, and in Livingston, Louisiana. The Hanford observatory houses one 4-km and one 2-km arm length interferometer, while the Livingston observatory houses a 4-km interferometer. During commissioning, LIGO collected science-quality data during four science runs (designated S1–S4), which were analysed by the LIGO Scientific Collaboration (LSC) to produce the first upper limits on gravitational waves from interferometric detectors. LIGO embarked on its fifth science run (S5) in November 2005; when it concluded in October 2007, LIGO had collected a year of data, coincident between its detectors, at its initial design sensitivity. During S5, the interferometers at their most sensitive reached a strain sensitivity to gravitational waves of $\sim 3 \times 10^{-23} \text{ Hz}^{-1/2}$ (Abbott *et al.* 2009e), which roughly translates to measurements on the order of 10^{-18} m. The VIRGO Collaboration has also constructed and operates a 3-km arm-length interferometer near Pisa, Italy, which is sensitive to gravitational waves at the $10^{-22} \text{ Hz}^{-1/2}$ level (Acernese *et al.* 2008).

There are also a number of sub-kilometre arm-length interferometers which continue to demonstrate advanced technologies. GEO600, a 600-m instrument near Hanover, Germany, has attained a best sensitivity of

$\sim 3 \times 10^{-22} \text{ Hz}^{-1/2}$ (Grote 2008). Located in Japan, the 300-m TAMA300 (Takahashi *et al.* 2008), and the 100-m cryogenic CLIO (Yamamoto *et al.* 2008) have respective best sensitivities of $\sim 10^{-21} \text{ Hz}^{-1/2}$ and $\sim 10^{-20} \text{ Hz}^{-1/2}$.

No detection of gravitational waves has been made to date. While the present-day initial generation of interferometers have demonstrated their ability to collect science-quality data over lengthy observation times, they are not sensitive enough to guarantee detection. A second generation of advanced interferometric detectors will feature order of magnitude improvements in sensitivity, and are expected to begin operation around the middle of the next decade. The Japanese LCGT project, and the Australian ACIGA consortium, also plan to construct advanced detectors.

2.4 Periodic gravitational waves from spinning neutron stars

Periodic gravitational waves are long-lived, quasi-monochromatic signals. They are anticipated to be generated by rapidly spinning neutron stars, provided that the star is not axially symmetric about its angular momentum vector.

In general, the spectrum of a periodic gravitational wave signal contains lines at three frequencies, which are functions of the neutron star's rotation frequency f_{rot} , and precession frequency f_{prec} (Zimmermann & Szedenits 1979, Zimmermann 1980, Van Den Broeck 2005, Vigelius & Melatos 2009a). Two lines, at $f_{\text{rot}} + f_{\text{prec}}$ and $2(f_{\text{rot}} + f_{\text{prec}})$, arise from the precessing motion, and are likely to be detectable only after improvements in sensitivity planned for advanced interferometers (Jones & Andersson 2002, Van Den Broeck 2005, Prix 2009, Abbott *et al.* 2007b). The third line, at $2f_{\text{rot}}$, is present only if the star is non-axisymmetric, i.e. if its three principal moments of inertia $I_1 < I_2 < I_3$ are all different. To date, searches for periodic gravitational waves have directly targeted only this line. The amplitudes of the plus and cross polarisations of this line are given by

$$A_+ = \frac{1}{2}h_0(1 + \cos^2 \iota), \quad A_\times = h_0 \cos \iota \quad (2.3)$$

where ι is the inclination angle of the total angular momentum vector to the line of sight. (See equation 5.3 for the relationship between these amplitudes and the amplitudes h_+ and h_\times of equation 2.1). The overall strain tensor amplitude,

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz}\epsilon}{D} f^2, \quad (2.4)$$

is given in terms of the gravitational wave frequency $f = 2f_{\text{rot}}$, the distance to the source D , the principal moment of inertia I_{zz} , and the equatorial ellipticity ϵ (Abbott *et al.* 2007b, Jaranowski *et al.* 1998). Equation 2.4 is used, along with energy conservation arguments (see section 8.2), to set indirect upper limits on the non-axisymmetric gravitational wave strain from known pulsars (e.g. Abbott *et al.* 2010), and from a Galactic population of spinning neutron stars (Abbott *et al.* 2007b, Knispel & Allen 2008).

The greatest source of uncertainty in estimating the gravitational wave strain is the possible range of ellipticities (Prix 2009, and references therein). Isolated neutron stars may be born with some non-axisymmetry, which becomes frozen into the crust as the star rapidly cools via neutrino emission (Ruderman 1969). The crust will likely crack and relax as the young neutron star spins down; the timescale over which this occurs is unclear, and is probably dependent on the maximum breaking strain of the crust (Abbott *et al.* 2007b, Prix 2009). Recent simulations by Horowitz & Kadau (2009) concluded that the breaking strain of a neutron star crust could be up to $\sim 10^{-1}$, an order-of-magnitude larger than previous estimates. Younger neutron stars (e.g. the central compact object in Cassiopeia A; see Chapters 8–9) are more likely to retain any non-axisymmetry, as there has been less time for the deformation to relax elastically or diffusively. In some circumstances, however, the crust may be melted from below by heat generated by r -mode oscillations (see below) within the star (Lindblom *et al.* 2000, Wu *et al.* 2001). Crust melting is likely to delay the formation of the crust in young, rapidly rotating neutron stars, and may very well affect the formation and retention of any non-axisymmetry.

Estimates of the ellipticities supportable by an isolated neutron star range over several orders of magnitude. The ellipticity supportable by a conventional neutron star crust is estimated by Ushomirsky *et al.* (2000) to be $\epsilon_{\text{max}} \lesssim 5 \times 10^{-7}$, using a breaking strain of 10^{-2} . Much larger ellipticities, however, can be supported by neutron stars with more exotic equations of state (Owen 2005). It is believed that the super-nuclear densities in the core may support some form of pure quark matter; proposed models include solid quark stars (Xu 2003), crystalline colour superconducting quark matter (Alford *et al.* 2004, Haskell *et al.* 2007, Lin 2007, Mannarelli *et al.* 2007, Knippel & Sedrakian 2009), and mixed quark-baryon cores (Glendenning 1992, Owen 2005). Using breaking strains of $\sim 10^{-2}$, these models estimate ellipticities of $10^{-6} \lesssim \epsilon_{\text{max}} \lesssim 10^{-4}$, and up to $\epsilon_{\text{max}} \lesssim 10^{-2}$ (Owen 2005, Haskell *et al.* 2007, Lin 2007). The ellipticity also depends on elasticity (Ushomirsky *et al.* 2000, Haskell *et al.* 2007).

In accreting neutron star systems, there are a number of possible scenarios whereby the accretion flow may create non-axisymmetric distortions. One

in particular is the subject of Chapters 3 and 4 of this thesis: the accreted plasma is channelled by the star's dipolar magnetic field onto the poles, where it builds up to form magnetically confined mountains (Melatos & Payne 2005). The pressure at the base of the mountain forces it to spread sideways, dragging with it frozen-in lines of magnetic flux.¹ Eventually, the increasing magnetic pressure exerted by the distorted field lines halts the spread of the mountain and a stable equilibrium is achieved. Because the magnetic axis is generally misaligned to the rotation axis, the system is nonaxisymmetric; Vigelius & Melatos (2009a) estimated ellipticities of $\epsilon_{\max} \lesssim 2 \times 10^{-5}$.

Another possible scenario is that the accreted matter locally compresses and heats the crust, leading to a non-uniform temperature distribution (Bildsten 1998). At crustal densities, the rate at which electrons are captured by protons to form neutrons is temperature dependent; as a result, the density increase associated with electron capture occurs at a depth dependent on the temperature. If the temperature distribution is non-axisymmetric and stable, a non-zero quadrupole moment is generated, and the ellipticity may reach up to the maximum predicted by Ushomirsky *et al.* (2000).

A strong magnetic field may also aid the generation of gravitational waves. The distortion in shape of the neutron star due to such a field has been studied by Bonazzola & Gourgoulhon (1996), Colaiuda *et al.* (2008), and Haskell *et al.* (2008); possible maximum ellipticities were found to be $\epsilon_{\max} \lesssim 10^{-6}$ – 10^{-5} at field strengths of $\sim 10^{15}$ G. Cutler (2002) showed that a large toroidal magnetic field will tend to squeeze the star at its equator to form a prolate shape, with the symmetry axis parallel to the angular momentum axis. This configuration is unstable until the symmetry axis realigns to be at right angles to the angular momentum axis, which is also optimal for gravitational wave emission (Owen 2006). The expected ellipticity depends on the magnetic field strength; for a magnetic field of 10^{15} G, one finds $\epsilon_{\max} \lesssim 10^{-6}$. Akgün & Wasserman (2008) studied the distortion due to toroidal magnetic fields of a type II superconducting star where the magnetic flux is concentrated in tubes which occupy a small fraction of the stellar volume; for magnetic fields $\sim 10^{15}$ G, they found $\epsilon \sim 10^{-9}$ – 10^{-8} . Recent studies of the neutron star magnetic fields by Braithwaite & Nordlund (2006) and Braithwaite & Spruit (2006) found the magnetic field lines form complicated toroidal geometries.

Periodic gravitational waves may also be generated by unstable fluid oscillations of the neutron star. At sufficiently high rotation rates, such that

¹ In the ideal-magnetohydrodynamic (ideal-MHD) limit, which holds if the Ohmic diffusion timescale is much greater than the accretion timescale, material cannot cross magnetic field lines. This assumption is relaxed in Vigelius & Melatos (2009b).

the rotational energy exceeds a certain fraction of the gravitational binding energy, the neutron star is deformed, for a few rotations, into a bar shape (Prix 2009, and references therein). This may produce a short-lived periodic gravitational wave signal, although its detectability is unclear.

More promising are the family of fluid oscillations known as r -modes (e.g. Andersson 1998, Kokkotas & Schmidt 1999, Stergioulas 2003, Andersson & Comer 2007), which are toroidal oscillations where the Coriolis force provides the restoring force. The frequencies of these modes are always less than the rotation frequency of the star; thus, they are prograde (i.e. rotate with the star) as seen by an inertial observer, and retrograde (i.e. counter-rotating) with respect to the star's co-rotating frame. This triggers the Chandrasekhar-Friedman-Schutz instability, which removes positive angular momentum from the star as gravitational waves. For r -modes, the instability is present at all rotation frequencies (Andersson 1998). It is expected to rapidly spin down young neutron stars (Lindblom *et al.* 1998), and produce detectable gravitational radiation many years after their formation (Owen *et al.* 1998, Bondarescu *et al.* 2009). The instability can, however, be suppressed by viscosity due to the boundary layer between core and crust (Bildsten & Ushomirsky 2000), or by hyperon interactions in the core (Lindblom & Owen 2002). Strange quark stars, however, are not subject to the same damping mechanisms (Madsen 2000).

Other fluid instabilities which may produce detectable gravitational wave signals include f -mode oscillations (e.g. Andersson *et al.* 2009) following magnetar flares, which are potentially detectable by initial LIGO (Abbott *et al.* 2008e, 2009k), and nonaxisymmetric Ekman flows excited by rotational glitches (van Eysden & Melatos 2008), which may be detectable by advanced interferometers.

It should be emphasised that the *maximum* possible ellipticities, listed above, imply nothing about the *average* ellipticities of spinning neutron stars in general, or of any one neutron star in particular. There is nothing to prevent the average ellipticity being much less than the maximum; thus, a non-detection of periodic gravitational waves from a particular neutron star cannot, at present, be used to exclude any one of the possible mechanisms described above. An all-sky search sensitive to a population of neutron stars may, in future, be able to make some statistical statement disfavouring a particular mechanism, if sufficiently stringent upper limits are obtained (and assuming no detection is made).

Determining a theoretical *minimum* ellipticity permitted by a particular mechanism would allow it to be ruled out by observational upper limits. Magnetic deformation, assuming the internal field to be of at least equal magnitude to the external field, sets a minimum ellipticity (Melatos 2000,

Cutler 2002), which is often very low: for the Crab pulsar, $\epsilon \sim 10^{-11}$ (Abbott *et al.* 2008c). In the case of accreting neutron stars, the flow of fresh accreted matter may replenish any natural decay of the non-axisymmetric distortion, and so it seems plausible that an equilibrium state, together with a detectable non-zero ellipticity, can be established. A highly accurate model of the physical processes involved would still be required to make an accurate prediction. A small step towards such an accurate model of magnetically confined mountains is taken in Chapters 3–4.

2.5 Summary

This chapter introduced the two areas of physics which are the focus of the remainder of this thesis; neutron stars, and gravitational waves. We first gave a brief history of the two fields. We introduced neutron stars, their formation in supernova remnants, pulsars, and the formation and evolution of millisecond pulsars. We then introduced gravitational waves, and efforts to detect them using kilometre-scale laser interferometers. Finally, we reviewed the physics involved in the generation of periodic gravitational waves by spinning neutron stars: an important link between the two phenomena, and of central concern to this thesis.

Chapter 3

A method of building stable mountains with sinking

In this chapter and Chapter 4, we study the formation of magnetic mountains on the surface of an accreting neutron star. The formation of these mountains and the associated burial of the star's magnetic field is important to models of the evolution of low-mass x-ray binaries and the creation of millisecond pulsars, as well as for the prospects of detecting gravitational waves from low-mass x-ray binaries. In section 3.1, we review previous work on this problem; in section 3.2, we present the numerical method used to build the magnetic mountain equilibria presented in Chapter 4.

See section 1.1 for author contributions and publications relevant to this chapter.

3.1 Magnetic burial

The magnetic dipole moment μ of a neutron star is observed to diminish in the long term as the star accretes (Taam & van den Heuvel 1986, van den Heuvel & Bitzaraki 1995), although Wijers (1997) argued that μ may also be a function of parameters other than the accreted mass M_a . The μ - M_a correlation has been ascribed to a number of physical mechanisms (Melatos & Phinney 2001, Cumming 2005). First, the magnetic field may be dissipated in the stellar crust by Ohmic decay, accelerated by heating as the accreted plasma impacts upon the star (Konar & Bhattacharya 1997, Urpin *et al.* 1998, Brown & Bildsten 1998, Cumming *et al.* 2004). Second, magnetic flux tubes may be dragged from the superconducting core by the outward motion of superfluid vortices, as the star spins down (Srinivasan *et al.* 1990, Ruderman *et al.* 1998, Konar & Bhattacharya 1999, Konenkov & Geppert

2001). Third, the magnetic field may be screened by accretion-induced currents within the crust (Bisnovatyi-Kogan & Komberg 1974, Blondin & Freese 1986, Lovelace *et al.* 2005). In particular, the field may be *buried* under a mountain of accreted plasma channelled onto the magnetic poles. When M_a is large enough, the mountain spreads laterally, transporting the polar magnetic flux towards the equator (Hameury *et al.* 1983, Romani 1990, Brown & Bildsten 1998, Cumming *et al.* 2001, Melatos & Phinney 2001, Choudhuri & Konar 2002, Payne & Melatos 2004, Zhang & Kojima 2006, Payne & Melatos 2007, Vigelius & Melatos 2008, 2009b).

Payne & Melatos (2004) computed the unique sequence of self-consistent, ideal-magnetohydrodynamic (ideal-MHD) equilibria that describes the formation of a polar mountain by magnetic burial as a function of M_a . They found that the accreted mountain is confined by the equatorially compressed magnetic field, which was unaccounted for in previous calculations, and that $10^{-5}M_\odot$ must be accreted to lower μ by 10%. Surprisingly, mountains are stable with respect to axisymmetric ideal-MHD perturbations; they oscillate globally in a superposition of acoustic and Alfvén modes but remain intact due to magnetic line-tying at the stellar surface (Payne & Melatos 2007). The same equilibria are susceptible to nonaxisymmetric, Parker-like instabilities (specifically the gravitationally driven, undular sub-mode), but the instability preserves a polar mountain when it saturates, despite reducing the mass ellipticity by $\sim 30\%$ (Vigelius & Melatos 2008). Recently, Vigelius & Melatos (2009b) considered resistive effects. They found that the mountain does not relax appreciably for realistic resistivities over the lifetime of a low- or high-mass X-ray binary, either by global diffusion, resistive g-mode instabilities, or reconnection in the equatorial magnetic belt. The Hall drift, which exerts a destabilising influence in isolated neutron stars (e.g. Rheinhardt & Geppert 2002), is unlikely to be important in accreting neutron stars due to crustal impurities (Cumming *et al.* 2004, Cumming 2005). Since non-ideal-MHD effects (e.g. resistivity, the Hall drift) have been found to be small when modelling the formation of a magnetic mountain on an accreting neutron star, the assumption that the neutron star plasma obeys ideal-MHD physics is a reasonable one, and is made throughout this work.

The investigations outlined in the previous paragraph suffer from two limitations. First, the mountain is assumed to rest upon a rigid surface. Under this assumption, the accreting plasma cannot sink into the stellar crust. This is unrealistic. During magnetic burial, frozen-in magnetic flux is redistributed slowly within the neutron star by the accreted plasma, as it sinks beneath the surface and spreads laterally. Choudhuri & Konar (2002) showed that the time-scale and end state of burial are tied to these slow interior motions. Second, the accreted plasma is assumed to satisfy an isothermal

equation of state. This is an accurate model only for neutron stars with low accretion rates $\dot{M}_a \lesssim 10^{-10} M_\odot \text{yr}^{-1}$; the thermodynamics of neutron stars accreting near the Eddington limit ($\sim 10^{-8} M_\odot \text{yr}^{-1}$) is more complicated, with a depth-dependent adiabatic index (Brown & Bildsten 1998, Brown 2000). The equation of state affects the growth rate of Parker-like instabilities (Kosiński & Hanasz 2006). In this work (Wette *et al.* 2010), we seek to overcome the first limitation.

3.2 Growing a realistically sized mountain by injection

In order to investigate how a magnetically confined mountain sinks into the stellar crust, we need a numerical method capable of building a stable mountain, with a realistic M_a , on top of a fluid base. The approach we take builds upon previous work by Payne & Melatos (2004, 2007) and Vigelius & Melatos (2008, 2009a). Here, as a service to the reader, we briefly recapitulate the physical arguments and key results from these previous papers, with references to the relevant sections and equations.

3.2.1 Previous work

In Payne & Melatos (2004), axisymmetric magnetic mountain equilibria are computed by solving an elliptic partial differential equation: the Grad-Shafranov equation describing hydromagnetic force balance in axisymmetric geometry [Payne & Melatos (2004), section 2.1 and equation (12)]. The calculation ensures that the mass-magnetic flux distribution $\partial M/\partial\psi$ is treated self-consistently: the final $\partial M/\partial\psi$ is equal to the initial $\partial M/\partial\psi$ together with the mass-flux distribution of the accreted matter, $\partial M_a/\partial\psi$, which is characterised by the parameter ψ_a [Payne & Melatos (2004), section 2.2 and equation (13)]. In the limit of small M_a , the final equilibrium flux solution is characterised by the ratio M_a/M_c , where the characteristic mass $M_c \propto M_\star R_\star^2 B_\star^2$ is the accreted mass required to halve μ [Payne & Melatos (2004), section 3.2 and equation (30); Payne & Melatos (2007), section 2.2 and equation (3)]. The characteristic mass contains the dependence of the equilibrium solution on the parameters of the neutron star, in particular the magnetic field strength B_\star . The Grad-Shafranov equilibria are computed using an iterative numerical solver [Payne & Melatos (2004), section 3.3]; this approach only converges numerically for low accreted masses $M_a \leq M_c \approx 10^{-4} M_\odot$ [Payne & Melatos (2004), section 3.4], and it fails to accommodate a fluid interior

within its fixed-boundary framework.¹

In Payne & Melatos (2007) and Vigelius & Melatos (2008), Grad-Shafranov equilibria are loaded into ZEUS, a multi-purpose, time-dependent, ideal-MHD solver (Stone & Norman 1992a,b, Hayes *et al.* 2006), and further evolved in axisymmetric (Payne & Melatos 2007) and three-dimensional geometries (Vigelius & Melatos 2008). The characteristic mass M_c is used to reduce the length scales of the simulated neutron star to circumvent numerical difficulties and render the simulations computationally tractable [Payne & Melatos (2007), section 3.3; Vigelius & Melatos (2008), section 2.3 and equation (6), and section 4.6]. Two approaches are explored to augmenting the mass of a Grad-Shafranov mountain, up to $M_a \lesssim 5.6M_c$: in the first approach, additional matter is injected through the outer boundary along the polar flux tube $0 \leq \psi \leq \psi_a$ [Payne & Melatos (2007), section 4.2]; in the second approach, the density of the mountain is uniformly increased at every point, while the magnetic field is preserved [Payne & Melatos (2007), section 4.4]. A plausible attempt to extend this latter approach to include sinking is outlined in Appendix 3.A.1; ultimately this attempt proved unsuccessful, and was abandoned. Instead, the method presented in this section uses ZEUS-MP (Hayes *et al.* 2006) to build magnetic mountain equilibria from scratch; this approach was first proposed in Vigelius & Melatos (2009a).

3.2.2 Outline of the method

The setup of the simulations presented here is described schematically in Figure 3.1. Three numerical experiments are performed: growing onto a hard surface, growing onto a soft surface by injecting matter from below, and repeating the latter experiment by injecting matter at some altitude.

Throughout this work, we adopt the viewpoint that the accreted matter and the mountain are one and the same; the accreted mass and the mass of the mountain are identical and are both denoted by M_a . This is a matter of terminology, not physics. There is no “hard edge” to the mountain; matter is accreted on all flux surfaces $0 \leq \psi \leq \psi_*$ [see equation 3.6 in section 3.2.4], not just on the polar cap $0 \leq \psi \leq \psi_a$, which contains $\sim 63\%$ of M_a . Under the assumption of ideal MHD, matter cannot spread across flux surfaces, i.e. there is no Ohmic diffusion. We also do not model the accreted matter once it has sunk beyond the crust, as do e.g. Choudhuri & Konar (2002); see the

¹ In Payne & Melatos (2004), the Grad-Shafranov equation is solved subject to Dirichlet and Neumann conditions at *fixed* boundaries. Mathematically, one can formulate a well-posed boundary-value problem for the Grad-Shafranov equation in the presence of a *free* boundary, e.g. the sinking base of a mountain; in practice, however, this is an extremely difficult problem to solve.

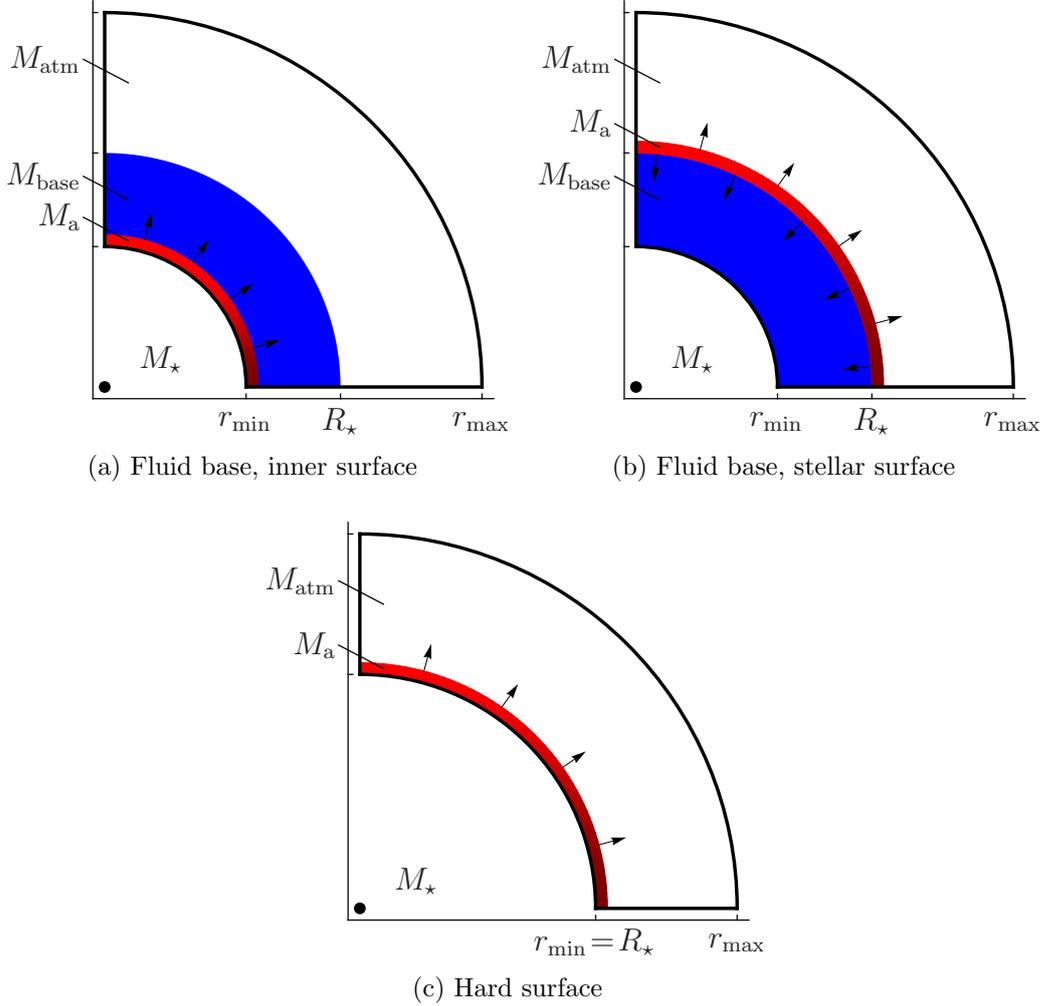


Figure 3.1: Diagrams illustrating schematically three mountain growth scenarios. The simulation region is bounded by $r_{\min} \leq r \leq r_{\max}$, $0 \leq \theta \leq \pi/2$ (thick lines), and represents a quadrant of the star. Boundary conditions assume symmetry about $\theta = 0$ and reflection at $\theta = \pi/2$. The surface of the star is located at $r = R_{\star}$. Three sub-regions are identified. The fluid interior beneath the surface, containing mass M_{base} , is shaded blue. The region where the mountain mass M_{a} is injected into the simulation is shaded in a light to dark red gradient; the colour intensity is proportional to the injected flux [see equation 3.6] as a function of θ . The outer atmosphere of the star, containing mass M_{atm} , is unshaded. A central gravitational point source is labelled with its mass M_{\star} . For mountains grown on a fluid base, M_{a} can be injected (a) at the inner boundary $r = r_{\min}$, or (b) at the stellar surface $r = R_{\star}$. For mountains grown on a hard surface, (c) there is no fluid interior; the inner boundary is identical to the stellar surface $r = r_{\min} = R_{\star}$ ($M_{\text{base}} = 0$). See the text in section 3.2.4.

discussion in section 4.3.

To simulate accretion, we inject matter from below, through the inner boundary of the simulation at $r = r_{\min}$. One might expect a realistic simulation of accretion to add matter from above, through the outer boundary $r = r_{\max}$. The two scenarios are, however, equivalent in ideal MHD; the magnetic field is frozen into the fluid, which is thus constrained to move along lines of magnetic flux. Provided that the simulation reaches equilibrium, it becomes inconsequential, with respect to ideal MHD, through which end of a flux tube matter is added. This is because matter cannot cross flux surfaces in ideal MHD, so the mass column $dM(\psi)$ between ψ and $\psi + \delta\psi$ adjusts to reach the same hydrostatic radial profile *in equilibrium*, whether it enters slowly from below or falls slowly from above. In the presence of gravity, which (in the case of a sinking mountain) induces steep density gradients in the fluid base, the results to be presented in section 4.2 confirm that this situation remains true; two different injection scenarios (described below) give ellipticities consistent to within 10%. There remains, however, the subtle and difficult question of irreversible magnetic reconnection at the grid corners, which remains unresolved (see the discussion in section 4.2.2).

In practice, it is advantageous to add matter through the inner boundary, because we wish to inject along particular flux tubes, and this is easiest to do at $r = r_{\min}$, where the magnetic footprints are fixed in place (unlike at $r = r_{\max}$). This constraint, known as magnetic line tying, contributes to the stability of the mountain (Goedbloed & Poedts 2004, Vigelius & Melatos 2008). It is well justified physically, provided that r_{\min} lies deep enough within the star, so that the fluid base (and frozen-in magnetic flux) remains relatively stationary, and is not significantly perturbed by the spreading and sinking of the mountain. This is the case if the mass M_{base} of the fluid base, initially in the region $r_{\min} < r < R_*$, is much greater than M_a . To confirm that the mountain does not greatly push the crustal material, we first calculate the fraction of M_{base} contained in each grid cell, and then determine the change in this quantity between the initial and final times of the simulation; this gives the change in the spatial distribution of M_{base} over the simulation, as a function of the grid cell. For all simulations with sinking, the median change in M_{base} , over all grid cells, is on average $\sim 10\%$; thus, the distribution of the fluid base does not change much during accretion. Recent molecular dynamics simulations of crystalline neutron matter, which predict a high breaking strain ~ 0.1 (Horowitz & Kadau 2009), also support the line-tying hypothesis.

When a mountain is grown onto a fluid base M_{base} , a difficulty arises. ZEUS-MP models a single fluid, with a unique velocity field (Hayes *et al.* 2006); there is no facility for simulating the movement of one fluid, the in-

jected mountain, with respect to another fluid, the stationary crust.² We are left with two alternatives: to assign the same velocity to the injected mountain and the crust (the behaviour of ZEUS-MP’s “inflow” boundary condition), or to assign a negligible or zero velocity to the injected mountain, in order to keep the base stationary. In the first case, ZEUS-MP fails catastrophically for desirable values of the injection velocity ($\gtrsim 5\%$ of the escape velocity). In the second case, which we study in section 4.2, mountains remain subterranean and never rise to the stellar surface $r = R_\star$. As a check, therefore, we examine two scenarios: injection at $r = r_{\min}$ and $r = R_\star$. We show in section 4.2 that the results in both scenarios are quantitatively alike, confirming their robustness.

3.2.3 Initial setup

The initial setup of our simulations closely follows Payne & Melatos (2007) and Vigelius & Melatos (2008). The setup of ZEUS-MP (see bibliography) is through a set of parameters which control: the geometry of the problem, the physical phenomena to be modelled (e.g. MHD, gravity), the simulation grid and its boundary conditions, the equation of state, and the choice of timestep. Appropriate values for these parameters are given in Payne & Melatos (2007), section 3 and appendix A1, and in Vigelius & Melatos (2008), sections 2.2–2.3 and appendix A.

To avoid numerical difficulties with steep magnetic field gradients, we simulate a scaled-down neutron star, where the mass M_\star and radius R_\star are artificially reduced, while the hydrostatic scale height $h_0 = c_s^2 R_\star^2 / GM_\star$ is kept constant (Payne & Melatos 2007). The down-scaling transformation preserves the equilibrium shape of the mountain exactly in the small- M_a limit (Payne & Melatos 2004, 2007) and has been validated approximately for $M_a \lesssim 20M_c$ (Vigelius & Melatos 2008). We use dimensionless units within ZEUS-MP, setting the isothermal sound c_s and gravitational constant G to unity, and adopting h_0 as the unit of length. Table 3.1 explains how to convert between an astrophysical neutron star, the scaled-down model, and dimensionless ZEUS-MP units.

The simulations are performed on an axisymmetric rectangular grid with N_r cells spaced logarithmically in r , and $N_\theta = 64$ cells spaced linearly in θ . The logarithmic spacing in r is determined by $\Delta r_{N_r-1} / \Delta r_0$, the ratio of the maximum to minimum radial grid spacing, as follows. The $N_r + 1$ radial cell boundaries $r_{\min} = r_0, r_1, r_2, \dots, r_{N_r} = r_{\max}$ are given by $r_{n+1} = r_n + \Delta r_n$,

² ZEUS-MP can track the concentrations of comoving components within the same fluid; we exploit this in section 3.2.4.

Table 3.1: Conversion of physical quantities into dimensionless variables in the simulations. Physical quantities are first converted to their values in the scaled-down model by multiplying by $[a(\text{Simulation})/a(\text{Astrophysical neutron star})]^n$, where $a = R_*/h_0$ parameterises the curvature down-scaling, and n is listed in column 4. Scaled-down physical quantities are then reexpressed in the dimensionless units of ZEUS-MP according to column 6. The table is divided into three horizontal parts containing: stellar parameters (Payne & Melatos 2004), simulation control parameters (see Figure 3.1), and simulation outputs.

Quantity	Symbol	Astrophysical neutron star	Down- scaling index n	Simulation	ZEUS-MP dimensionless unit
scaling ratio	a	1.9×10^4	none	50	none
stellar mass	M_*	$1.4M_\odot$	2	$10^{-5}M_\odot$	$M_0 = c_s^2 h_0 / G = 8.1 \times 10^{24}$ g
stellar radius	R_*	10^6 cm	1	2.7×10^3 cm	$h_0 = c_s^2 R_*^2 / GM_* = 54$ cm
stellar magnetic field	B_*	10^{12} G	none		$B_0 = c_s^2 / G^{1/2} h_0 = 7.2 \times 10^{17}$ G
isothermal sound speed	c_s	10^8 cm s $^{-1}$	none		c_s
critical mass	M_c	$1.2 \times 10^{-4}M_\odot$	4	$6.1 \times 10^{-15}M_\odot$	M_0
inner boundary	r_{\min}		1	see Table 4.1	h_0
outer boundary	r_{\max}		1	$1.2R_*$	h_0
accreted mass	M_a		4	see Table 4.1	M_0
mass of outer atmosphere	M_{atm}		4	$5 \times 10^{-6}M_\odot$	M_0
mass of fluid base	M_{base}		4	$10M_a$	M_0
mountain density	ρX_a		1		$\rho_0 = M_0/h_0^3 = 5.2 \times 10^{19}$ g cm $^{-3}$
magnetic field	\mathbf{B}		none		B_0
ellipticity	ϵ		2		none
time	t		none		$t_0 = h_0/c_s = 5.4 \times 10^{-7}$ s

where

$$\sum_{n=0}^{N_r-1} \Delta r_n = r_{\max} - r_{\min}, \quad (3.1)$$

$$\frac{\Delta r_{n+1}}{\Delta r_n} = \left(\frac{\Delta r_{N_r-1}}{\Delta r_0} \right)^{1/(N_r-1)}. \quad (3.2)$$

The values of $\Delta r_{N_r-1}/\Delta r_0$ used in the simulations presented in Chapter 4 are given in Table 4.1. They are chosen large enough to concentrate grid resolution near the inner boundary, but small enough to ensure reasonable run times. We set $r_{\max} = 1.2R_\star = 60h_0$ to give the mountain ample room to expand without meeting the outer boundary, and stipulate reflecting boundary conditions at $\theta = 0$ and $\theta = \pi/2$, “inflow” boundary conditions at $r = r_{\min}$, and “outflow” boundary conditions at $r = r_{\max}$; more details can be found in Payne & Melatos (2007). The magnetic field is initially that of a dipole, and B_\star is its magnitude at the polar surface.

A gravitational point source M_\star is placed at $r = 0$, and self-gravity is ignored. The density field is initialised to be the static atmosphere of an isothermal fluid with no self-gravity:

$$\rho(t = 0, r) = \rho_{R_\star} \exp \left[\frac{GM_\star}{c_s^2} \left(\frac{1}{r} - \frac{1}{R_\star} \right) \right]. \quad (3.3)$$

Ideally the region $r > R_\star$ should start evacuated, but ZEUS-MP requires the density to be nonzero everywhere, so we set $M_{\text{atm}} = 5 \times 10^{-6} M_\odot$ (approximately 4% of the mass of the smallest mountain; see Table 4.1). Integrating equation 3.3 over the region $r > R_\star$ (see Figure 3.1) fixes the density at the stellar surface ρ_{R_\star} in terms of M_{atm} . In contrast, we require the mass of the fluid base M_{base} (when a soft surface is being modelled) to be much larger than the mass of the mountain, as discussed in section 3.2.2. In all runs, we choose $M_{\text{base}}/M_a \approx 10$. Integrating equation 3.3 over the region containing M_{base} then fixes r_{\min} .

3.2.4 Injection procedure

ZEUS-MP’s “inflow” boundary condition permits injection at the edge of the simulation volume. To enable injection at $r = R_\star$, as in Figure 3.1b, we implement a more flexible custom procedure, and use the built-in “inflow” condition only to tie the magnetic flux at $r = r_{\min}$. We describe the procedure briefly below; further details are in Appendix 3.A.2.

We inject mass M_a into an injection region $r_{\text{inj}} < r < r_{\text{inj}} + \delta r_{\text{inj}}$, $0 < \theta < \pi/2$, over a time interval $0 < t < T_a$. (We set $\delta r_{\text{inj}} = 0.1h_0$ throughout.)

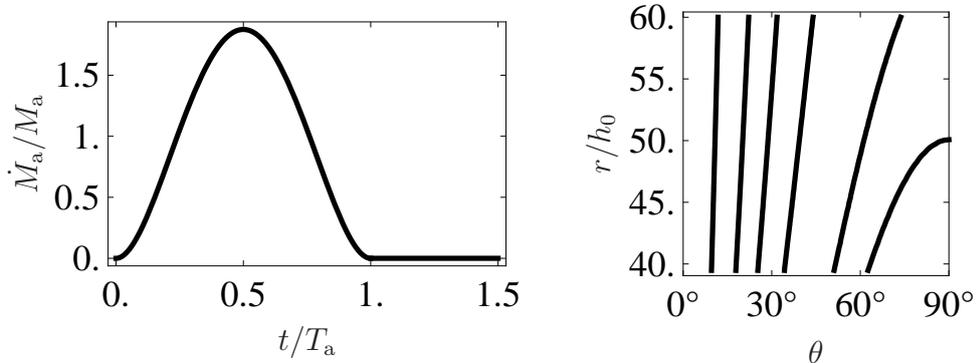


Figure 3.2: (left) The accretion rate $\dot{M}_a(t)$, given by equation 3.5. (right) The initial mass-flux distribution $\partial M/\partial \psi$, given by equation 3.6. Contours are at (right to left) 0.05, 0.1, 0.3, 0.5, 0.7, and 0.9 of the maximum.

The flux of accreted matter at time t entering a point (r, θ) in the injection region is given by

$$\frac{\partial^3 M_a}{\partial t \partial r \partial \theta}(t, r, \theta) \propto \dot{M}_a(t) \frac{\partial M_a}{\partial \psi}(r, \theta). \quad (3.4)$$

where we choose

$$\dot{M}_a(t) \propto t^2(T_a - t)^2, \quad (3.5)$$

and

$$\frac{\partial M_a}{\partial \psi}(r, \theta) \propto \exp(-bR_* r^{-1} \sin^2 \theta). \quad (3.6)$$

The normalisation of equation 3.4 is chosen so that, for each simulation, the mass of the mountain is equal to M_a at time $t = T_a$, i.e. $M_a(t = T_a) \equiv M_a$. After time $t = T_a$, no further mass is added, but we evolve the system up to $t = t_{\max} = 1.5T_a$ to test the stability of the mountain obtained.

Equation 3.5 determines the rate of accretion; it is plotted in Figure 3.2 (left). The functional form was chosen to ensure numerical stability in ZEUS-MP, and has no particular astrophysical justification, except to ensure that a mountain builds up to its target mass smoothly over the time scale T_a . For this reason, it is a smooth bell-shaped function, designed to avoid any discontinuity in the accretion rate, which might excite undesired oscillations in the fluid or provoke numerical instabilities.

Equation 3.6 gives a mass-flux distribution consistent with that of Payne & Melatos (2004); it is plotted in Figure 3.2 (right). It does not attempt to model the interaction of the accreted matter with the magnetosphere, from

which the mass-flux distribution would originate; instead, it is chosen such that the majority ($\sim 63\%$) of the accreted matter falls on the polar cap $0 \leq \psi \leq \psi_a$. The parameter $b = \psi_*/\psi_a = 3$ determines the polar cap radius $R_* \sin^{-1}(b^{-1/2})$. It is determined astrophysically by disk-magnetosphere force balance, and is related to the stellar magnetic field via $b \propto \psi_* \propto B_*$ (Payne & Melatos 2004). In this work, however, we treat b (and therefore B_*) as a free parameter, and do not attempt a self-consistent solution of the disk-magnetosphere system (e.g. Romanova *et al.* 2008). With this freedom, b is chosen unrealistically large to preserve numerical stability (Payne & Melatos 2007); more realistic values for b are investigated in Payne & Melatos (2004).

We use ZEUS-MP's multi-species tracking facility (Hayes *et al.* 2006) to record, throughout the simulation, the fraction of the density, $0 \leq X_a(t, r, \theta) \leq 1$, that originates from accretion (i.e. added at $t > 0$ via the injection procedure), as opposed to from the initial configuration at $t = 0$. This allows us to track the spread of the mountain independently of the motion of the remaining (displaced) stellar matter.

We require that the mountain grows quasistatically, in the sense that the accretion timescale T_a is always much greater than $t_{\text{Alfvén}}$, the characteristic pole-equator crossing time of an Alfvén wave. Following Vigelius & Melatos (2008), we compute the crossing time at $t = 0$, $r = R_*$: from the Alfvén speed $v_{\text{Alfvén}} = B_*/(4\pi\rho_{R_*})^{1/2} \approx 0.2c_s$ (see Table 1), we obtain $t_{\text{Alfvén}} = \pi R_*/(2v_{\text{Alfvén}}) \approx 400t_0$. The condition $t_{\text{Alfvén}} \ll T_a$ is verified by comparison with the values for T_a listed in Table 4.1. The condition also implies that the magnetostatic limit always holds: the ratio $B_*M_a/\dot{M}_a \gg \pi^{3/2}R_*\rho_{R_*}^{1/2} \approx 2 \times 10^8 \text{ G s}$, and from Tables 1 and 2, $B_*M_a/\dot{M}_a \approx B_*T_a \gtrsim 3 \times 10^9 \text{ G s}$.

For mountains grown on a hard surface, we additionally set the velocity $\mathbf{v}(t, r, \theta)$ within the injection region, such that the accreted matter is always given a fixed speed $v_{\text{inj}} = 10^{-4}c_s$ parallel to the magnetic field $\mathbf{B}(t, r, \theta)$. The value of v_{inj} should be a small fraction of the escape speed $v_{\text{esc}} \approx 4.1c_s$, so there is negligible mass lost through the outer boundary (see section 4.1.2). We find that setting the direction of the injection velocity \mathbf{v} carefully is critical to ZEUS-MP's numerical stability close to the injection region, but does not otherwise affect the simulations.

3.3 Summary

This chapter reviewed the physics of magnetic burial, and presented a numerical method capable of building stable magnetic mountains with the possibility of sinking. We reviewed previous work on the numerical simulation of magnetic mountains, discussed our strategy of injecting the mountain from

below, and presented the initial setup of the simulations and the custom injection procedure. Chapter 4 will use this method to simulate sinking magnetic mountains.

3.A Additional material

This appendix includes additional material pertaining to this chapter.

3.A.1 Matching a fluid base to a Grad-Shafranov mountain

We attempted to incorporate a fluid base into the framework of Payne & Melatos (2007) and Vigelius & Melatos (2008) in the following ad-hoc manner. Starting with a Grad-Shafranov equilibrium loaded into ZEUS-MP, we extend the inner simulation boundary, initially at $r = R_*$, inwards to create a region $r_{\min} < r < R_*$, containing the fluid base. The magnetic field \mathbf{B} in this region is initialised to a dipole. At $r = R_*$, B_r matches perfectly, but B_θ is discontinuous [see Figure 2 of Payne & Melatos (2004)]. The initial density $\rho(t = 0, r, \theta)$ is chosen to match the Grad-Shafranov density $\rho_{\text{GS}}(r, \theta)$ at $r = R_*$, and to match an isothermal, non-self-gravitating profile within $r < R_*$. A number of ad-hoc choices of $\rho(t = 0, r, \theta)$ were tried, e.g. the maximum of $\rho_{\text{GS}}(R_*, \theta)$ and $\rho'_{R_*} \exp[GM_*(r^{-1} - R_*^{-1})/c_s^2]$, with $\rho'_{R_*} = \min_\theta \rho_{\text{GS}}(R_*, \theta)$. When the combined Grad-Shafranov mountain and fluid base are evolved in ZEUS-MP, the results are undesirable. Except when $R_* - r_{\min} \ll h_0$, the fluid base is sufficiently far from equilibrium to completely disrupt the Grad-Shafranov mountain, which collapses over a short time-scale $\sim t_0$.

3.A.2 Custom injection procedure

We add a new subroutine to ZEUS-MP which is called at the beginning of each time-step δt . Within the subroutine, the density $\rho(t, r, \theta)$, mountain concentration $X_a(t, r, \theta)$, and velocity $\mathbf{v}(t, r, \theta)$ of a grid cell within the injec-

tion region (at point (r, θ) with size $\delta r \times \delta \theta$) are updated, as follows:

$$\rho(t + \delta t, r, \theta) = \rho(t, r, \theta) + \delta \rho(t, r, \theta), \quad (3.7)$$

$$X_a(t + \delta t, r, \theta) = \frac{\rho(t, r, \theta) X_a(t, r, \theta) + \delta \rho(t, r, \theta)}{\rho(t + \delta t, r, \theta)}, \quad (3.8)$$

$$\begin{aligned} \mathbf{v}(t + \delta t, r, \theta) &= v_{\text{inj}} \frac{\mathbf{B}(t, r, \theta)}{|\mathbf{B}(t, r, \theta)|} X_a(t + \delta t, r, \theta) \\ &+ \mathbf{v}(t, r, \theta) [1 - X_a(t + \delta t, r, \theta)]. \end{aligned} \quad (3.9)$$

The density increment is given by

$$\delta \rho(t, r, \theta) = \frac{M_a}{2\pi \delta r \delta \theta} \mathcal{I}(t, r, \theta); \quad (3.10)$$

the factor of 2π comes from the size of the grid cell in the ϕ dimension. The function

$$\mathcal{I}(t, r, \theta) = \frac{1}{\mathcal{N}} \int_{t_1}^{t_2} dt \int_{r_1}^{r_2} dr r^2 \int_{\theta}^{\theta + \delta \theta} d\theta \sin \theta \frac{\partial^3 M_a}{\partial t \partial r \partial \theta}(t, r, \theta) \quad (3.11)$$

integrates the injected flux given by equation (3.4); the constant

$$\mathcal{N} = \int_0^{T_a} dt \int_{r_{\text{inj}}}^{r_{\text{inj}} + \delta r_{\text{inj}}} dr r^2 \int_0^{\pi/2} d\theta \sin \theta \frac{\partial^3 M_a}{\partial t \partial r \partial \theta}(t, r, \theta) \quad (3.12)$$

ensures the correct normalisation. The times

$$t_1 = \min(t, T_a), \quad (3.13)$$

$$t_2 = \min(t + \delta t, T_a), \quad (3.14)$$

give the intersection of the current time-step with the injection time interval, and the radii

$$r_1 = \min[\max(r, r_{\text{inj}}), r_{\text{inj}} + \delta r_{\text{inj}}], \quad (3.15)$$

$$r_2 = \min[\max(r + \delta r, r_{\text{inj}}), r_{\text{inj}} + \delta r_{\text{inj}}], \quad (3.16)$$

give the intersection of the grid cell with the injection region.

Chapter 4

Sinking of a mountain on an accreting neutron star

In Chapter 3, we reviewed previous work on the problem of magnetic burial and the formation of magnetic mountains, and presented a numerical method capable of building magnetic mountain equilibria, and that can incorporate sinking. In this chapter, we present the results of simulations, using this method, of mountains grown onto a hard surface and sinking into a soft fluid base. In section 4.1, we tabulate the simulations and perform a number of validation checks. In section 4.2, the structure of mountains grown on hard and soft surfaces are compared, to determine the role of sinking, and the resulting mass quadrupole moment is evaluated as a function of M_a for hard and soft surfaces. A comparison with the results of Choudhuri & Konar (2002) is presented in section 4.3; the implications for gravitational wave emission from rapidly rotating accretors (e.g. low-mass X-ray binaries) are discussed in section 4.4.

See section 1.1 for author contributions and publications relevant to this chapter.

4.1 Simulations of magnetic mountains with sinking

Table 4.1 lists the parameters of the simulations presented in this chapter. Mountains grown on a hard surface are labelled $\mathcal{H}(M_a/M_c)$. Mountains grown on a fluid base are labelled $\mathcal{S}(r_{\text{inj}}, M_a/M_c)$, where the injection radius r_{inj} may be either r_{min} or R_* . The parameters of each run are chosen to grow a mountain with a particular target mass, M_a . We choose four values for

Table 4.1: Simulations of magnetic mountains presented in this chapter. The accompanying parameters are: the target accreted mass M_a , in units of M_c and M_\odot ; the number of grid cells in the r direction N_r ; the radius of the inner radial boundary r_{\min} ; the ratio of the maximum to minimum radial grid spacing $\Delta r_{N_r-1}/\Delta r_0$; the injection radius r_{inj} ; the injection velocity v_{inj} ; the injection time T_a ; and the total (successfully completed) simulation time t_{max} .

Simulation	$\frac{M_a}{M_c}$	$\frac{M_a}{M_\odot}$	N_r	$\frac{r_{\min}}{h_0}$	$\frac{\Delta r_{N_r-1}}{\Delta r_0}$	r_{inj}	$\frac{v_{\text{inj}}}{c_s}$	$\frac{T_a}{t_0}$	$\frac{t_{\text{max}}}{T_a}$
$\mathcal{H}(1)$	1	1.2×10^{-4}	64	50.0	200	$r_{\min} = R_\star$	10^{-4}	5×10^3	1.5
$\mathcal{H}(10)$	10	1.2×10^{-3}	64	50.0	200	$r_{\min} = R_\star$	10^{-4}	5×10^3	1.5
$\mathcal{H}(10^2)$	100	1.2×10^{-2}	64	50.0	200	$r_{\min} = R_\star$	10^{-4}	5×10^3	1.5
$\mathcal{H}(10^3)$	1000	1.2×10^{-1}	64	50.0	200	$r_{\min} = R_\star$	10^{-4}	5×10^3	1.5
$\mathcal{S}(r_{\min}, 1)$	1	1.2×10^{-4}	96	44.7	306	r_{\min}	0	5×10^3	1.5
$\mathcal{S}(r_{\min}, 10)$	10	1.2×10^{-3}	112	42.8	344	r_{\min}	0	5×10^3	1.5
$\mathcal{S}(r_{\min}, 10^2)$	100	1.2×10^{-2}	120	41.1	378	r_{\min}	0	5×10^3	1.5
$\mathcal{S}(r_{\min}, 10^3)$	1000	1.2×10^{-1}	128	39.5	410	r_{\min}	0	5×10^3	1.5
$\mathcal{S}(R_\star, 1)$	1	1.2×10^{-4}	96	44.7	306	R_\star	0	5×10^3	1.5
$\mathcal{S}(R_\star, 10)$	10	1.2×10^{-3}	112	42.8	344	R_\star	0	1×10^4	1.5
$\mathcal{S}(R_\star, 10^2)$	100	1.2×10^{-2}	120	41.1	378	R_\star	0	2×10^4	1.5
$\mathcal{S}(R_\star, 10^3)$	1000	1.2×10^{-1}	128	39.5	410	R_\star	0	8×10^4	0.35

Table 4.2: Important physical parameters of accreting neutron stars (top part), and a summary of the results of the simulations presented in this work (bottom part).

Quantity	Value/Range	Reference
accreted mass	$10^{-4} - 0.8M_{\odot}$	1, 2, 4
accretion timescale	$10^4 - 10^6$ yr	1, 7
density of crust	$10^9 - 10^{14}$ g cm $^{-3}$	5, 10
depth of crust	~ 1000 m	5, 10
initial magnetic field	$10^{12} - 10^{13}$ G	3, 8, 9
temperature	$10^8 - 10^9$ K	6, 10
ellipticity	$5 \times 10^{-5} - 2 \times 10^{-4}$	11
effect of sinking	ϵ reduced by 25 – 60 %	11

References: 1. Taam & van den Heuvel (1986); 2. van den Heuvel & Bitzaraki (1995); 3. Hartman *et al.* (1997); 4. Wijers (1997); 5. Brown & Bildsten (1998); 6. Brown (2000); 7. Cumming *et al.* (2001); 8. Arzoumanian *et al.* (2002); 9. Faucher-Giguère & Kaspi (2006); 10. Chamel & Haensel (2008); 11. This work.

M_a in the range $10^{-4} - 10^{-1}M_{\odot}$. These values are chosen to demonstrate the ability of the injection procedure to generate stable mountains over a wide range of masses. This range also encompasses the range of M_a of real accreting neutron stars (see Table 4.2). The main source of uncertainty is the accretion efficiency (van den Heuvel & Bitzaraki 1995), which may be as low as $\sim 5\%$ (Tauris *et al.* 2000); this is reflected in the chosen range of M_a .

Simulations were performed on the **xe** cluster of the Australian NCI National Facility (see bibliography). The CPU time required for each run was, on average, $\sim 10^{-2}$ seconds per grid cell per unit t_0 of simulation time. For $\mathcal{S}(R_{\star})$ -type simulations, one must scale T_a with M_a to prevent numerical instabilities. Even so, run $\mathcal{S}(R_{\star}, 10^3)$ does not complete; ZEUS-MP aborts at $t \approx 0.35T_a$, when the adaptive time-step shrinks below its allowed minimum. Figure 4.1 shows that this behaviour arises when B_r diverges at $r \lesssim r_{\text{inj}}$ along the boundary $\theta = 0$: for $M_a = 10^2M_c$, B_r threatens to break out for $t \lesssim 0.5T_a$ but ultimately settles down to the equilibrium configuration before $t = T_a$, whereas for $M_a = 10^3M_c$, it grows uncontrollably up to the time of failure.

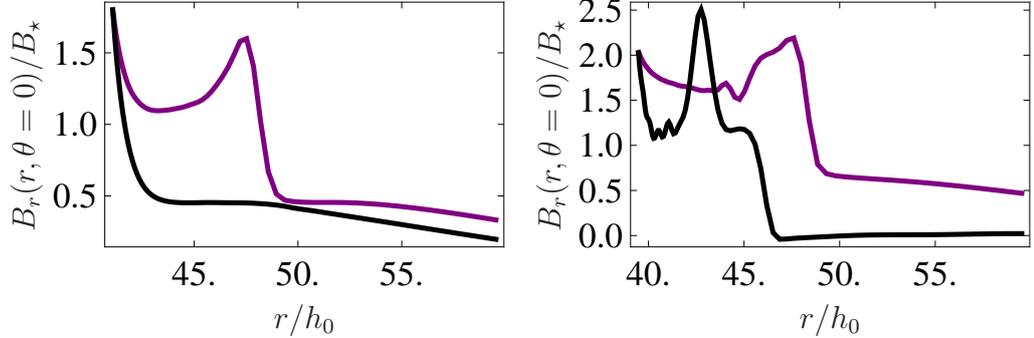


Figure 4.1: Radial component of the magnetic field B_r along $\theta = 0$: for simulation $\mathcal{S}(R_*, 10^2)$ (left), at $t = 0.45T_a$ (purple) and $t = 0.85T_a$ (black); and for simulation $\mathcal{S}(R_*, 10^3)$ (right), at $t = 0.24T_a$ (purple), and its time of failure $t = 0.35T_a$ (black).

4.1.1 Verification

We first check that, for each mountain, (i) we accumulate the correct total mass M_a , with minimal loss through the outer boundary; (ii) the mass above the surface, M_{atm} , remains much smaller than M_a ; and (iii) for mountains with sinking, the mass in the fluid base, M_{base} , remains large compared to M_a , so that the magnetic line-tying condition at $r = r_{\text{min}}$ is a good approximation. Figures 4.2a, 4.2c, and 4.2e show $M_a(t) = \int_V dV \rho X_a$ integrated over the simulation volume V at time t . We see that the mountains achieve their target mass, which remains in the simulation for $t > T_a$. The injected mass $M_a(t)/M_a$, found by integrating equation 3.5 with respect to time, is plotted alongside in grey; the two curves overlap. Figure 4.2b shows M_{atm} for the hard-surface experiment; it is always small. Figures 4.2d and 4.2f show M_{base} for the soft-surface experiments; it always exceeds $\approx 10M_a$, as desired. For all simulations where $M_{\text{base}} > 0$, the total fraction of M_{base} lost through the outer boundary is $\ll 0.01\%$, consistent with Vigelius & Melatos (2008).

We next check that accretion takes place in the magnetostatic limit, i.e. that the total kinetic energy $E_k = \int_V dV \rho |\mathbf{v}|^2 / 2$ is small compared to the total magnetic energy $E_m = \int_V dV |\mathbf{B}|^2 / (8\pi)$. Figure 4.3 shows the ratio of E_k to E_m as a function of time. We see that E_k/E_m tends to increase with M_a but typically never rises above 1%, except in $\mathcal{H}(10^3)$ and the incomplete run $\mathcal{S}(r_{\text{min}}, 10^3)$. After accretion stops, at $t = T_a$, E_k/E_m typically falls to less than 10^{-4} .

Magnetic field transport in ZEUS-MP is divergence-free by construction (Hayes *et al.* 2006), but it is worth checking whether this property

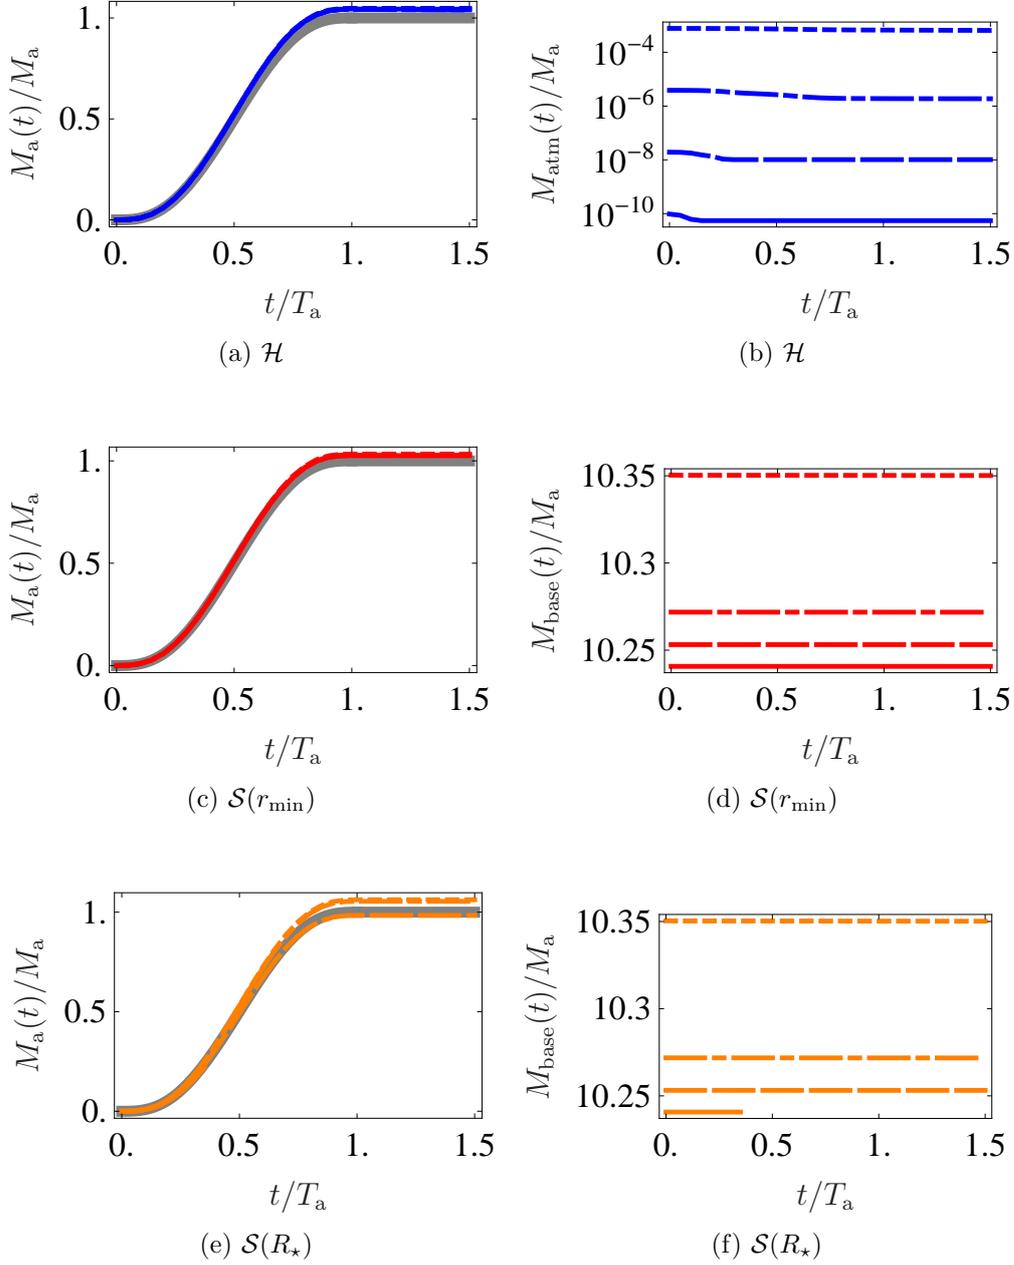


Figure 4.2: Accreted mass $M_a(t)$ and masses in the stellar atmosphere $M_{\text{atm}}(t)$ and fluid base $M_{\text{base}}(t)$, plotted in black, for $M_a/M_c = 1$ (dotted), 10 (dot-dashed), 10^2 (dashed), and 10^3 (solid). The injected mass from equation 3.5 is over-plotted in grey. The labels beneath each panel indicate a hard-surface (\mathcal{H} , blue) or soft-surface [$\mathcal{S}(r_{\min})$, red, or $\mathcal{S}(R_*)$, orange] run; see Table 4.1. The short solid orange line at the bottom of Figure (f) is from the aborted run $\mathcal{S}(R_*, 10^3)$

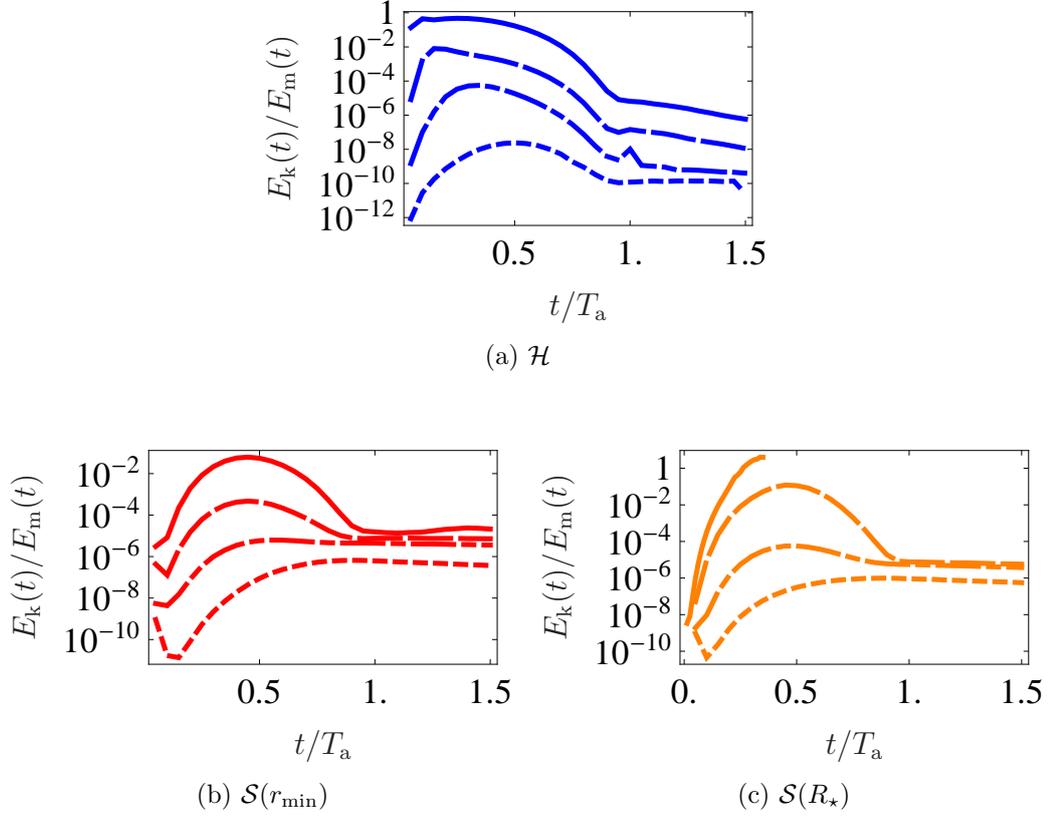


Figure 4.3: Kinetic energy $E_k(t)$, normalised by the magnetic energy $E_m(t)$, for $M_a/M_c = 1$ (dotted), 10 (dot-dashed), 10^2 (dashed), and 10^3 (solid). The labels beneath each panel indicate a hard-surface (\mathcal{H} , blue) or soft-surface [$\mathcal{S}(r_{\min})$, red, or $\mathcal{S}(R_*)$, orange] run; see Table 4.1.

is preserved by the injection algorithm. We find that the mean value of $|\nabla \cdot \mathbf{B}| / \sum_i (|B_i|/dx^i)$ is initially $\lesssim 6 \times 10^{-3}$, and increases by a factor of 3.5 at most over the run. Figure 4.4 shows contours of the normalised $|\nabla \cdot \mathbf{B}|$ for an illustrative mountain.

4.1.2 Illustrative example

We choose mountain $\mathcal{H}(10^2)$, grown on a hard surface with $M_a = 10^2 M_c$, to illustrate the general evolution of a magnetic mountain during accretion. Figure 4.5 shows contours of the mountain density ρX_a , normalised by the initial surface density ρ_{R_*} , at four different times. Figure 4.6 shows the magnetic flux ψ , normalised by $\psi_* = B_* R_*^2/2$, at the same times. Matter

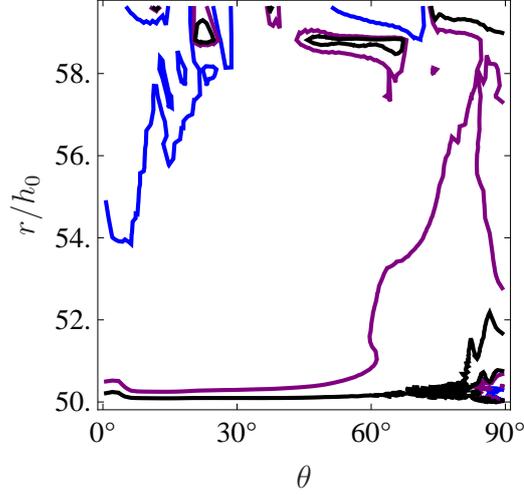


Figure 4.4: Contours of the absolute, normalised divergence of \mathbf{B} of mountain $\mathcal{H}(10^3)$ at $t = T_a$. Contour levels are at 5×10^{-2} (blue), 10^{-2} (purple), 5×10^{-3} (black).

is added predominately at the pole, as determined by $dM/d\psi$. In the early stages of accretion ($t = 0.1T_a$), the magnetic field is only slightly disturbed. As accretion progresses, the mountain spreads towards the equator, dragging the frozen-in magnetic field with it. The angular span of the ψ contours is compressed from $\sim 70^\circ$ [Figure 4.6a] to $\sim 20^\circ$ [Figure 4.6b]. At the half-way point ($t = 0.5T_a$), the flux is significantly displaced from its initial configuration, but remains anchored to the inner boundary at $r = r_{\min}$, demonstrating magnetic line tying. We see, in the $\psi/\psi_\star = 1.2 \times 10^{-2}$ contour, the early formation of the magnetic “tutu” configuration, observed in Payne & Melatos (2004, 2007) for $M_a = 10^{-5}M_\odot$.

At $t = T_a$, the mountain reaches its target mass ($M_a = 10^2M_c$ in Figures 4.5 and 4.6). Despite sliding towards the equator, the accreted matter still exhibits a noticeable variation in density with respect to θ ; a polar mountain is formed. The tutu configuration of the magnetic field is clearly visible; see for comparison Figures 2 and 4(a) of Payne & Melatos (2004). This equilibrium state remains largely unchanged when we run the simulation for an additional $0.5T_a$, during which no further mass is added.

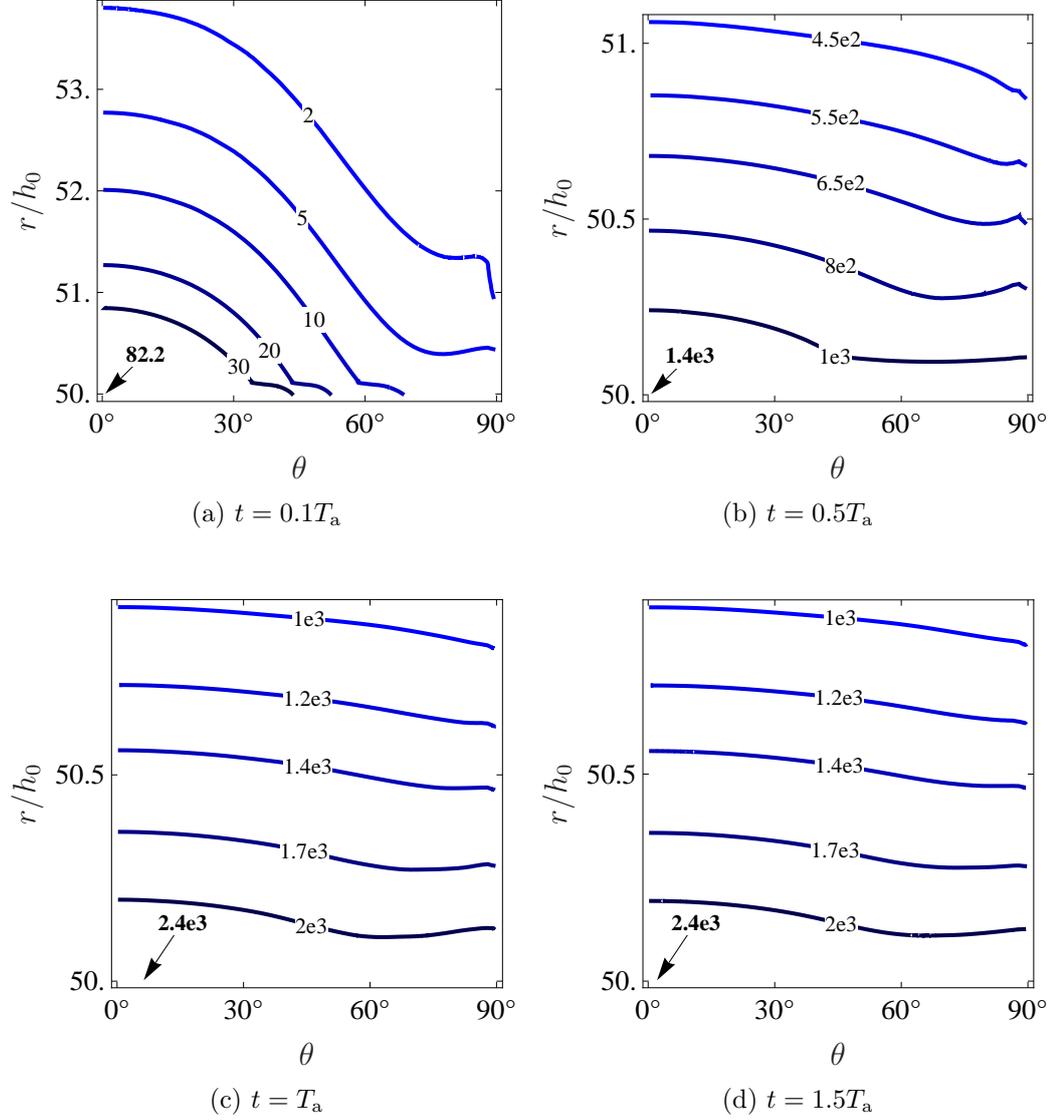


Figure 4.5: Hydromagnetic structure of hard-surface mountain $\mathcal{H}(10^2)$ at times $t/T_a = 0.1, 0.5, 1.0,$ and 1.5 : contours of accreted density $\rho X_a/\rho_{R_*}$ (blue). The maximum is indicated with a small arrow and labelled in bold. Colour intensities are linear in $\rho X_a/\rho_{R_*}$. Note that the scale of the r axis differs between the figures.

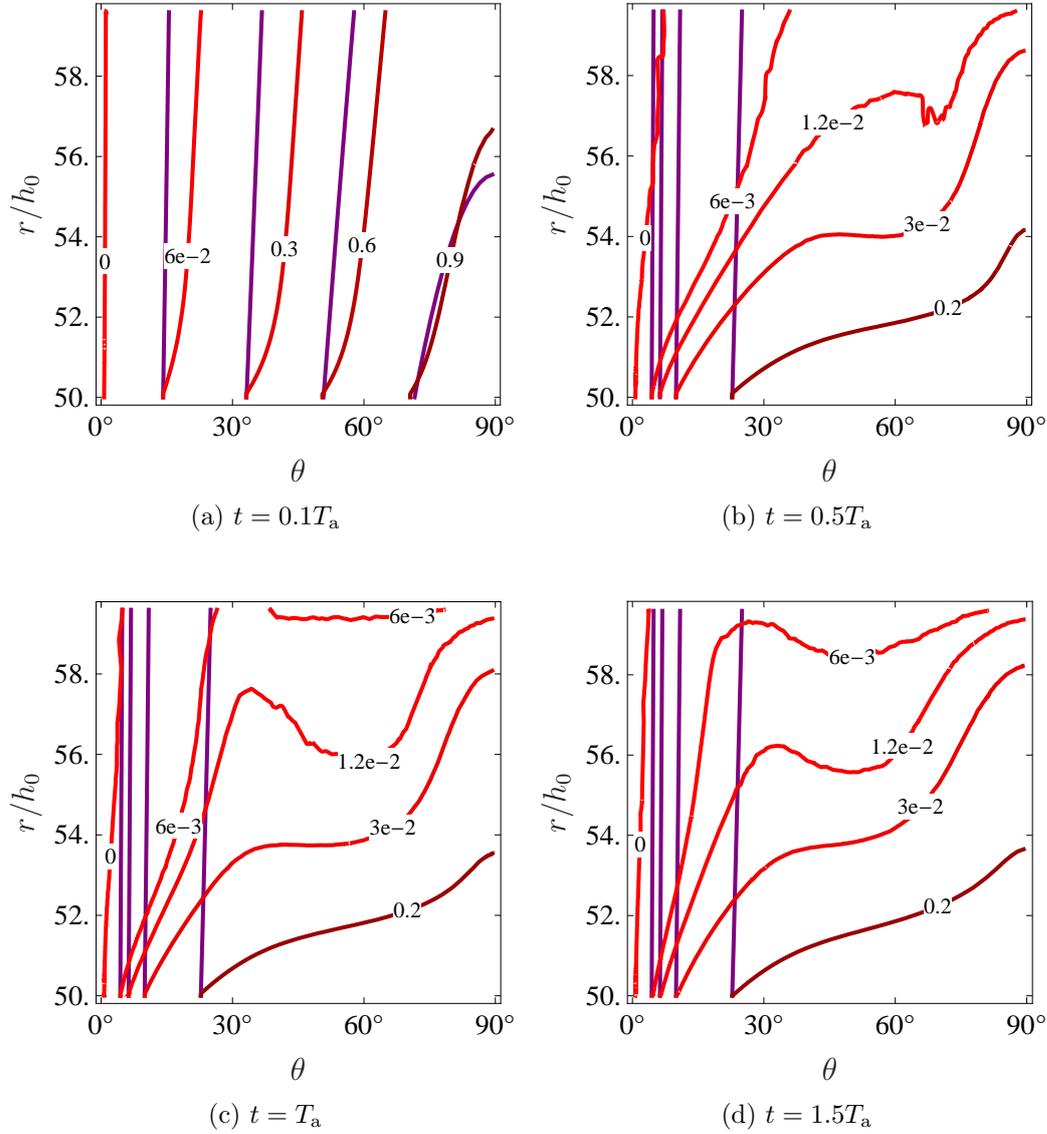


Figure 4.6: Hydromagnetic structure of hard-surface mountain $\mathcal{H}(10^2)$ at times $t/T_a = 0.1, 0.5, 1.0,$ and 1.5 : contours of magnetic flux ψ/ψ_* , at the labelled times (red), and at $t = 0$ (purple). The $t > 0$ contours meet their $t = 0$ equivalents on the left vertical axis. Colour intensities are linear in ψ/ψ_* . The change in the contour scale between Figure (a) and Figures (b)–(d) is due to the difference in ψ between $t = 0$, where the magnetic field is dipolar, and subsequent times when the magnetic field is distorted.

4.2 Comparison of mountains grown on hard and soft bases

We next compare the hard-surface equilibrium state, illustrated in Figures 4.5c and 4.6c, with the two experiments where we include sinking. Figures 4.7 and 4.8 show contours of ρX_a and ψ respectively in each of the three scenarios, with $M_a = 10M_c$. The hard-surface mountain [Figure 4.7a] spreads appreciably, and the magnetic flux [Figure 4.8a] is significantly displaced towards the equator. The density contour $\rho X_a/\rho_{R_\star} = 10$ begins a distance $\sim 4h_0 = 215$ cm (see Table 3.1) above the injection radius $r = r_{\min} = R_\star$ at the pole and sinks below the equator to $\sim 75\%$ of the polar height of the mountain. In contrast, the same contour of the sinking mountain grown at $r = r_{\min}$ [Figure 4.7b] begins $\sim 0.6h_0$ above the injection radius $r = r_{\min}$ at the pole and sinks below the equator to just $\sim 19\%$ of the polar height of the mountain. From the $\rho X_a/\rho_{R_\star} = 10^{-3}$ contour, we see that the accreted matter is confined to $r - r_{\min} \lesssim 0.7h_0$ above the inner boundary at the pole and $r - r_{\min} \lesssim 0.15h_0$ at the equator.

In short, the sunk mountain grown at $r = r_{\min}$ hugs the inner boundary and pole and resembles the initial mass distribution seen in Figure 4.5a. This is not surprising. Matter is fed in at $r = r_{\min}$ with zero velocity, as discussed in section 3.2.2. It expands outward due to the pressure gradient created as matter piles up at the injection radius; since we are injecting quasistatically, the pressure gradient is small. On the other hand, the weight of the massive overburden ($M_{\text{base}} \approx 10M_a$) presses down on the added material. The magnetic flux is displaced [Figure 4.8b], but much less than for the hard mountain. Field lines remain tied to the inner boundary, bending away in its immediate vicinity (because the slug of injected matter does not rise). Above this layer, the field lines of the initial and final states remain largely parallel.

4.2.1 Sinking scenarios

The structure of the sunk mountain grown at $r = R_\star$ [Figure 4.7c] differs from the other two cases. The contour $\rho X_a/\rho_{R_\star} = 10$, tracked above, starts at the pole, remains virtually flat at $\sim 2h_0$ below the injection radius, bends sharply inward near the equator, moves directly toward the inner boundary, then curves back towards the pole, crossing it again at $\sim 6h_0$ below the injection radius. For the previous two mountains, grown from $r = r_{\min}$, the angular variation in density increases with altitude. Here the reverse is true: the angular variation density decreases with increasing r , up until $r \lesssim R_\star$,

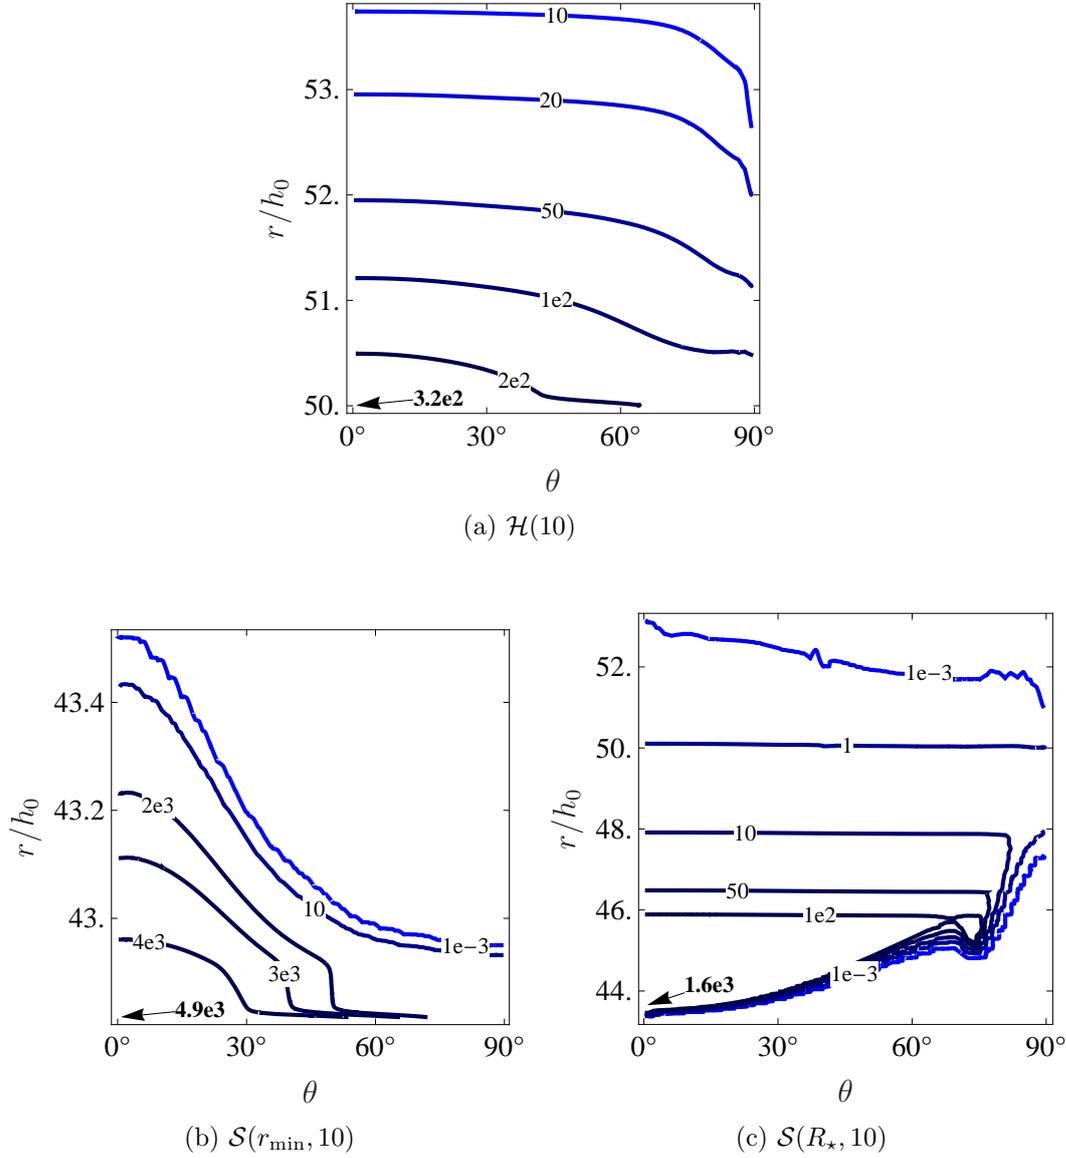


Figure 4.7: Comparison of the hydromagnetic structure of hard- and soft-surface mountains: contours of accreted density $\rho X_a / \rho_{R_*}$ of mountains $\mathcal{H}(10)$, $\mathcal{S}(r_{\min}, 10)$, and $\mathcal{S}(R_*, 10)$, at time $t = T_a$. Colour intensities are logarithmic to base 10 in $\rho X_a / \rho_{R_*}$; other details are as for Figure 4.5. Note that the scale of the r axis differs between the figures.

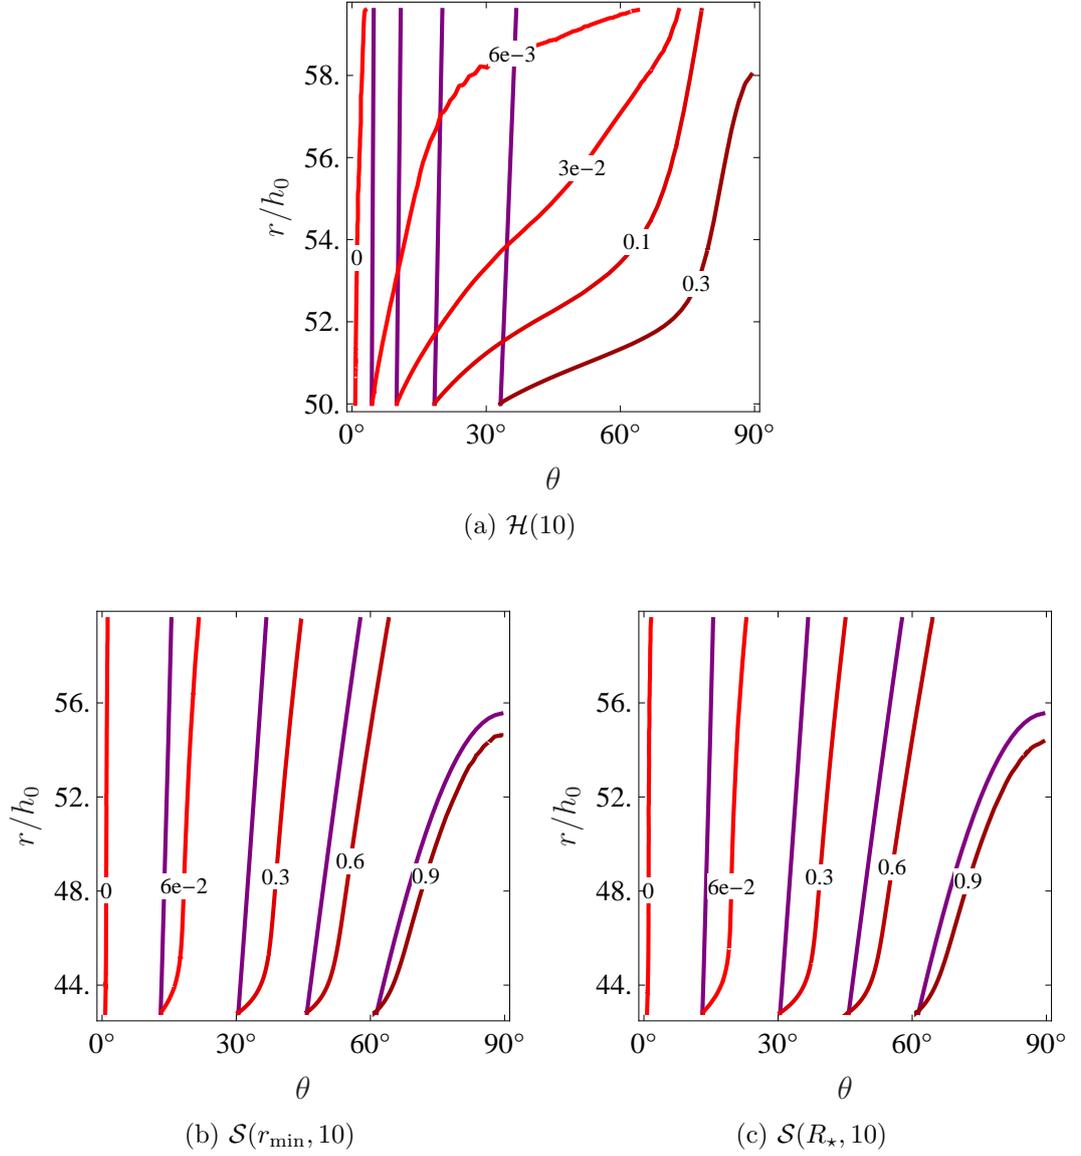


Figure 4.8: Comparison of the hydromagnetic structure of hard- and soft-surface mountains: contours of magnetic flux ψ/ψ_* of mountains $\mathcal{H}(10)$, $\mathcal{S}(r_{\min}, 10)$, and $\mathcal{S}(R_*, 10)$, at time $t = T_a$. Details are as for Figure 4.6. Note that the scale of the r axis differs between the figures.

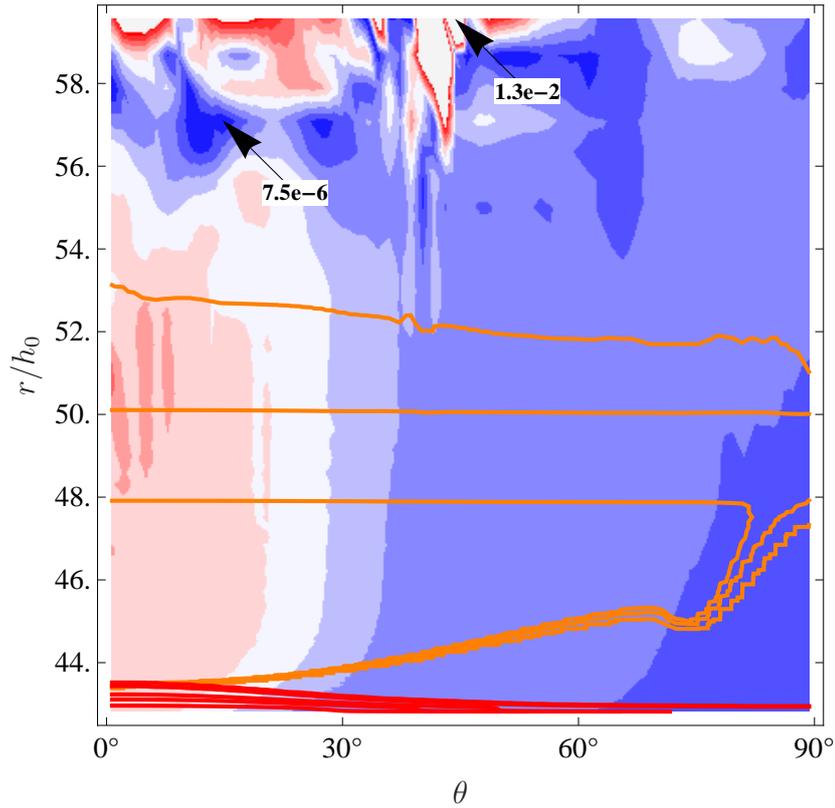


Figure 4.9: Shaded contours of the absolute, normalised difference in total density ρ between $\mathcal{S}(r_{\min}, 10)$ and $\mathcal{S}(R_x, 10)$, overlaid with the ρX_a (accreted only) contours of Figures 4.7b (red) and 4.7c (orange). Colouring ranges from blue (minimum difference) to white (halfway between the minimum and maximum differences) to red (maximum difference). The minimum and maximum density differences are indicated with small arrows and labelled in bold.

with the greatest variation within $\sim 4h_0$ of the inner boundary. While the two sinking scenarios differ in their final distributions of accreted (as opposed to total) density, their final distributions of magnetic flux [Figures 4.8b and 4.8c] are very similar.

The mountain sinks three times further into the fluid base at the pole than at the equator. This is consistent with how mass is injected according to equation 3.6; the input flux is ~ 20 times greater at the pole than at the equator. In addition, the magnetic field guides accreted matter sideways as field lines flatten across the surface towards the equator, whereas matter at the pole can sink inward readily along almost vertical flux tubes [e.g. the contour $\psi/\psi_\star = 6 \times 10^{-2}$ in Figure 4.8c].

The density contours bunch together along the underside of the mountain, spanning five orders of magnitude; the injected matter does not sink below this boundary. The lowest of the bunched density contours, $\rho X_a/\rho_{R_\star} = 10^{-3}$, never reaches the inner boundary; the mountain is floating in isostatic equilibrium with the surrounding fluid base. The $\rho X_a/\rho_{R_\star} = 10^{-3}$ contour rises to only $\sim 1h_0$ above R_\star at the equator and $\sim 3h_0$ at the pole; in contrast it sinks to $R_\star - r \lesssim 6h_0$ at the pole. The path of this contour in the region $r > R_\star$, if overlaid on Figure 4.7a, would trace densities between 2×10^{-3} and $50\rho_{R_\star}$; the density of accreted matter above R_\star is much less for the sunk mountain than for the hard-surface mountain. Finally, note that 0.3% of the mountain mass is above the stellar surface. Compared to the other two scenarios, the structure of $\mathcal{S}(R_\star, 10)$ is perhaps more reminiscent of an “iceberg”.

Ultimately we are interested in the final distribution of the total mass ρ , that is, the accreted matter, ρX_a , plus the fluid base it displaces, $\rho(1 - X_a)$. Does injection at r_{\min} or R_\star make a difference? Figure 4.9 displays the absolute, normalised difference $|\rho_{r_{\min}} - \rho_{R_\star}|/|\rho_{r_{\min}} + \rho_{R_\star}|$ between the total densities in the two sinking scenarios as a shaded plot. The largest differences occur at $r > R_\star$, where there is little mass, and are therefore unimportant; these differences are simply spurious numerical artifacts. For $r < R_\star$, the difference peaks near the pole but remains less than $\sim 0.8\%$. In other words, despite the difference in the final distribution of ρX_a between the two injection scenarios (emphasised by the overlaid contours), the final distribution of ρ is essentially the same. Injecting at r_{\min} or R_\star makes no difference, because the soft base readjusts in each case to yield the same overall equilibrium state. This is an important result. It confirms the robustness of the injection method and the argument presented in section 3.2.2: in ideal MHD, the equilibrium state is independent of precisely where matter is initially injected. In practice, injection at r_{\min} seems preferable, because it reduces the simulation time (see Table 4.1) and improves numerical stability, as illustrated by the failure of

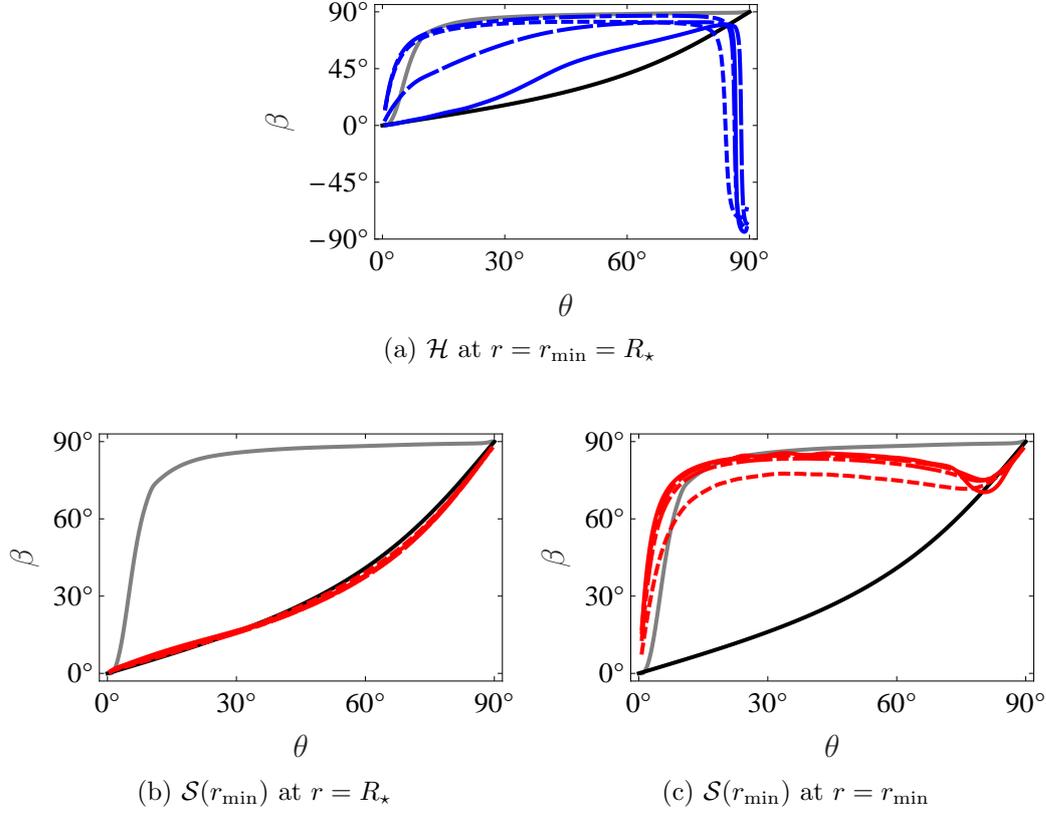


Figure 4.10: Angle between the magnetic field \mathbf{B} and the radial unit vector, plotted versus colatitude θ at (a) $r = r_{\min} = R_{\star}$, (b) $r = R_{\star}$, and (c) $r = r_{\min}$, for $M_a/M_c = 1$ (dotted), 10 (dot-dashed), 10^2 (dashed), and 10^3 (solid). Results from this work are plotted in colour. Plotted in monochrome are $\beta(r, \theta)$ for a Payne & Melatos (2004) Grad-Shafranov mountain (grey), and for a dipole (black).

run $\mathcal{S}(R_{\star}, 10^3)$.

4.2.2 Magnetic line tying

Finally, we investigate the assumption of magnetic line tying. In Figure 4.10 we plot the angle $\beta(r, \theta) = \sin^{-1}(B_{\theta}/|\mathbf{B}|)$ between the magnetic field \mathbf{B} and the radial unit vector as a function of θ at $r = r_{\min}$ and $r = R_{\star}$. We also plot β at the inner boundary of a Grad-Shafranov mountain with $M_a = M_c$ (Payne & Melatos 2004), and β for a dipole (independent of radius). The $\mathcal{S}(R_{\star})$ mountains give the same results as $\mathcal{S}(r_{\min})$.

The hard-surface mountains in Figure 4.10a behave like the Grad-Shafranov

mountain at low M_a but become increasingly dipolar as M_a increases. This is expected; at low M_a , the accreted mass stays close to the pole and distorts the magnetic field there. As M_a increases, the mountain spreads over a greater volume, and the magnetic field is distorted less at any particular point. The sign inversion close to the equator may be caused by numerical reconnection, or by the reflective boundary condition at $\theta = \pi/2$; further tests with a resistive ideal-MHD solver are needed to make sure.

The soft-surface mountains in Figure 4.10b are dipolar at $r = R_*$, as expected, but at $r = r_{\min}$ [Figure 4.10c] they more closely resemble the Grad-Shafranov mountain. One might expect β to closely match a dipole at $r = r_{\min}$, given that the magnetic field lines are tied there, and we choose $M_{\text{base}} \gg M_a$ in order to minimise sideways fluid displacements at the bottom of the soft base. It is unclear whether the magnetic distortions are artificial, because the injected slug of matter enters from below and cannot expand upwards to match accretion from above (see section 4.3 for further discussion). Alternatively, kinks in the magnetic field may be communicated rapidly down to arbitrary depths by Alfvén waves, even though the Alfvén speed $\propto \rho^{-1/2}$ decreases rapidly with depth. If so, the high breaking strain of the solid, conducting crust (Horowitz & Kadau 2009) assumes even greater importance in enforcing line tying.

We argued, in section 3.2.2, that the final equilibrium state (i.e. of the mountain plus the fluid base) is independent of r_{inj} ; this has been demonstrated in Figure 4.9. In general, however, a given total M_a and injected mass flux $\partial M_a / \partial \psi$ does not define a *unique* ideal-MHD equilibrium. Matter injected from above spreads sideways faster than it sinks, like a layered cocktail drink, while a slug of matter injected from below forces the base sideways without much movement at the surface [compare Figures 4.7c and 4.7b]. Conceivably, therefore, ZEUS-MP may converge on different equilibria depending on r_{inj} . The results of section 4.2 engender confidence that the final equilibrium state does not depend on r_{inj} ; the issue is not definitively settled, however, for the following subtle reason.¹

Consider a polar field line in Figure 4.6. As accretion proceeds, it bends towards the equator until it touches the corner $(r, \theta) = (r_{\max}, \pi/2)$. At that point, it instantaneously snaps through some nonzero angle, from $B_r \neq 0$ (free boundary at $r = r_{\max}$) to $B_r = 0$ (reflecting boundary at $\theta = \pi/2$). Effectively, this corresponds to a dissipative, reconnection-like event occurring just *outside* the simulation volume, artificially pinching off magnetic loops.²

¹Sterl Phinney, private communication.

²The effect can be magnified in ZEUS-MP by increasing the cell size close to $r = r_{\max}$; eventually ZEUS-MP aborts when B_r diverges close to the $(r_{\max}, \pi/2)$ corner.

Such a process is irreversible. Furthermore, it acts differently on the sequence of quasistatic equilibria that ZEUS-MP hypothetically passes through during slow accretion from above and below, because sideways spreading happens at different altitudes in the two cases.

In the runs presented in this work, the density in the vicinity of the corner $(r_{\max}, \pi/2)$ is tiny, as is the mass efflux through the boundary $r = r_{\max}$ (see section 4.1.1). One can therefore argue plausibly that the irreversible dissipation at $(r_{\max}, \pi/2)$, while it exists in principle, does not significantly affect the final state. There is a chance, however, that if one adds material slowly from above, reconnection (where numerical or real) pinches off one small magnetic loop after another at the equator, as in the Earth's magnetotail. Resistive MHD simulations by Vigelius & Melatos (2009b) do not show such behaviour, but they mostly started from preformed Grad-Shafranov equilibria instead of growing the mountain from scratch. A more careful consideration of this issue is required for future simulations.

4.2.3 Mass quadrupole moment

The distorted hydromagnetic equilibria in section 4.2 have an associated mass quadrupole moment, with principal axis along the pre-accretion magnetic axis, which is quantified in terms of the poloidal ellipticity³

$$\epsilon = \frac{\pi}{I_{zz}} \int_{r_{\min}}^{r_{\max}} dr r^4 \int_0^{\pi/2} d\theta \sin \theta (3 \cos^2 \theta - 1) \rho(t, r, \theta), \quad (4.1)$$

with $I_{zz} = 2M_{\star}R_{\star}^2/5$. Figure 4.11 shows ϵ with respect to time as the mountain grows. All mountains achieve a nonzero ellipticity at $t = T_a$, which decreases negligibly thereafter, confirming the mountains are stable; extending the simulation time to confirm the mountains' stability was not attempted due to computational constraints. The time taken for the hard-surface mountains to converge to their equilibrium values of ϵ decreases with M_a , from $\sim T_a$ ($M_a = 1M_c$) to $\sim 0.1T_a$ ($M_a = 10^3M_c$); this is consistent with the decreased confinement of the mountain by the magnetic field, i.e. its increased ability to spread. The ellipticities of the soft-surface mountains for the two injection scenarios are virtually identical; even the incomplete run $\mathcal{S}(R_{\star}, 10^3)$ closely follows $\mathcal{S}(r_{\min}, 10^3)$ up until failure. As in section 4.2, the final density distribution is independent of the injection procedure.

³ We note that the ellipticity calculated here is the poloidal ellipticity $\epsilon = (I_3 - I_1)/I_3$ appropriate for an axisymmetric object with principal moments of inertia $I_1 = I_2 < I_3$ (Vigelius & Melatos 2008). It is not the same as the equatorial ellipticity $\epsilon = (I_2 - I_1)/I_3$ appropriate for a nonaxisymmetric object with $I_1 < I_2 < I_3$, which is used elsewhere in this thesis, in particular in the context of gravitational waves.

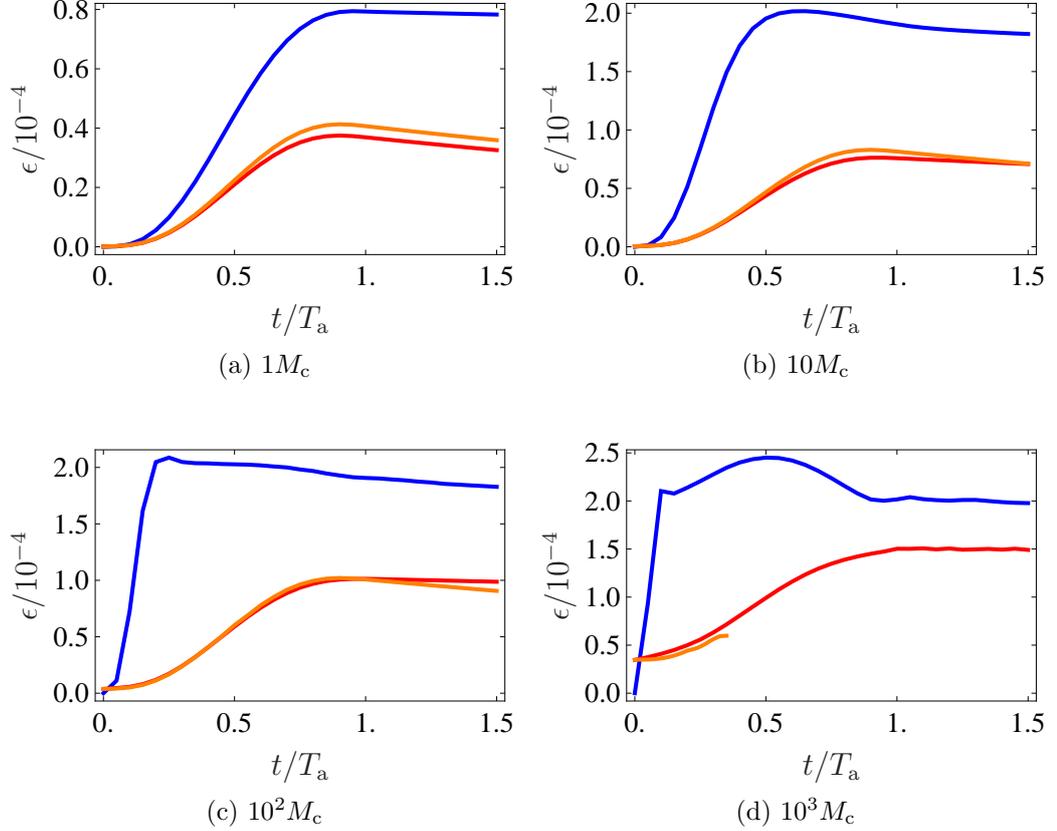


Figure 4.11: Ellipticity as a function of time, for $M_a/M_c = 1, 10, 10^2$, and 10^3 (top left to bottom right), and for mountains \mathcal{H} (blue), $\mathcal{S}(r_{\min})$ (red), and $\mathcal{S}(R_*)$ (orange).

The origin of the uneven behaviour of the ellipticity of $\mathcal{H}(10^3)$ [Figure 4.11d] is unknown; we note, however, that its functional dependence on t is similar to that of \dot{M}_a [equation 3.5 and Figure 3.2], and therefore it is likely that the rise and fall of the ellipticity is due to the reconfiguration of the $\mathcal{H}(10^3)$ mountain in response to the changing accretion rate. We note similar undulations in the kinetic energies [Figures 4.3b and 4.3c] and, to a lesser degree, in the ellipticity of $\mathcal{H}(10)$ [Figure 4.11b].

Figure 4.12 shows ϵ (in black) at $t = T_a$ as a function of M_a . The ellipticities of the hard-surface mountains rise by a factor of ~ 2.5 as M_a increases from $1M_c$ to $10M_c$ and flattens thereafter, rising by a further 5% as M_a increases from $10M_c$ to 10^3M_c . Importantly, accreting further matter does not reduce ϵ ; the mountain does not smooth itself out. The ellipticities of the

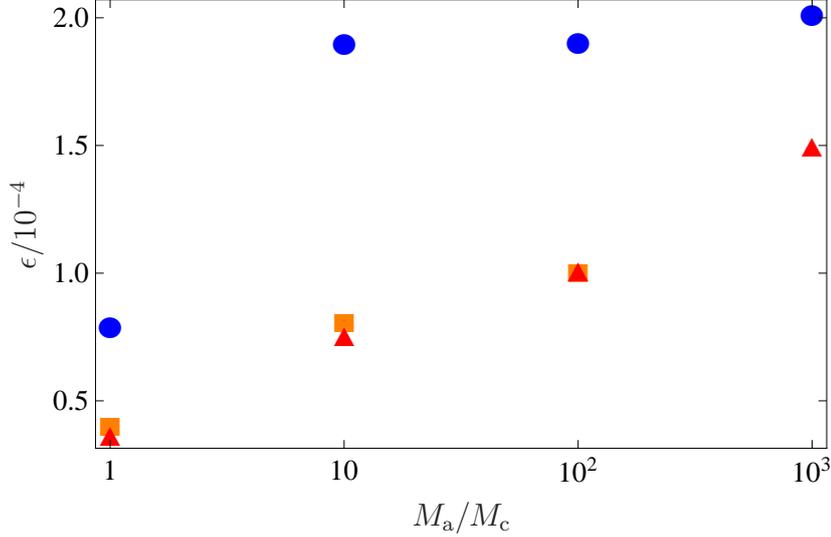


Figure 4.12: Ellipticity as a function of M_a , at time $t = T_a$, for mountains \mathcal{H} (circles), $\mathcal{S}(r_{\min})$ (triangles), and $\mathcal{S}(R_*)$ (squares).

soft-surface mountains rise by $\sim 60\%$ per decade in M_a .

Figure 4.12 clearly quantifies the effect of sinking: ϵ decreases, relative to the hard-surface scenario, by $\sim 50\%$ at $M_a = 1M_c$, $\sim 60\%$ at $M_a = 10M_c$, and $\sim 25\%$ at $M_a = 10^3M_c$.

4.3 Comparison with Choudhuri & Konar

Choudhuri & Konar (2002) developed a kinematic model of accretion, which treats sinking in a different (but complementary) way to this work. An axisymmetric magnetic field is evolved under the influence of a *prescribed* velocity field, which models the flow of accreted matter from pole to equator, where it submerges and moves towards the core (see their Figure 1). Ohmic diffusion is included, but, for a subset of the results (where the resistivity $\eta = 0.01$), it is negligible, permitting a direct comparison with this work.

Figure 5 of Choudhuri & Konar (2002) shows the evolved configuration of an initially dipolar field which permeates the entire star. We compare to Figures 4.8b and 4.8c of this work. In both models, the magnetic field is distorted significantly by accreted matter spreading towards the equator ($r_m < r < r_s$ in Choudhuri & Konar; the entire simulation in this work). In Choudhuri & Konar’s work, the magnetic field is completely submerged beneath the surface and confined to the zone where the submerged accreted

matter flows back towards the pole. In this work, magnetic field lines still penetrate the surface, implying less effective screening. Within the core, Choudhuri & Konar’s magnetic field remains relatively undisturbed. Magnetic line-tying is not enforced, but the prescribed radial flow within the core naturally restricts the sideways displacement of the magnetic field there. If there were sideways motion of the matter within the core, it would modify the degree of magnetic screening, but neither our simulations nor the results of Choudhuri & Konar show evidence for such motion. Extending our simulations deeper into the star to include the core and explore this possibility properly would be a technical challenge; for instance, we would need to incorporate a more realistic equation of state and track even more disparate equilibrium time-scales.

4.4 Implications for gravitational waves

To explain the narrow range in the rotation frequencies of low-mass x-ray binaries (Chakrabarty *et al.* 2003), it is proposed that the stars radiate angular momentum in gravitational waves at a rate which balances the accretion torque (Wagoner 1984, Bildsten 1998). Magnetic mountains are one of a number of physical mechanisms proposed for the associated permanent quadrupole; see Vigeliu & Melatos (2009a) and references therein. The relationship between ϵ and the rotation frequency f predicted by torque balance is $f \propto \epsilon^{-2/5}$. Thus, the 25% to 60% reduction in ϵ due to sinking calculated in this work increases f by 12% to 44%, all other things being equal. This goes some way towards bringing magnetic mountain ellipticities down to a level consistent with the data, but there is still a long way to go. Observations to date have found $45 \text{ Hz} < f < 620 \text{ Hz}$ for burst oscillation sources and $182 \text{ Hz} < f < 598 \text{ Hz}$ for accreting millisecond pulsars, implying $6.6 \times 10^{-9} \lesssim \epsilon \lesssim 4.6 \times 10^{-6}$ and $7.2 \times 10^{-9} \lesssim \epsilon \lesssim 1.4 \times 10^{-7}$ respectively. Conversely, the ellipticities of sunk mountains calculated in this work, $3.5 \times 10^{-5} \lesssim \epsilon \lesssim 1.5 \times 10^{-4}$, imply $11 \text{ Hz} \lesssim f \lesssim 20 \text{ Hz}$. Clearly, other relaxation mechanisms, like Ohmic diffusion, must also be playing an important role in reducing ϵ , as the observed f require.

The reduction in ϵ by sinking also reduces the gravitational wave strain (e.g. Abbott *et al.* 2007b), $h \propto \epsilon f^2$, by 6% to 17%. This is unlikely, by itself, to rule out the detection of gravitational waves from low-mass x-ray binaries by ground-based interferometric detectors; assuming the signal can be coherently integrated, the loss in h can be compensated for by an increase in the observation time $\propto h^{-2}$ of 13% to 45%. Other difficulties associated with the detection of gravitational waves from low-mass x-ray binaries, such

as poorly known orbital parameters and accretion-induced phase wandering (Watts *et al.* 2008), are likely to be more important.

4.5 Summary

In this chapter, we simulated the growth of magnetically confined mountains on an accreting neutron star, using the numerical procedure presented in Chapter 3, under the two scenarios where the mountain sits on a hard surface and sinks into a soft, fluid base. In the latter scenario, we confirmed that the final equilibrium state is independent of the altitude where matter is injected. We found that the ellipticity of a hard-surface mountain does not increase appreciably for $M_a \gtrsim 10M_c$, saturating at $\sim 2 \times 10^{-4}$, whereas the ellipticity of a soft-surface mountain continues to increase from $M_a = 10M_c$ to 10^3M_c . Sinking reduces the ellipticity by up to 60% relative to the hard-surface value.

Chapter 5

Gravitational wave data analysis

The search for gravitational wave signals in the output of ground-based interferometric detectors, such as LIGO and VIRGO, presents many significant challenges. For example, the detectors are much more sensitive to innumerable sources of noise than they are to the expected sources of gravitational waves; noise sources must therefore be identified, characterised, and, if possible, eliminated. Noise may mimic the form of a plausible signal, rendering the two indistinguishable; the detector output must then be checked against data quality vetoes, or else the output of another widely-separated gravitational wave detector. Any first detection of gravitational waves is likely to be marginal, if not controversial; to have any confidence in the statistical significance of a candidate signal, a thorough understanding of the statistics of the analysis method is essential.

In this chapter, we summarise previous and contemporary searches for gravitational waves using interferometric detectors (in section 5.1), and review data analysis techniques pertaining to periodic gravitational wave signals (in section 5.2). The latter section is intended to provide the background for Chapters 6 and 9.

5.1 Summary of gravitational waves searches

Anticipated sources of gravitational waves are generally divided into four categories, based on the properties of their signals. Each category of signals requires a different data analysis approach. Short-lived signals, which last on the order of seconds, consist of compact binary coalescences (section 5.1.1), and bursts (section 5.1.2). Long-lived signals, which persist over a typical

observing time, are divided into the stochastic background (section 5.1.3), and periodic gravitational waves (section 5.2).

5.1.1 Compact binary coalescences

A binary system comprising two compact objects will radiate away orbital angular momentum as gravitational radiation; the two compact objects will spiral towards each other, and eventually merge. The gravitational signal emitted during these compact binary coalescences (CBCs) is generally divided into three phases: the inspiral phase, before the merger; the merger itself; and the ringdown phase, during which the newly-formed compact object reaches equilibrium. The inspiral and ringdown waveforms are modelled to high accuracy using a variety of analytic methods, including post-Newtonian expansion and black hole perturbation theory (e.g. Blanchet 2006). Improved understanding of the merger phase, including the computation of its waveform, has come through recent advances in numerical relativity (e.g. Baker *et al.* 2007, Cadonati *et al.* 2009).

A typical analysis (e.g. Abbott *et al.* 2009g,h) uses matched filtering of the detector output against a bank of possible template waveforms, producing a set of triggers. A subset of the data, termed the playground, is set aside to tune the analysis pipeline. The triggers are searched for coincidence between detectors, and subjected to a number of consistency tests designed to eliminate noise-generated signals. To determine the background rate of triggers originating from noise, data from two detectors are time-shifted, with respect to each other, by an interval greater than the light travel time between the detectors; any remaining coincident triggers must therefore be due to noise. Coincident triggers in the un-time-shifted data are assessed against the background rate of triggers to determine whether they are detection candidates. In the absence of a detection, 90% confidence upper limits are set on the rate of CBCs per year per unit L_{10} ; the latter unit equals 10^{10} times the blue luminosity of the Sun, and is a measure of the number of galaxies within the sensitive reach of the search. Upper limits are expressed with respect to three canonical types of binary system: binary neutron stars, binary black holes, and black hole-neutron star binaries.

The most recent LIGO search for CBCs (Abbott *et al.* 2009i) used S5 data to set upper limits on the rate of neutron star, black hole, and black hole-neutron star binary coalescences, where the total mass of the binary is $\lesssim 35M_{\odot}$, of 1.4×10^{-2} , 7.3×10^{-4} , and $3.6 \times 10^{-3} \text{ yr}^{-1} L_{10}^{-1}$ respectively. These rates are only a few orders of magnitude above optimistic astrophysical estimates. Using S4 data, Abbott *et al.* (2008g) set upper limits of ~ 0.3 – $1.0 \text{ yr}^{-1} L_{10}^{-1}$ on binary black holes with total binary masses of ~ 30 – $80M_{\odot}$,

and Abbott *et al.* (2009g) set an upper limit of $1.6 \times 10^{-3} \text{ yr}^{-1} L_{10}^{-1}$ on binary black holes with total binary masses of $85\text{--}390M_{\odot}$.

5.1.2 Bursts

Burst searches are designed to be sensitive to a wide class of un-modelled transient gravitational wave signals, including from binary black hole mergers and core collapse supernovae. They face many of the same issues as CBC searches, in particular that the detector noise is a rich source of plausible gravitational wave signals. The relaxed assumption of an arbitrary waveform presents an additional challenge to designing efficient and robust search algorithms (e.g. Searle *et al.* 2008). Often, an algorithm will decompose the detector output into a time-frequency plane, then search for clusterings of points where a particular statistic exceeds the background. Other methods use three or more widely-separated detectors to reconstruct the signal polarisations and its sky position. Two recent LIGO searches during S5 set upper limits, at 90% confidence, on the burst rate per year, in the limit of strong signals: a rate of 3.6 bursts per year for source frequencies below 2 kHz (Abbott *et al.* 2009f), and 5.4 bursts per year for source frequencies in the range 1–6 kHz (Abbott *et al.* 2009j).

Bursts of gravitational waves which are detectable by LIGO are very likely to be accompanied by electromagnetic counterparts. Multi-messenger astronomy aims to cross-correlate gravitational wave triggers with real-time electromagnetic observations (Kanner *et al.* 2008), and similarly to follow up triggers from these sources with gravitational wave searches (Abbott *et al.* 2008b). Several LIGO burst searches have already been carried out using γ -ray triggers. In Abbott *et al.* (2008f), gravitational waves associated with γ -ray bursts (GRBs) have been searched for in the S2–S4 data; best 90% upper limits on gravitational wave strain were on the order of $10^{-21} \text{ Hz}^{-1/2}$. Using S5 data, Abbott *et al.* (2008d) ruled out, with 99% confidence, the possibility that GRB 070201 originated from a CBC located in the Andromeda galaxy (as was suggested by its electromagnetically reconstructed position) with component masses of $1\text{--}3M_{\odot}$ and $1\text{--}40M_{\odot}$. A number of collaborations are planned between the LIGO Scientific Collaboration and a number of electromagnetic and neutrino observatories (Shawhan 2009).

Abbott *et al.* (2008e) searched for gravitational waves associated with bursts from soft gamma repeaters (SGRs), and placed 90% upper limits on their emitted gravitational wave energy (assuming isotropic emission and a source distance of 10 kpc) of 3×10^{45} to 9×10^{52} erg, within the predictions of some theoretical SGR models. A specific analysis targeting the SGR 1900+14 storm of March 29, 2006, using a more sensitive technique, improved the

upper limits on gravitational wave energy emission during the storm to 2×10^{45} to 6×10^{50} erg, comparable to the energies in SGR giant flares (Abbott *et al.* 2009k).

5.1.3 The stochastic background

The stochastic gravitational wave background is the broadband, cumulative signal from a number of long-lived astrophysical and cosmological sources. Astrophysical sources include the superposition of periodic signals from spinning neutron stars, low-mass x-ray binaries, and supernovae. Cosmological sources include primordial spacetime fluctuations amplified by inflation in the early Universe, and cosmic strings. The gravitational wave strain of the stochastic background is far below the strain noise of even advanced ground-based detectors. To improve the signal-to-noise ratio, detector outputs are cross-correlated and integrated over long observation times. Additional environmental monitoring data is used to identify any instrumental lines which may be cross-correlated between detectors.

Upper limits on the stochastic gravitational wave background are given in terms of the dimensionless $\Omega_{\text{GW}}(f)$, the gravitational wave energy density per logarithmic frequency per the critical energy density of the Universe (Abbott *et al.* 2004a). Assuming the Hubble constant is $72 \text{ km s}^{-1} \text{ Mpc}^{-1}$, LIGO S4 data has been used to set 90% upper limits of $\Omega_{\text{GW}}(f) < 6.5 \times 10^{-5}$ over 51–150 Hz (Abbott *et al.* 2007c), and LIGO S5 data combined with Virgo data has set 95% upper limits of $\Omega_{\text{GW}}(f) < 6.9 \times 10^{-6}$ around 100 Hz (Abbott *et al.* 2009c). This latest result improves upon limits deduced from the cosmic microwave background and Big Bang nucleosynthesis. Abbott *et al.* (2007d) searched for an anisotropic background in S4 data, and Abbott *et al.* (2007a) searched for the background around 900 Hz by correlating the LIGO Livingston interferometer with the ALLEGRO resonant mass detector.

5.1.4 Periodic gravitational waves

Long-lived, quasi-monochromatic gravitational wave signals are known as periodic (also continuous or persistent) gravitational waves. They are anticipated to be generated by rapidly spinning neutron stars (see section 2.4). Searches for periodic gravitational waves have, to date, divided into three categories: searches targeting known non-accreting pulsars, searches directed at the low-mass x-ray binary Scorpius X-1, and all-sky searches for as-yet undiscovered isolated neutron stars.

The first search for gravitational waves from known pulsars targeted the pulsar J1939+2134, using data from LIGO's S1 science run (Abbott *et al.*

2004b). Subsequent searches targeted 28 pulsars using data from S2 (Abbott *et al.* 2005b), 78 pulsars using S3 and S4 data (Abbott *et al.* 2007e), and 116 pulsars using data from S5 (Abbott *et al.* 2010). The gravitational wave signal is assumed to be phase-locked to the pulsars' electromagnetic signal, which are known from radio and (in one case) x-ray observations. The observed frequency and spindown of the radio signal also sets an indirect limit, based on energy conservation, on the gravitational wave strain (see section 8.2). The first pulsar for which this indirect limit was reached by a gravitational wave search was the Crab pulsar (Abbott *et al.* 2008c). In the most recent analysis (Abbott *et al.* 2010), the gravitational wave strain of the Crab is limited to be, at best, on the order of 10% of the indirect upper limit, and the energy radiated in gravitational waves by the Crab is limited to less than 2% of its total rotational energy. The measured upper limits of the pulsars PSR J0537-6910 and PSR J1952+3252 are also very close to their indirect limits.

Indirect upper limits on gravitational waves from accreting neutron stars are set by balancing the energy lost due to gravitational waves with the energy gained from accretion. The latter quantity is determined from the star's observed x-ray luminosity. Magnetic mountains (Chapters 3–4) are one possible mechanism by which accreting neutron stars may emit gravitational radiation. The highest such indirect limit, that of Scorpius X-1, is reachable only by second-generation detectors. Two LIGO searches have targeted Scorpius X-1: a search of S2 data using the \mathcal{F} statistic (Abbott *et al.* 2007b), and a search of S4 data using a cross-correlation technique (Abbott *et al.* 2007d).

All-sky searches for unknown isolated neutron stars are computationally limited by the very large number of templates needed to cover the sky, and a broad range of frequencies and frequency derivatives. LIGO searches to date have taken a number of different approaches. The S2 data was searched using the single-detector \mathcal{F} statistic (Abbott *et al.* 2007b), and the Hough transform method (Abbott *et al.* 2005a). Abbott *et al.* (2009a) conducted a multi-detector \mathcal{F} statistic search of S4 data using Einstein@Home, a very large-scale distributed computing project. The S4 data was also searched using three different semi-coherent methods: PowerFlux, StackSlide, and the Hough transform (Abbott *et al.* 2008a). An all-sky analysis of the complete S5 data set is in progress; thus far, searches of early S5 data have been completed using PowerFlux (Abbott *et al.* 2009b) and Einstein@Home (Abbott *et al.* 2009d).

5.2 Periodic gravitational wave data analysis

This section reviews the data analysis of periodic gravitational waves. We present the model of the gravitational wave signal in section 5.2.1, and a fully coherent analysis method, the \mathcal{F} statistic, in section 5.2.2. In section 5.2.3 we review issues regarding the computational cost of periodic gravitational wave searches.

5.2.1 The signal model

The signal of a periodic gravitational wave is modelled analytically to high accuracy. A full description of the signal model is given in Jaranowski *et al.* (1998), and further analysed in Jaranowski & Królak (1999, 2000). This section presents an outline of its derivation.

We first define a number of reference frames. The gravitational wave frame, which is identical to the inertial frame defined by the transverse traceless gauge (section 2.3), is denoted $(\hat{x}_w, \hat{y}_w, \hat{z}_w)$. The \hat{z}_w axis points in the direction of travel, and the orientation of the \hat{x}_w - \hat{y}_w plane about the \hat{z}_w axis is determined by the polarisation angle ψ . In this frame, the metric perturbation \mathbf{H}_w is given by

$$\mathbf{H}_w(t) = h_+(t)\mathbf{H}_w^+ + h_\times(t)\mathbf{H}_w^\times, \quad (5.1)$$

where the polarisation bases are given by

$$\mathbf{H}_w^+ = \hat{x}_w \otimes \hat{x}_w - \hat{y}_w \otimes \hat{y}_w = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (5.2a)$$

$$\mathbf{H}_w^\times = \hat{x}_w \otimes \hat{y}_w + \hat{y}_w \otimes \hat{x}_w = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (5.2b)$$

and \otimes denotes the tensor product. The components of each polarisation are modelled by

$$h_+(\tau) = A_+ \cos \Phi(\tau), \quad h_\times(\tau) = A_\times \sin \Phi(\tau), \quad (5.3)$$

where A_+ and A_\times are their respective amplitudes, and $\Phi(\tau)$ is the phase of the gravitational wave with respect to the neutron star proper time τ .

In the reference frame of an interferometric detector, denoted $(\hat{x}_d, \hat{y}_d, \hat{z}_d)$, the vectors

$$\hat{x}_d = \hat{n}_d^1 = (1, 0, 0), \quad \hat{n}_d^2 = (\cos \zeta, \sin \zeta, 0) \quad (5.4)$$

point along the interferometer arms, which are separated by an angle ζ . The response of the detector to the metric perturbation (Finn & Chernoff 1993), as measured with respect to its local time t , is equal to

$$h(t) = \text{tr} [\mathbf{R}_d^T \mathbf{H}_d(t)] , \quad (5.5)$$

where the response matrix is given by (Finn 2009)

$$\mathbf{R}_d = \frac{1}{2} (\hat{n}_d^1 \otimes \hat{n}_d^1 - \hat{n}_d^2 \otimes \hat{n}_d^2) , \quad (5.6)$$

and \mathbf{H}_d denotes the metric perturbation in the detector frame.

In propagating from the neutron star to the detector, the gravitational wave signal is modulated both in amplitude and in phase. The phase modulation results from converting the neutron star proper time to the detector time; it is a function of the sidereal and orbital motion of the Earth, the linear motion of the star (and its orbital motion, if it is in a binary), and a number of relativistic corrections. The amplitude modulation arises from transforming from the wave frame to the detector frame; it is only dependent on the Earth's sidereal rotation.

The gravitational wave phase, as measured in the neutron star rest frame, is modelled as a Taylor expansion:

$$\Phi(\tau) = \Phi_{\tau_0} + 2\pi \sum_{n=0}^N f_w^{(n)} \frac{(\tau - \tau_0)^{n+1}}{(n+1)!} , \quad (5.7)$$

where $f_w^{(n)}$ are the time derivatives of the frequency evaluated in the neutron star rest frame at an initial time τ_0 . The relationship between the neutron star proper time and the detector time is given by

$$\tau - \tau_0 = \sqrt{1 - \frac{v_{\text{ns}}^2}{c^2}} (t_{\text{ns}} - t_{\text{ns},0}) , \quad (5.8)$$

$$t_{\text{ssb}} = t_{\text{ns}} + \frac{|\vec{r}_{\text{ns}}(t_{\text{ns}})|}{c} + \Delta , \quad (5.9)$$

$$t = t_{\text{ns}} + \frac{|\vec{r}_{\text{ns}}(t_{\text{ns}}) - \vec{r}_d(t)|}{c} + \Delta . \quad (5.10)$$

Equation 5.8 accounts for the time dilation observed due to the speed v_{ns} of the neutron star with respect to the centre of mass of the solar system, commonly known as the solar system barycentre (SSB). Equations 5.9 and 5.10 account for the Doppler modulation of the emitted phase when observed in an Earth-based detector. The co-ordinate times at the neutron star, the SSB,

and at Earth are given by t_{ns} , t_{ssb} , and t respectively. The vectors \vec{r}_{ns} and \vec{r}_{d} denote the positions of the neutron star and of the detector with respect to the SSB reference frame. The Δ denotes additional relativistic terms due to the wave's propagation through the solar system (Taylor & Weisberg 1989, Abbott *et al.* 2007b): the Einstein delay due to gravitational red shift and time dilation from the planets, and the Shapiro delay due to the waves passing through curved spacetime close to the Sun. If the neutron star is in a binary system, additional corrections are required.

Together, equations 5.8–5.10 give the phase at the detector:

$$\Phi(t) = \Phi(\tau(t)). \quad (5.11)$$

For isolated neutron stars, we can assume $v_{\text{ns}} \approx 0$ (Jaranowski & Królak 1999). It follows that τ , t_{ns} , and t_{ssb} now differ only by a constant term (Abbott *et al.* 2007b), and that we need only consider the motion of the Earth with respect to the SSB, given by \vec{r}_{d} . The observed phase $\Phi(t)$ simplifies to

$$\Phi(t) = \phi_0 + 2\pi \sum_{n=0}^N f^{(n)} \frac{(t-t_0)^{n+1}}{(n+1)!} + 2\pi \frac{\hat{n} \cdot \vec{r}_{\text{d}}(t)}{c} \sum_{n=0}^N f^{(n)} \frac{(t-t_0)^n}{n!}, \quad (5.12)$$

where the vector \hat{n} points from the SSB to the neutron star. The spindown parameters $f^{(n)}$ are the time derivative of the frequency evaluated at $t = t_0$ in an inertial frame at rest with respect to the SSB. They coincide with the intrinsic spindowns $f_{\text{w}}^{(n)}$ only when $v_{\text{ns}} = 0$. The number of required spindowns N depends on the age of the neutron star and the time it is observed for. Typically, it is only necessary to include the initial frequency $f \equiv f^{(0)}$, and the first spindown $\dot{f} \equiv f^{(1)}$. Searches for young neutron stars, such as the search described in Chapters 8 and 9, may also require a second spindown $\ddot{f} \equiv f^{(2)}$.

The transformation from the wave frame $(\hat{x}_{\text{w}}, \hat{y}_{\text{w}}, \hat{z}_{\text{w}})$ to the detector frame $(\hat{x}_{\text{d}}, \hat{y}_{\text{d}}, \hat{z}_{\text{d}})$ can be written as

$$(\hat{x}_{\text{d}}, \hat{y}_{\text{d}}, \hat{z}_{\text{d}}) = \mathbf{M}_{\text{cd}}(t)(\hat{x}_{\text{c}}, \hat{y}_{\text{c}}, \hat{z}_{\text{c}}) = \mathbf{M}_{\text{wd}}(t)(\hat{x}_{\text{w}}, \hat{y}_{\text{w}}, \hat{z}_{\text{w}}), \quad (5.13)$$

where

$$\mathbf{M}_{\text{wd}}(t) = \mathbf{M}_{\text{cd}}(t)\mathbf{M}_{\text{wc}}. \quad (5.14)$$

The intermediate frame, $(\hat{x}_{\text{c}}, \hat{y}_{\text{c}}, \hat{z}_{\text{c}})$, is the celestial sphere frame, where \hat{z}_{c} points towards the celestial north pole and \hat{x}_{c} towards the vernal point. The two matrices \mathbf{M}_{wc} and $\mathbf{M}_{\text{cd}}(t)$ are Euler rotations:

$$\mathbf{M}_{\text{wc}} = \mathbf{R}_{\text{zxx}} \left(-\psi, -\frac{\pi}{2} - \delta, \frac{\pi}{2} - \alpha \right), \quad (5.15)$$

$$\mathbf{M}_{\text{cd}}(t) = \mathbf{R}_{\text{zxx}} \left(\Phi_{\text{s}}(t) + \frac{\pi}{2}, \frac{\pi}{2} - \lambda, \gamma - \frac{\zeta}{2} \right), \quad (5.16)$$

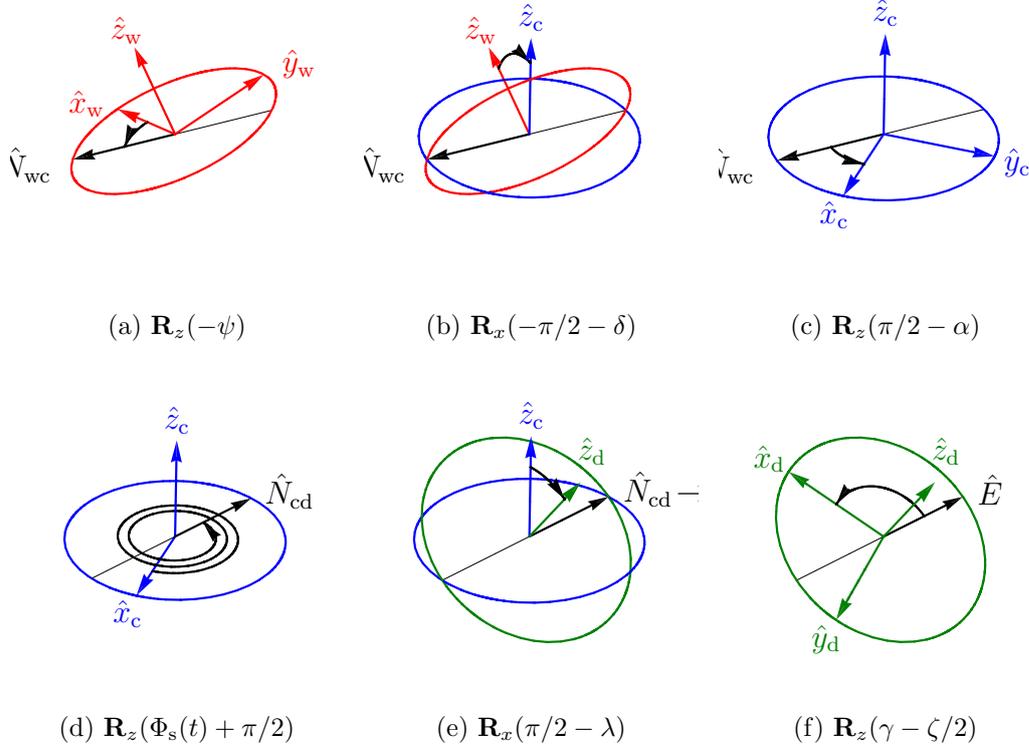


Figure 5.1: The transformation from the wave frame (red) to the detector frame (green), via the celestial sphere frame (blue). The top row corresponds to equation 5.15, and the bottom row to equation 5.16. The example rotations are for $\psi = 0.25\pi$, $\delta = -0.15\pi$, $\alpha = 0.25\pi$, $\Phi_s = 4.35\pi$, $\lambda = 0.3\pi$, $\gamma = 0.9\pi$, and $\zeta = \pi/2$. The vectors $\hat{N}_{wc} = \hat{z}_c \times \hat{z}_w / |\hat{z}_c \times \hat{z}_w|$, $\hat{N}_{cd} = \hat{z}_c \times \hat{z}_d / |\hat{z}_c \times \hat{z}_d|$, and \hat{E} is the local detector East.

where

$$\mathbf{R}_{zxz}(\theta_1, \theta_2, \theta_3) = \mathbf{R}_z(\theta_3)\mathbf{R}_x(\theta_2)\mathbf{R}_z(\theta_1) \quad (5.17)$$

is given in terms of the rotations

$$\mathbf{R}_x(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & -\sin \theta & \cos \theta \end{pmatrix}, \quad \mathbf{R}_z(\theta) = \begin{pmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (5.18)$$

The angles introduced in equations 5.15–5.16 are: the right ascension α and declination δ of the source; the local sidereal time $\Phi_s(t)$ and latitude λ of the detector; and γ , the angle of the bisector of the detector arms with respect to local East. The composition of the rotations which transform $(\hat{x}_w, \hat{y}_w, \hat{z}_w)$

to $(\hat{x}_d, \hat{y}_d, \hat{z}_d)$ is illustrated in Figure 5.1. The metric perturbation in the detector frame is now given by

$$\mathbf{H}_d(t) = \mathbf{M}_{\text{wd}}(t)\mathbf{H}_w(t)\mathbf{M}_{\text{wd}}^T(t). \quad (5.19)$$

The signal waveform, or template, is often written in the form

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t), \quad (5.20)$$

where $F_+(t)$ and $F_\times(t)$ are known variously as the antenna response or beam patterns functions. Expression for $F_+(t)$ and $F_\times(t)$ are given in, e.g., Bonazzola & Gourgoulhon (1996), Jaranowski *et al.* (1998), and Srivastava & Sahay (2002). They may be written compactly in terms of the polarisation bases and the detectors response, as follows:

$$F_{+,\times}(t) = \text{tr} [\mathbf{R}_d^T \mathbf{H}_d^{+,\times}(t)], \quad (5.21)$$

where

$$\mathbf{H}_d^{+,\times}(t) = \mathbf{M}_{\text{wd}}(t)\mathbf{H}_w^{+,\times}\mathbf{M}_{\text{wd}}^T(t), \quad (5.22)$$

and the superscript $^{+,\times}$ denotes a choice of either $+$ or \times . We note that the polarisation bases $\mathbf{H}_d^{+,\times}(t)$ are time-dependent, and are functions of both parameters of the source (α, δ, ψ) and of the detector $(\Phi_s(t), \lambda, \gamma, \zeta)$.

We may equivalently define the beam pattern functions using quantities with respect to the celestial sphere frame, thus:

$$F_{+,\times}(t) = \text{tr} [\mathbf{R}_c^T(t) \mathbf{H}_c^{+,\times}]. \quad (5.23)$$

In this frame, the polarisation bases are

$$\mathbf{H}_c^{+,\times} = \mathbf{M}_{\text{wc}}\mathbf{H}_w^{+,\times}\mathbf{M}_{\text{wc}}^T, \quad (5.24)$$

and the response matrix

$$\mathbf{R}_c(t) = \frac{1}{2} \left[\hat{n}_c^1(t) \otimes \hat{n}_c^1(t) - \hat{n}_c^2(t) \otimes \hat{n}_c^2(t) \right] \quad (5.25)$$

is given by the interferometer arms

$$\hat{n}_c^k(t) = \mathbf{M}_{\text{cd}}^T(t)\hat{n}_d^k. \quad (5.26)$$

We note that, in equation 5.23, the polarisation bases $\mathbf{H}_c^{+,\times}$ are now time-independent and functions only of the source parameters.

It is useful to expand the response matrix $\mathbf{R}_c(t)$ in terms of its time dependence, which comes from the local sidereal time $\Phi_s(t)$ of the detector.

It turns out that $\mathbf{R}_c(t)$ can be written in terms of five time-independent matrices $\mathbf{R}_c[n]$, with index $-2 \leq n \leq 2$, as follows:

$$\mathbf{R}_c(t) = \mathbf{R}_c[0] + \sum_{n=1}^2 \left\{ \mathbf{R}_c[n] \cos n\Phi_s(t) + \mathbf{R}_c[-n] \sin n\Phi_s(t) \right\}. \quad (5.27)$$

If we define the coefficients

$$\bar{F}_{+, \times}[n] = \text{tr} \left[\mathbf{R}_c^T[n] \mathbf{H}_c^{+, \times} \right], \quad (5.28)$$

then the beam patterns may also be written as

$$F_{+, \times}(t) = \bar{F}_{+, \times}[0] + \sum_{n=1}^2 \bar{F}_{+, \times}[n] \cos n\Phi_s(t) + \bar{F}_{+, \times}[-n] \sin n\Phi_s(t). \quad (5.29)$$

We see that the beam patterns have components which are periodic over one sidereal day ($|n| = 1$) and one half sidereal day ($|n| = 2$), in addition to a constant component ($n = 0$).

Using equation 5.29, the signal waveform $h(t)$ may also be written as

$$h(t) = \sum_{n=-2}^2 \left\{ \bar{h}[1, n] \cos \left(\Phi(t) + n\Phi_s(t) \right) + \bar{h}[-1, n] \sin \left(\Phi(t) + n\Phi_s(t) \right) \right\}, \quad (5.30)$$

with the coefficients

$$\bar{h}[k, n] = \frac{1}{1 + |\text{sign } n|} \left\{ A_+ \text{sign} \left(k + \text{sign}(n) + 1 \right) \bar{F}_+ \left[k|n| \right] - A_\times \text{sign} \left(k + \text{sign}(n) - 1 \right) \bar{F}_\times \left[-k|n| \right] \right\}. \quad (5.31)$$

It is straightforward, from equation 5.30, to obtain the Fourier transform and power spectrum of the signal. The amplitude modulation generates four sidebands of the signal frequency observed in the detector, at ± 1 and ± 2 times the sidereal frequency; the phase modulation and any intrinsic spindown of the signal will generate additional sidebands.

5.2.2 The \mathcal{F} statistic

In common with the stochastic background, the gravitational wave strain from sources of periodic gravitational waves is much weaker than the background detector noise. Nevertheless, because the signal model is accurately

known, it can be used to filter the data for any variation that matches the desired waveform, and thus improve the signal-to-noise ratio. This is known as coherent matched filtering, and is generally the most sensitive method of extracting a known signal, given a fixed observation time. In the context of periodic gravitational wave data analysis, the method is known as the \mathcal{F} statistic. It was first derived by Jaranowski *et al.* (1998) for a single detector, and later extended to multiple detectors by Cutler & Schutz (2005). Its derivation is outlined in this section.

Consider the output of a gravitational wave detector, $x(t)$. We first hypothesise that the output contains only detector noise $n(t)$, which is Gaussian (with zero mean), and stationary in time. The probability of $x(t)$, under this hypothesis, is given by a multivariate Gaussian distribution:

$$p(x|\text{noise}) = \mathcal{N} \exp \left[-\frac{1}{2}(x||x) \right], \quad (5.32)$$

with the appropriate normalisation \mathcal{N} . The inner product is defined as

$$(x||y) = 4\Re \left[\int_{-\infty}^{\infty} df \tilde{x}(f) S_n^{-1}(f) \tilde{y}(f)^* \right], \quad (5.33)$$

where $S_n(f)$ is the one-sided power spectral density of the detector noise, and $\tilde{}$ denotes the Fourier transform, $*$ the complex conjugate, and \Re the real part. We now hypothesise that $x(t)$ also contains, in addition to noise, a gravitational wave signal $h(t, \vec{A}, \vec{\lambda})$. The signal template is a function of the amplitude (or extrinsic) parameters \vec{A} , and the Doppler (or intrinsic) parameters $\vec{\lambda}$; the reason for the distinct sets of parameters will be seen shortly. The probability of $x(t)$ is now equivalent to assuming that $x(t) - h(t, \vec{A}, \vec{\lambda}) = n(t)$ contains only detector noise, and is given by

$$p(x|\text{signal+noise}) = \mathcal{N} \exp \left[-\frac{1}{2}(x - h||x - h) \right]. \quad (5.34)$$

The figure of merit used in matched filtering is the likelihood ratio Λ , which is commonly written as the log-likelihood

$$\ln \Lambda = \ln \frac{p(x|\text{signal+noise})}{p(x|\text{noise})} \quad (5.35)$$

$$= (x||h) - \frac{1}{2}(h||h). \quad (5.36)$$

It gives the relative probability of $x(t)$ containing a periodic gravitational wave signal (the hypothesis of equation 5.34), as opposed to containing only

noise (the hypothesis of equation 5.32). The likelihood test compares the value of Λ , computed for a particular signal template $h(t, \vec{A}, \vec{\lambda})$, against a threshold Λ^* . If Λ exceeds Λ^* , a gravitational wave signal matching $h(t, \vec{A}, \vec{\lambda})$ has been detected. (Of course, in the real world of non-Gaussian noise, imperfect experiments and bug-prone software, this is only a first step; only after many further investigations and consistency checks would a detection *really* be claimed!) The value of Λ^* is set by: the false alarm rate α , the probability that a claimed detection is spurious; and the false dismissal rate β , the probability that a genuine signal is ignored. By the Neyman-Pearson lemma, the likelihood test is an optimal (or most powerful) test, in that it minimises the false dismissal rate at a given false alarm rate.

The derivation of the \mathcal{F} statistic presented here assumes a frequentist approach to the statistics of signal analysis. In the Bayesian approach, the likelihood ratio becomes the Bayes factor

$$\ln \Lambda + \ln \frac{\text{p}(\text{signal+noise})}{\text{p}(\text{noise})}, \quad (5.37)$$

where the new second term encodes the prior probabilities of the signal and of the noise. Prix & Krishnan (2009) have re-interpreted the \mathcal{F} statistic within a Bayesian framework, and shown that the Bayesian \mathcal{F} statistic, by taking the priors into account, is more efficient in certain circumstances.

The detection of a signal proceeds by trying to maximise Λ with respect to the signal parameters $\vec{A}, \vec{\lambda}$. Following Jaranowski *et al.* (1998), the signal model, described in section 5.2.1, may be written in the form

$$h(t, \vec{A}, \vec{\lambda}) = \sum_{i=1}^4 A_i h_i(t, \vec{\lambda}), \quad (5.38)$$

where the amplitude parameters

$$A_1 = A_+ \cos \phi_0 \cos 2\psi - A_\times \sin \phi_0 \sin 2\psi, \quad (5.39a)$$

$$A_2 = A_+ \cos \phi_0 \sin 2\psi + A_\times \sin \phi_0 \cos 2\psi, \quad (5.39b)$$

$$A_3 = -A_+ \sin \phi_0 \cos 2\psi - A_\times \cos \phi_0 \sin 2\psi, \quad (5.39c)$$

$$A_4 = -A_+ \sin \phi_0 \sin 2\psi + A_\times \cos \phi_0 \cos 2\psi, \quad (5.39d)$$

are functions of the signal amplitudes A_+ and A_\times , the polarisation angle ψ , and the initial signal phase ϕ_0 . Since $h(t, \vec{A}, \vec{\lambda})$ is linear in the A_i , it is straightforward to find values for them that maximise Λ : they are given by

$$A_i^{\max} = \sum_{j=1}^4 [\mathbf{M}^{-1}]_{ij}(x||h_j), \quad (5.40)$$

where the matrix \mathbf{M} has elements $M_{ij} = (h_i \| h_j)$, and \mathbf{M}^{-1} is its inverse. Because the inner product is typically computed over many cycles of the signal, only the components $M_{11} \approx M_{33}$, $M_{22} \approx M_{44}$, and $M_{12} = M_{21} \approx M_{34} = M_{43}$ are non-zero.

The log-likelihood, when maximised over \vec{A} , gives the \mathcal{F} statistic:

$$2\mathcal{F}(\vec{\lambda}) = \sum_{i=1}^4 A_i^{\max}(x \| h_i) = \sum_{i,j=1}^4 [\mathbf{M}^{-1}]_{ij}(x \| h_i)(x \| h_j). \quad (5.41)$$

Expanding this expression gives (Jaranowski *et al.* 1998)

$$2\mathcal{F} = \frac{2}{D} \left\{ B \left[(x \| h_1)^2 + (x \| h_3)^2 \right] + A \left[(x \| h_2)^2 + (x \| h_4)^2 \right] - 2C \left[(x \| h_1)(x \| h_2) + (x \| h_3)(x \| h_4) \right] \right\}, \quad (5.42)$$

where

$$\begin{aligned} A &\approx 2(h_1 \| h_1) \approx 2(h_3 \| h_3), & B &\approx 2(h_2 \| h_2) \approx 2(h_4 \| h_4), \\ C &\approx 2(h_1 \| h_2) \approx 2(h_3 \| h_4), & D &= AB - C^2. \end{aligned} \quad (5.43)$$

The \mathcal{F} statistic is a function of the remaining Doppler parameters $\vec{\lambda}$. These parameters include the sky position, the initial frequency and spindowns, and (if applicable) binary orbital parameters. In contrast to the \vec{A} parameters, the signal template $h(t, \vec{A}, \vec{\lambda})$ is not linear in the $\vec{\lambda}$ parameters; it is not practicable, therefore, to analytically find the maximum of the \mathcal{F} statistic with respect to $\vec{\lambda}$. Instead, we decide on a parameter space that encompasses values of $\vec{\lambda}$ that are of interest, and then select a finite number of $\vec{\lambda}$ that span the extent of the parameter in some way (see Chapter 7). Finally, we compute the \mathcal{F} statistic for each chosen $\vec{\lambda}$, and record values which are statistically significant. The above process constitutes a gravitational wave search.

In the absence of a signal, the four inner products $(x \| h_i)$ are Gaussian random variables, and are correlated due to the non-zero off-diagonal terms in \mathbf{M} . The correlation can be removed, however, via a suitable linear transformation (Jaranowski *et al.* 1998), and $2\mathcal{F}$ can then be written as the sum of squares of four independent Gaussian variables. The probability distribution of $2\mathcal{F}$ is, therefore, a χ^2 distribution with four degrees of freedom. Due to this property of $2\mathcal{F}$, quoted values of the \mathcal{F} statistic, here and in the literature (e.g. Abbott *et al.* 2007b), will almost always refer to the value of $2\mathcal{F}$, i.e. twice the value.

by partitioning the data $x(t)$ into intervals over which the noise can be regarded as constant, and then weighting each interval according to its noise (Cutler & Schutz 2005). With these assumptions, the inner products reduce to

$$(x||y) \approx \frac{2}{S_n(f)} \int_{-T/2}^{T/2} dt x(t)y(t), \quad (5.46)$$

$$(\vec{x}||\vec{y}) \approx \sum_{\text{detectors}} (x||y), \quad (5.47)$$

where f_0 is the initial frequency of the signal. For simplicity we consider only the single-detector case. We substitute the signal waveform, given by equation 5.30, into equation 5.45; the result is

$$\begin{aligned} \rho^2 = \sum_{m,n=-2}^2 \left\{ \bar{h}[1,m]\bar{h}[1,n] \left(\cos(\Phi + m\Phi_s) \middle| \cos(\Phi + n\Phi_s) \right) + \right. \\ \left. 2\bar{h}[1,m]\bar{h}[-1,n] \left(\cos(\Phi + m\Phi_s) \middle| \sin(\Phi + n\Phi_s) \right) + \right. \\ \left. \bar{h}[-1,m]\bar{h}[-1,n] \left(\sin(\Phi + m\Phi_s) \middle| \sin(\Phi + n\Phi_s) \right) \right\}. \quad (5.48) \end{aligned}$$

We model the sidereal motion of the Earth by $\Phi_s(t) = \phi_s + \Omega_s t$, where ϕ_s is an initial phase and Ω_s is the sidereal frequency. Since the noise is independent of frequency, the spindown of the signal is unimportant, and we model the signal phase by $\Phi(t) \approx 2\pi f_0 t$. Substituting into equation 5.48, and noting that $2\pi f_0 \gg \Omega_s$, the inner products evaluate to

$$\left(X(\Phi + m\Phi_s) \middle| Y(\Phi + n\Phi_s) \right) = \mathcal{S} \begin{cases} \cos[(n-m)\phi_s] & X = Y, \\ \sin[(n-m)\phi_s] & X \neq Y, \end{cases} \quad (5.49)$$

where X, Y denote either cos or sin, and

$$\mathcal{S} = \frac{T}{S_n} \text{sinc} \left[\frac{1}{2}(n-m)\Omega_s T \right]. \quad (5.50)$$

The dominant contributions to the value of ρ^2 come from the terms in equation 5.48 where $m = n$, and thus $\mathcal{S} \propto T$. Using only these terms, we can write

$$\rho^2 \approx \frac{T}{S_n} \sum_{n=-2}^2 (\bar{h}[1,n]^2 + \bar{h}[-1,n]^2). \quad (5.51)$$

The remaining terms, where $m \neq n$, generate sinusoidal perturbations with periods inversely proportional to $|n - m|$. Figure 5.2a shows ρ^2 plotted for

a week-long observation time; the contributions from the $m \neq n$ terms to the overall signal-to-noise ratio are small. Figure 5.2b shows examples of the probability distribution of $2\mathcal{F}$; both the mean and variance of the distributions increase with ρ^2 .

5.2.3 Computational cost

The computational cost of a search, i.e. the total time taken to run the analysis software, is an important factor in the design of search for periodic gravitational waves (e.g. Abbott *et al.* 2007b, 2008a). It is calculated by multiplying together the number of templates to be computed, the observation time of the data, and the computational cost per template per unit observation time. For a fully coherent search for a non-binary periodic source, the total computational cost generally scales with (e.g. Whitbeck 2006, Abbott *et al.* 2007b)

$$T_{\text{span}} \underbrace{f_{\text{max}}^2 T_{\text{span}}^2}_{\alpha, \delta} \underbrace{f_{\text{max}} T_{\text{span}}}_f \underbrace{f_{\text{max}} T_{\text{span}}^2}_f \underbrace{f_{\text{max}} T_{\text{span}}^3}_{\dot{f}} \cdots \underbrace{f_{\text{max}} T_{\text{span}}^{n+1}}_{f^{(n)}} \cdots, \quad (5.52)$$

where f_{max} is the maximum frequency searched, and T_{span} is the time difference between the last and first time-stamps of the data being searched. The parameters under the braces denote that the under-braced factor is included only if the parameter is searched over. The first unbraced T_{span} denotes the scaling with the observation time $T_{\text{obs}} = \eta T_{\text{span}}$. We assume the duty cycle $\eta \lesssim 1$ is roughly constant, and therefore does not scale significantly with T_{span} .

We briefly justify the origin of the scaling factors in equation 5.52; see sections 7.3.3 and 8.4.1 for further discussion regarding calculating the number of templates. As the number of cycles of a periodic signal increases with T_{span} , the template density must also increase, so that the phases of signal may still be matched by a sufficiently nearby template in the parameter space. If f_{max} is increased, and thus the range of searched frequencies, the parameter space of the spindown parameter also increases, as the values of the spindown parameter scale with the frequency (Jaranowski *et al.* 1998).

The computational cost of a coherent search for gravitational waves from a known pulsar (where the sky position, frequency and spindown are known) scales, according to equation 5.52, approximately with T_{span} ; i.e. the cost of processing the data set scales with its length. For multi-template coherent searches, however, the cost rapidly increases. For an all-sky search, with one spindown parameter, the cost scales as $f_{\text{max}}^4 T_{\text{span}}^6$; for a directed search, where only the frequency and spindowns are searched over, the cost scales

as $f_{\max}^2 T_{\text{span}}^4$ for one spindown, and $f_{\max}^3 T_{\text{span}}^7$ for two spindowns. In contrast, the improvement in sensitivity only scales as $T_{\text{obs}}^{1/2} \propto T_{\text{span}}^{1/2}$. In summary, the steep increase of the computational cost with T_{span} , and the relatively meagre increase in sensitivity, is a limiting factor for all multi-template periodic gravitational wave searches (Abbott *et al.* 2005a, 2007b, 2008a, 2009a,d).

There are, however, many innovative and successful approaches to addressing this limitation. One approach is to simply spread the computational cost over multiple CPUs. In practice, all gravitational wave searches are run on large-scale computer clusters with an ever-increasing number of CPUs. The ultimate fulfilment of this approach is the Einstein@Home distributed computing project; the search of the S4 data was distributed over approximately 100000 volunteered computers and took ~ 13000 CPU years (Abbott *et al.* 2009a).

A second approach is to attempt to find a less expensive, or more sensitive, data analysis technique. One technique is to discard the phase of the gravitational wave data, typically by calculating the power over short time intervals, which results in a semi-coherent search (section 6.1). While the sensitivity of semi-coherent methods increases with $T_{\text{obs}}^{1/4}$, more slowly than for a coherent search, the computational cost scaling with T_{span} is virtually eliminated; all the under-braced T_{span} are replaced with $T_{\text{coh}} \ll T_{\text{span}}$, the length of the time intervals over which the power is calculated. Overall, therefore, a semi-coherent search can often be more sensitive because it is able to search longer data sets at reasonable computational cost. In Chapter 6, we consider the PowerFlux semi-coherent method in further detail, and attempt to generalise it to estimate multiple parameters of the gravitational wave signal. Hierarchical schemes, which alternate coherent and semi-coherent stages, provide further gains in sensitivity (e.g. Brady & Creighton 2000, Cutler *et al.* 2005)

A third approach to reducing the computational cost is to decrease the required number of templates. The overall template density is determined by the maximum allowable mismatch, denoted μ_{\max} , between any periodic signal and its nearest template in the parameter space. A careful analysis of the function dependence of the mismatch on the difference in parameters of signal and template leads to distance function, or metric, on the parameter space, which determines how closely the templates need to be spaced. With the distance determined, it remains to decide on the geometry of the placement of templates with respect to each other; the mathematics of lattices and sphere coverings describes how to place the templates as thinly as possible, while ensuring the maximum allowed mismatch is bounded by μ_{\max} . Chapter 7 describes an algorithm which, following the above procedure, generates an optimally-placed template bank.

5.3 Summary

This chapter introduced the field of gravitational wave data analysis. We first summarised the results of previous and present-day searches for four categories of gravitational wave signals: compact binary coalescences, bursts, the stochastic background, and periodic gravitational waves. This last category is of particular interest to this thesis; therefore, we presented a detailed review of the analytic model of periodic gravitational wave signals, and the coherent matched filtering technique used to detect them. Finally, we discussed the often prohibitive computational cost of periodic gravitational waves; this issue provides the motivation for the work presented in the next two chapters of this thesis.

Chapter 6

Parameter estimation using generalised PowerFlux

In this chapter, we consider one approach to the issue of the high computational cost of fully coherent, multi-template searches for periodic gravitational waves, as discussed in section 5.2.3. This approach is to use semi-coherent data analysis methods, which we introduce in section 6.1. We then consider the semi-coherent PowerFlux method (Dergachev & Riles 2005) in further detail. We present a derivation of this method in section 6.2, and generalise the method to extract additional parameters of the gravitational wave signal in section 6.3. In section 6.4, we compare the detection efficiencies of the original and generalised PowerFlux methods.

See section 1.1 for author contributions and publications relevant to this chapter.

6.1 Semi-coherent data analysis

Semi-coherent data analysis methods are similar to coherent matched filtering (section 5.2.2); they both attempt to match the output of a detector to a template of the periodic gravitational wave signal. They differ in one important respect: whereas matched filtering uses both the amplitude and phase of the detector data, semi-coherent methods use only the amplitude information. For the same observation time, semi-coherent methods are thus less sensitive than matched filtering, due to the lost phase information. On the other hand, the computational cost of semi-coherent methods do not increase as rapidly as for matched filtering, as the observation time is increased. Semi-coherent methods can therefore be used for all-sky broadband searches

over long observation times, whereas an equivalent coherent search would be computationally prohibitive.

The basic procedure of the semi-coherent method is to divide the detector data into short segments, and take the discrete Fourier transform of each; the resulting data are known as Short Fourier Transforms (SFTs). Power as a function of frequency and time, computed from the SFTs, is then stacked and slid in frequency so that the frequency bins containing a particular signal are added together, taking account of the frequency evolution due to Doppler modulation and any intrinsic spindown.

The time span of the data segments used to generate the SFTs is designated T_{SFT} , and is typically chosen to be 30 minutes. Within this time span, signals from an isolated source with frequencies less than ~ 1000 Hz will not shift in frequency by more than half the width of a SFT frequency bin (Abbott *et al.* 2007b). For sources in binary systems, a shorter T_{SFT} can be chosen dependent on the size of the expected orbital Doppler modulation.

Three semi-coherent methods, which have been used in all-sky searches using LIGO data, are PowerFlux, StackSlide, and the Hough transform (Abbott *et al.* 2008a, and references therein). The StackSlide method gives a sum of the power weighted by detector noise, while the PowerFlux and Hough methods give sums weighted by both detector noise and the beam patterns; in addition, the Hough transform replaces the power with a binary count determined by whether the power in each SFT exceeds a set threshold. Of the three, in general the PowerFlux method is the most sensitive.

6.2 A derivation of the PowerFlux method

The PowerFlux method was first derived by Dergachev & Riles (2005). Here, we present an alternative derivation (Mendell & Wette 2008), which is also the starting point for the generalisations presented in the next section.

It follows from the definition of T_{SFT} that we can treat the beam patterns F_+ , F_\times , and the signal frequency at the detector, as constant over the duration of an SFT. Ignoring losses due to the difference between the signal frequency and central frequency of the nearest SFT bin, the normalised power of the signal is given by

$$\frac{2|\tilde{h}|^2}{T_{\text{SFT}}} = \frac{A_+^2 F_+^2 + A_\times^2 F_\times^2}{2} T_{\text{SFT}}, \quad (6.1)$$

where \tilde{h} is the discrete Fourier transform of $h(t)$ divided by the sample rate of the data, and it is understood that F_+ and F_\times are evaluated at the midpoint of each SFT. Equation 6.1 represents the expected power of an elliptically

polarised signal in one SFT. We now label the SFTs using an index α , and consider searching for a linearly polarised signal (i.e. we choose $A_{\times} = 0$) which is present in the sequence of SFTs. We represent the SFT data by \tilde{x}_{α} , and the power taken from the SFT bin closest in frequency to the expected signal by

$$P_{\alpha} = \frac{2|\tilde{x}_{\alpha}|^2}{T_{\text{SFT}}}. \quad (6.2)$$

A natural way to estimate A_{+}^2 , analogous to χ^2 minimisation, is to find the value of A_{+}^2 that minimises

$$g = \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[P_{\alpha} - \frac{A_{+}^2 F_{+\alpha}^2}{2} T_{\text{SFT}} \right]^2. \quad (6.3)$$

This quantity is a weighted sum of the squared difference between the power of the signal (equation 6.2) and of the data (equation 6.1 with $A_{\times} = 0$), where the weights S_{α} are the one-sided power spectral densities of the noise for the frequency bins used in each corresponding SFT. We set

$$\frac{\partial g}{\partial A_{+}^2} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{A_{+}^2 F_{+\alpha}^2}{2} T_{\text{SFT}} \right) F_{+\alpha}^2 T_{\text{SFT}} \right] = 0 \quad (6.4)$$

and solve for A_{+}^2 , obtaining

$$A_{+}^2 = 4 \sum_{\alpha} \frac{F_{+\alpha}^2 |\tilde{x}_{\alpha}|^2}{S_{\alpha}^2 T_{\text{SFT}}^2} \bigg/ \sum_{\alpha} \frac{F_{+\alpha}^4}{S_{\alpha}^2}. \quad (6.5)$$

Equation 6.5 is the PowerFlux detection statistic given in Dergachev & Riles (2005), arrived at by maximising the signal-to-noise ratio. The two approaches appear to be equivalent (Abbott *et al.* 2008a). To distinguish it from the generalisations to follow, we refer to this method as *linear PowerFlux*. To extend the above derivation to circularly polarised signals ($A_{+} = A_{\times}$), we replace $F_{+\alpha}^2$ with $(F_{+\alpha}^2 + F_{\times\alpha}^2)$ in equation 6.3; this gives what we will refer to as the *circular PowerFlux* method.

To be sensitive to elliptically polarised signals ($A_{+} \neq A_{\times} \neq 0$), the implementation of PowerFlux used in LIGO searches (Abbott *et al.* 2008a, 2009b) uses linear PowerFlux, with a search over discrete values of the polarisation angle ψ , together with circular PowerFlux. This approach, however, does not directly estimate the amplitudes of elliptically polarised signals, nor does it directly recover ψ .

6.3 Generalisations of PowerFlux

In this section, we investigate whether the PowerFlux method can be extended to extract additional amplitude and polarisation information, while still using only the power computed from SFTs.

Mendell & Wette (2006, Section 2) investigated using the real and imaginary parts of the SFTs to estimate A_+ , A_\times , and ψ , using the method in Jaranowski *et al.* (1998). The method works well for long time-span SFTs, but fails to remain robust when $T_{\text{SFT}} \sim 30$ min. This is perhaps not surprising. All the parameter estimation methods investigated in Mendell & Wette (2006), in Mendell & Wette (2008), and presented in this chapter, assume that the beam patterns functions, F_+ and F_\times , are linearly independent over the data set from which the amplitude parameters are estimated. While this assumption is true over the time-span of a large set of SFTs (which is used by all of the other investigated methods), it is only marginally valid over the time-span of a single 30-min SFT. Additionally, the search over frequency would begin to scale in the same manner in a coherent search. We do not consider this approach further.

6.3.1 Estimation of A_+^2 and A_\times^2

We first extend PowerFlux to simultaneously estimate A_+^2 and A_\times^2 . The natural generalisation of equation 6.3 is

$$g = \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[P_{\alpha} - \frac{A_+^2 F_{+\alpha}^2 + A_\times^2 F_{\times\alpha}^2}{2} T_{\text{SFT}} \right]^2. \quad (6.6)$$

Following the same minimisation procedure as in section 6.2, we obtain

$$\frac{\partial g}{\partial A_+^2} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{A_+^2 F_{+\alpha}^2 + A_\times^2 F_{\times\alpha}^2}{2} T_{\text{SFT}} \right) F_{+\alpha}^2 T_{\text{SFT}} \right] = 0, \quad (6.7a)$$

$$\frac{\partial g}{\partial A_\times^2} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{A_+^2 F_{+\alpha}^2 + A_\times^2 F_{\times\alpha}^2}{2} T_{\text{SFT}} \right) F_{\times\alpha}^2 T_{\text{SFT}} \right] = 0. \quad (6.7b)$$

We then solve for A_+^2 and A_\times^2 ; the result is

$$A_+^2 = \frac{4}{\mathcal{D}} \left[\sum_{\alpha} \frac{F_{\times\alpha}^4}{S_{\alpha}^2} \sum_{\alpha} \frac{F_{+\alpha}^2 |\tilde{x}_{\alpha}|^2}{S_{\alpha}^2 T_{\text{SFT}}^2} - \sum_{\alpha} \frac{F_{+\alpha}^2 F_{\times\alpha}^2}{S_{\alpha}^2} \sum_{\alpha} \frac{F_{\times\alpha}^2 |\tilde{x}_{\alpha}|^2}{S_{\alpha}^2 T_{\text{SFT}}^2} \right], \quad (6.8a)$$

$$A_\times^2 = \frac{4}{\mathcal{D}} \left[\sum_{\alpha} \frac{F_{+\alpha}^4}{S_{\alpha}^2} \sum_{\alpha} \frac{F_{\times\alpha}^2 |\tilde{x}_{\alpha}|^2}{S_{\alpha}^2 T_{\text{SFT}}^2} - \sum_{\alpha} \frac{F_{+\alpha}^2 F_{\times\alpha}^2}{S_{\alpha}^2} \sum_{\alpha} \frac{F_{+\alpha}^2 |\tilde{x}_{\alpha}|^2}{S_{\alpha}^2 T_{\text{SFT}}^2} \right], \quad (6.8b)$$

where

$$\mathcal{D} = \sum_{\alpha} \frac{F_{+\alpha}^4}{S_{\alpha}^2} \sum_{\alpha} \frac{F_{\times\alpha}^4}{S_{\alpha}^2} - \left[\sum_{\alpha} \frac{F_{+\alpha}^2 F_{\times\alpha}^2}{S_{\alpha}^2} \right]^2. \quad (6.9)$$

We refer to this method as *generalised PowerFlux I*. Note that, for $\mathcal{D} \neq 0$, $F_{+\alpha}^2$ and $F_{\times\alpha}^2$ must be linearly independent functions of α , which should be true over a reasonable number of SFTs. A natural detection statistic is $A_+^2 + A_{\times}^2$. Its evaluation requires 2.5 times as many summations as for linear PowerFlux, and we must still include a search over discrete values of ψ .

6.3.2 Estimation of A_+^2 , A_{\times}^2 , and ψ

To further generalise PowerFlux to directly estimate ψ , we note that the antenna patterns F_+ and F_{\times} can be written in terms of two functions, a and b , which are independent of ψ (Jaranowski *et al.* 1998):

$$F_+(\psi, t) = \sin \zeta (a(t) \cos 2\psi + b(t) \sin 2\psi), \quad (6.10a)$$

$$F_{\times}(\psi, t) = \sin \zeta (b(t) \cos 2\psi - a(t) \sin 2\psi). \quad (6.10b)$$

where ζ is the angle between the interferometer arms. The normalised signal power is then written as

$$\frac{2|\tilde{h}_{\alpha}|^2}{T_{\text{SFT}}} = \frac{\mathcal{A}a_{\alpha}^2 + \mathcal{B}b_{\alpha}^2 + \mathcal{C}a_{\alpha}b_{\alpha}}{2} T_{\text{SFT}}, \quad (6.11)$$

where the coefficients are

$$\mathcal{A} = \sin^2 \zeta (A_+^2 \cos^2 2\psi + A_{\times}^2 \sin^2 2\psi), \quad (6.12a)$$

$$\mathcal{B} = \sin^2 \zeta (A_+^2 \sin^2 2\psi + A_{\times}^2 \cos^2 2\psi), \quad (6.12b)$$

$$\mathcal{C} = \sin^2 \zeta (A_+^2 - A_{\times}^2) 2 \cos 2\psi \sin 2\psi. \quad (6.12c)$$

Following the minimisation procedure again, we redefine

$$g = \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[P_{\alpha} - \frac{\mathcal{A}a_{\alpha}^2 + \mathcal{B}b_{\alpha}^2 + \mathcal{C}a_{\alpha}b_{\alpha}}{2} T_{\text{SFT}} \right]^2; \quad (6.13)$$

the \mathcal{A} , \mathcal{B} , and \mathcal{C} which minimise g are then found by solving

$$\frac{\partial g}{\partial \mathcal{A}} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{\mathcal{A}a_{\alpha}^2 + \mathcal{B}b_{\alpha}^2 + \mathcal{C}a_{\alpha}b_{\alpha}}{2} T_{\text{SFT}} \right) a_{\alpha}^2 T_{\text{SFT}} \right] = 0, \quad (6.14a)$$

$$\frac{\partial g}{\partial \mathcal{B}} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{\mathcal{A}a_{\alpha}^2 + \mathcal{B}b_{\alpha}^2 + \mathcal{C}a_{\alpha}b_{\alpha}}{2} T_{\text{SFT}} \right) b_{\alpha}^2 T_{\text{SFT}} \right] = 0, \quad (6.14b)$$

$$\frac{\partial g}{\partial \mathcal{C}} = - \sum_{\alpha} \frac{1}{S_{\alpha}^2} \left[\left(P_{\alpha} - \frac{\mathcal{A}a_{\alpha}^2 + \mathcal{B}b_{\alpha}^2 + \mathcal{C}a_{\alpha}b_{\alpha}}{2} T_{\text{SFT}} \right) a_{\alpha} b_{\alpha} T_{\text{SFT}} \right] = 0. \quad (6.14c)$$

It is then straightforward to solve equations 6.12a-6.12c for the amplitudes A_+^2 , A_\times^2 , and the polarisation angle ψ .

This method is referred to as *generalised PowerFlux II*. It requires computing 4 times as many sums as the linear PowerFlux method; because a search over discrete values of ψ is no longer required, however, the computational cost of this method may be less overall. An alternate method of estimating A_+^2 , A_\times^2 , and ψ , which uses the output of the linear PowerFlux method at several fixed values of ψ , is given in Appendix A of Abbott *et al.* (2008a).

6.4 Comparison of the methods

In this section, we compare the original PowerFlux methods presented in section 6.2 to the generalised methods derived in section 6.3. Using Monte Carlo simulations, we determine the distributions of the estimated parameters (section 6.4.1) and the relative detection efficiencies (section 6.4.2). The detection statistic used for each method is $A_+^2 + A_\times^2$, with $A_\times^2 = 0$ in the case of linear PowerFlux. For simplicity, we do not restrict the estimations of A_+^2 and A_\times^2 to the physical region $A_+^2 \geq 0$ and $A_\times^2 \geq 0$, nor check if the actual minimum value of g is on the boundary of this region. An implementation of the StackSlide method (Abbott *et al.* 2008a) is also included in the comparison.

We first perform searches on 10000 sets of SFTs, containing only randomly generated Gaussian noise, and obtain, for each method, a distribution of their detection statistic in the absence of a signal. From these distributions we determine the thresholds for a 1% false alarm rate appropriate for each method. We then search 3000 sets of SFTs, each containing noise plus a simulated periodic signal from an isolated source. The strength of the injected signals are given by h_0/\mathcal{H} , where for convenience we define the normalisation $\mathcal{H} = (S/T_{\text{SFT}})^{1/2}$, and S is the one-sided power spectral density of the noise. The strain amplitude h_0 and inclination angle ι are for a nonaxisymmetric source (equation 2.3). For each of the 3000 simulated signals, $\cos \iota$ was chosen at random from the range $[-1, 1]$, resulting in elliptically polarised signals; ψ was also chosen at random from the range $[-\pi/4, \pi/4]$.

Other parameters of the signal are its sky position (right ascension and declination) and its frequency at the solar system barycentre. The right ascension has a minimal effect and is thus set to zero; results are found for a selection of values of the declination. The signal frequency is a constant, i.e. there is no spindown of the source. We use the detector configuration of the LIGO Hanford observatory, as given in Jaranowski *et al.* (1998). Finally, we

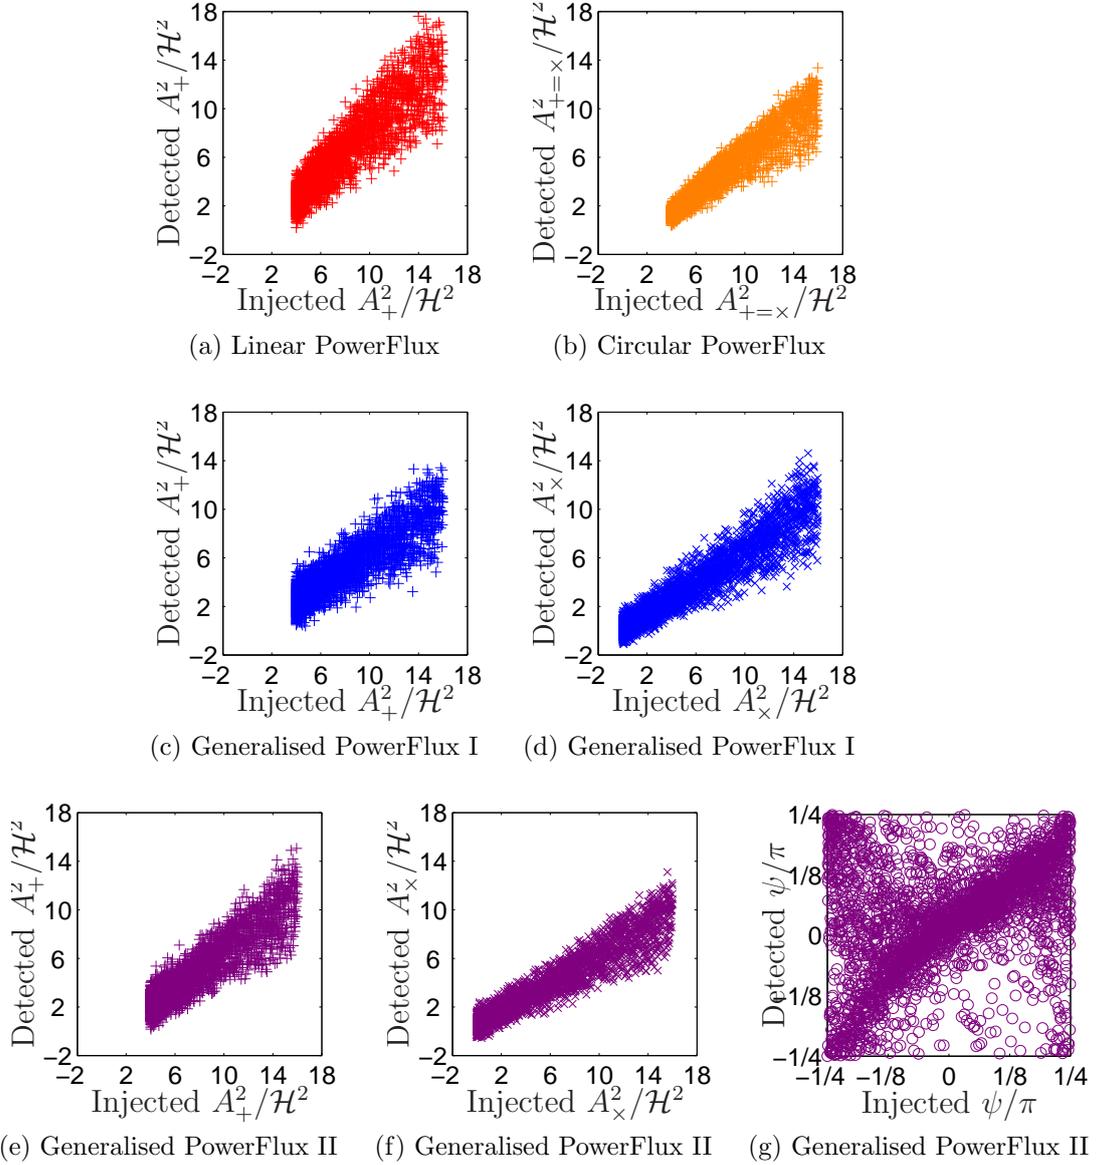


Figure 6.1: Injected versus detected normalised squared amplitudes, for $h_0/\mathcal{H} = 4$. The noise-only mean is first subtracted from the detected squared amplitudes. The methods used to estimate the parameters are linear (red) and circular (orange) PowerFlux, and generalised PowerFlux I (blue) and II (purple).

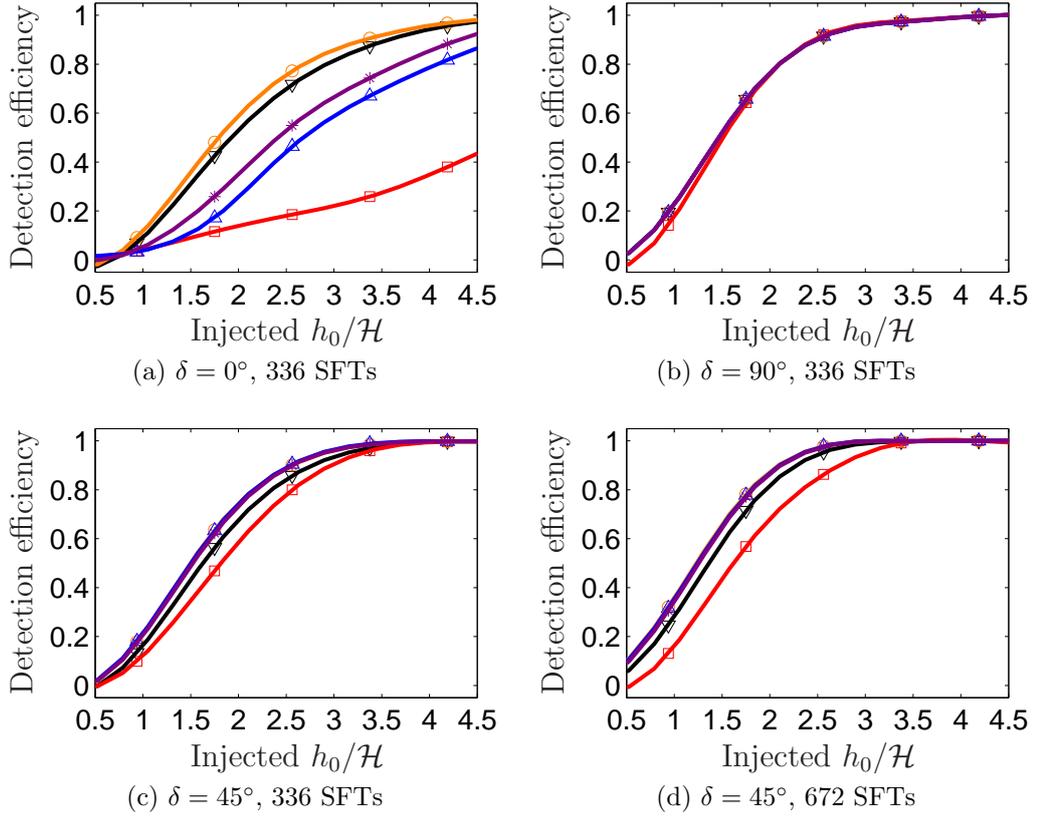


Figure 6.2: Detection efficiency, versus the normalised injected amplitude, of linear (red squares) and circular (orange circles) PowerFlux, generalised PowerFlux I (blue triangles) and II (purple stars), and StackSlide (black inverted triangles). The four plots differ in the declination of the injected signals, and in the number of SFTs, as labelled. The solid curves are spline fits to the data points. Errors in the determined efficiencies are less than 3%.

vary the number of SFTs.

A realistic search would include a template bank of sky positions, frequencies, spindowns, and (in some cases) polarisation angles. Here, we do not consider any mismatch in sky position or spindowns, i.e. we assume that, in these parameters, the search template is perfectly matched to the simulated signals. We do, however, include a random mismatch between the frequencies of the template and of the signal of up to half an SFT bin width; and, for searches requiring a value for ψ , we include a random mismatch in ψ between the template and the signal of up to $\pi/16$. The mismatched frequency and ψ (if needed) were used to compute the search template, along with the perfectly matched sky position and spindown. Thus, the simulations include losses due to mismatch between the template and the signal.

After searching the 3000 sets of SFTs, we obtain, for each method, a distribution of their detection statistic in the presence of a signal of normalised amplitude h_0/\mathcal{H} . The detection efficiency is then calculated as the fraction of the distribution of the detection statistic which falls above the appropriate 1% false alarm rate threshold. We repeat the above process of generating and searching 3000 sets of SFTs for increasing values of the normalised amplitude. The final result is the detection efficiency of each method as a function of h_0/\mathcal{H} .

To verify the results presented in the section, the authors of Mendell & Wette (2008) each independently wrote a MATLAB script to implement the above procedure. The two scripts produce detection efficiency curves, for the same input parameters, which are identical to within the uncertainties expected from the finite number of searches performed. This gives us confidence that the two implementations, and the results they give, are correct.

6.4.1 Parameter distributions

Figure 6.1 shows, for a search of 3000 sets of SFTs injected with signals of strength $h_0/\mathcal{H} = 4$, distributions of the injected versus the detected squared amplitudes obtained by the methods presented in sections 6.2–6.3. The mean value of the squared amplitudes estimated for noise alone is first subtracted from the detected squared amplitudes; both the injected and detected squared amplitudes are then normalised by \mathcal{H} . The detected amplitudes are typically smaller than the injected amplitudes, due to the loss of power caused by the mismatch between the signal and the template used to find it. The width of the distributions are due to noise; the circular PowerFlux and generalised PowerFlux II methods have, by visual inspection, the narrowest distributions.

Figure 6.1g shows the injected versus the detected polarisation angle as

estimated by generalised PowerFlux II. The estimation can be quite poor, for example, when noise causes the detected ψ to appear in the wrong quadrant.

6.4.2 Detection efficiencies

Figure 6.2 shows the detection efficiencies of the methods versus h_0/\mathcal{H} , for a selection of values of the number of SFTs and of the declination. As the source moves away from zero declination, the detection efficiencies increase, and the differences between the methods become smaller. The relative efficiencies of StackSlide, linear PowerFlux, and the generalised PowerFlux methods can change with respect to each other, but we note that circular PowerFlux typically remains the most efficient method. This is good news, since this method is already being used by the LSC. We note that the detection efficiencies increase with the number of SFTs, as expected, but the relative efficiencies of the methods do not change significantly.

One possible explanation for the efficiency of circular PowerFlux is as follows. We note that the weights used by the Hough transform (Abbott *et al.* 2008a) were proportional to $F_+^2 + F_\times^2$; these weights were shown to be optimal in an average sense (Sintes & Krishnan 2006, Krishnan & Sintes 2007). Next, we note that $A_+^2 F_+^2 + A_\times^2 F_\times^2$, when averaged over ψ , can be factored as $(F_+^2 + F_\times^2)(A_+^2 + A_\times^2)/2$, i.e. as the antenna pattern for circular polarisation multiplied by the sum of the squared amplitudes. This holds true even for an elliptically polarised signal. Since our simulations perform a Monte Carlo average over ψ , this may explain why circular PowerFlux performs the best.

6.5 Summary

In this chapter, we generalised the PowerFlux semi-coherent method to estimate the amplitude and polarisation parameters of periodic gravitational wave signals. We compared the parameter estimation and detection efficiencies of the generalised PowerFlux methods, thus obtained, against the standard PowerFlux methods, using simulated signals injected into Gaussian noise. While the relative performance of the various methods depends on the declination of the injected signals, in general we found that the standard circular PowerFlux method is the most efficient.

Chapter 7

Template bank generation using optimal lattices

In Chapter 6, we considered the use of semi-coherent data analysis methods as an approach to mitigating the high computational cost of coherent searches for periodic gravitational waves over large parameter spaces. In this chapter, we consider a different approach to the same issue. We present an algorithm that, given a space of template parameters, generates a bank of templates with the following properties: that the number of templates is minimised, reducing the computational cost of the search; and that any signal within the space is matched, up to a certain tolerance, by a template in the bank.

We first introduce a number of important concepts, which form the basis of the algorithm. In section 7.1, we review the mathematics of lattices and sphere coverings; in section 7.2, we introduce the metric of the template parameter space. The algorithm is presented in section 7.3. In section 7.4 we present an implementation of the algorithm, and test its performance.

The template bank generation algorithm is utilised in the gravitational wave search for Cassiopeia A presented in Chapters 8 and 9. In particular, section 8.4.1 discusses the application of the algorithm to the parameter space of this search.

7.1 Lattices and sphere coverings

This section provides a brief introduction to the mathematics of lattices and sphere coverings. For a comprehensive reference, see Conway & Sloane (1988a). Many of the concepts described in this section are illustrated, using the two-dimensional hexagonal lattice, in Figure 7.1.

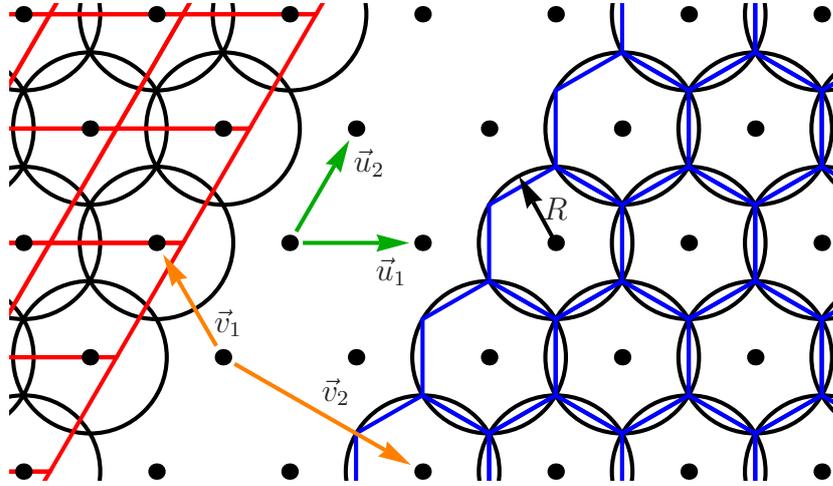


Figure 7.1: The hexagonal lattice (black points), and its covering spheres (black circles). Two possible generators of the lattice are $\mathbf{G} = (\vec{u}_1 \vec{u}_2)$ (green vectors), and $\mathbf{G}' = (\vec{v}_1 \vec{v}_2)$ (orange vectors). The radius of the covering spheres is R . Two choices for the fundamental region are the parallelograms (red) and the hexagons (blue); both have the same area $\approx 2.6R^2$. The hexagons are also the Voronoi cells.

7.1.1 Lattices

A lattice is a collection of points which are generated by tiling a single building block repeatedly across a space in all directions (Conway & Sloane 1988d). Each copy of the building block either contains a single lattice point, or else intersects multiple lattice points at its boundary, with the fractions of each point within the building block adding up to a single lattice point. If two points are in the lattice, then their sum and difference (computed using vectors to each point, with respect to a common origin) will also be in the lattice. In this way, lattices can be generated using only a minimal basis set of lattice points.

More formally, we define a n -dimensional lattice \mathcal{L} to be the set of points

$$\mathcal{L} = \{ \mathbf{G} \vec{k} \}, \quad (7.1)$$

where \vec{k} is any n -dimensional vector with integer entries (Conway & Sloane 1988d). The lattice exists either in an n -dimensional space \mathbb{R}^n (when $m = n$), or else in an n -dimensional sub-space of a larger m -dimensional space \mathbb{R}^m (when $m > n$). The matrix \mathbf{G} is known as the generator matrix, and has dimensions $m \times n$, where $m \geq n$. The columns of \mathbf{G} are the basis vectors

$\vec{u}_1, \dots, \vec{u}_n$ of the lattice.¹

The building block is also known as the fundamental region. One possible fundamental region is the parallelotope defined by the basis vectors of the lattice. Another choice is the Voronoi cell, which is the region centred on each lattice point that contains those points at least as close to the central lattice point as to any other point in the lattice. The volume of the fundamental region, regardless of its shape, is given by

$$V = \sqrt{\det \mathbf{G}^T \mathbf{G}}, \quad (7.2)$$

where ^T denotes transpose. The quantity $\mathbf{G}^T \mathbf{G}$ is known as the $(n \times n)$ Gram matrix of the lattice, and its determinant is known as the lattice determinant or discriminant.

Two lattices $\mathcal{L}, \mathcal{L}'$ are equivalent if their respective generator matrices \mathbf{G}, \mathbf{G}' are related by

$$\mathbf{G}' = s\mathbf{Q}\mathbf{G}\mathbf{Z} \quad (7.3)$$

where s is a non-zero constant, \mathbf{Q} is a real orthogonal matrix, and \mathbf{Z} is a matrix with integer entries and determinant ± 1 . Equation 7.3 allows the following operations to be performed on lattices: scaling (via s), rotations (via \mathbf{Q}), and reflections and basis vector permutations (via \mathbf{Z}).

7.1.2 Sphere coverings

A sphere covering is a set of n -dimensional spheres, each centred on a member of a set of points, denoted \mathcal{P} , in an n -dimensional space \mathbb{R}^n (Conway & Sloane 1988c). The radius of the spheres, known as the covering radius R , is chosen to be the largest possible distance from any point in the space to its nearest point in \mathcal{P} . As a consequence, every point in \mathbb{R}^n is contained in at least one of the spheres; conversely, any set of spheres with radius less than R will *not* contain every point in \mathbb{R}^n .

An important property of a sphere covering is its thickness Θ , which is defined to be the average number of spheres enclosing any point in \mathbb{R}^n . The thickness is a measure of the efficiency of the covering: the closer Θ is to 1, the fewer the number of spheres per unit volume required to cover \mathbb{R}^n . Where the points \mathcal{P} form a lattice \mathcal{L} , the thickness of the associated sphere covering is given by

$$\Theta = \frac{V_n R^n}{V}, \quad (7.4)$$

¹ This work uses column vectors to represent lattice points; Conway & Sloane (1988a) uses row vectors. Thus, e.g. \mathbf{G} in this work is \mathbf{G}^T in Conway & Sloane.

Table 7.1: Thicknesses of a selection of lattice sphere coverings, up to $n = 6$. The listed properties are with respect to other coverings of the same dimension.

n	Name	Thickness		Properties
		Θ	θ	
1	$A_1^* \cong \mathbb{Z}$	1.000	0.500	Thinnest possible covering
2	A_2^*	1.209	0.385	Thinnest possible covering
	\mathbb{Z}^2	1.571	0.500	
3	A_3^*	1.464	0.349	Thinnest possible lattice covering
	\mathbb{Z}^3	2.721	0.650	
4	A_4^*	1.766	0.358	Thinnest possible lattice covering
	\mathbb{Z}^4	4.935	1.000	
5	A_5^*	2.124	0.404	Thinnest possible lattice covering
	\mathbb{Z}^5	9.195	1.747	
6	L_6^c	2.465	0.477	Thinnest lattice covering known
	A_5^*	2.551	0.494	
	\mathbb{Z}^6	17.44	3.375	

where V_n is the volume of a n -dimensional unit sphere, and V is the volume of the fundamental region (equation 7.2). The thickness is a property of the lattice, and is therefore independent of any particular choice of \mathbf{G} . It is often convenient to use the normalised thickness

$$\theta = \frac{\Theta}{V_n}, \quad (7.5)$$

which can be arranged to give the covering radius

$$R = \left(\theta \sqrt{\det \mathbf{G}^T \mathbf{G}} \right)^{1/n}. \quad (7.6)$$

It is desirable to know what is the thinnest possible covering in any given dimension; this is often referred to as the sphere covering problem. The current state of knowledge, for $n \leq 6$, is summarised in Table 7.1; for results in higher dimensions, see e.g. Conway & Sloane (1988a), Pridmore (2007b), and Schuermann & Vallentin (2009). The thinnest possible coverings are known for $n \leq 2$, and the thinnest possible lattice coverings are known for $n \leq 5$ (Conway & Sloane 1988c); in each case they are the coverings of the A_n^* lattices, which have normalised thicknesses

$$\theta(A_n^*) = \sqrt{n+1} \left[\frac{n(n+2)}{12(n+1)} \right]^{n/2}. \quad (7.7)$$

In two and three dimensions, A_2^* and A_3^* are more commonly known as the hexagonal and body-centred cubic lattice respectively. For $n = 1$, A_1^* is equivalent to the integer lattice \mathbb{Z} ; as n increases, the A_n^* covering becomes markedly thinner than that of \mathbb{Z}^n (Table 7.1). The thinnest possible coverings are not known for $n > 5$. The A_n^* lattice remains one of the thinnest known lattices for $5 < n < 23$ (Prix 2007b), although a number of lattices exist that are thinner, such as the L_6^c lattice (Schuermann & Vallentin 2008, 2009).

A sphere covering is similar to, but distinct from, a sphere *packing*. Every point in \mathbb{R}^n is contained in at *least* one covering sphere, but only in at *most* one packing sphere (Conway & Sloane 1988d). As its name implies, a sphere packing is obtained by e.g. packing ball bearings into a box.

A sphere covering can be thought of as representing the infinite number of points of \mathbb{R}^n by a finite subset, \mathcal{P} , subject to some resolution determined by R . To generate the template bank, we need to represent the infinite number of possible templates in the parameter space by a smaller subset (the template bank), subject to the requirement that any signal is matched (by a template in the bank) with up to a maximum tolerated loss in its signal-to-noise ratio. It seems likely, therefore, that a sphere covering could be utilised in constructing a bank of templates. In particular, a sphere covering which has the minimum possible thickness (which we shall refer to as optimal), then the equivalent template banks will cover the parameter space with a minimum number of templates, as desired.

7.2 The parameter space metric

To use a sphere covering to generate a template bank, we need to be able to transform the spacing of the templates into distances between lattice points. The template spacing is determined by how far apart in parameter space a signal can be from its nearest template before the loss in signal-to-noise ratio, due to the mismatch between signal and templates, exceeds a given maximum. A useful way of quantifying this to introduce a metric (Balasubramanian *et al.* 1996, Owen 1996) on the parameter space, from which distances between template points can be computed. The distance is chosen to be proportional to the loss in signal-to-noise ratio, if we consider one template to be a signal, against which the second template is being matched. The metric may be derived from the detection statistic, used in the search, or from another quantity proportional to the signal-to-noise ratio (see discussion in Prix 2007a).

To illustrate the construction of a parameter space metric, we use the

\mathcal{F} -statistic (section 5.2.2) as an example. The expectation value of $2\mathcal{F}$,

$$E\left\{2\mathcal{F}[x(t), h(t, \vec{A}^{\max}, \vec{\lambda})]\right\} \equiv 2\bar{\mathcal{F}}(\vec{\lambda}) \quad (7.8)$$

indicates the likelihood of having detected, in the data stream $x(t)$, a signal matching the template $h(t, \vec{A}^{\max}, \vec{\lambda})$, with maximised-over amplitude parameters \vec{A}^{\max} and searched-over Doppler parameters $\vec{\lambda}$. We are interested in $2\bar{\mathcal{F}}$ as a function only of the $\vec{\lambda}$, and therefore suppress the appearance of other parameters.

Suppose that the data contains a signal perfectly matching the template $h(\vec{\lambda})$; then $2\bar{\mathcal{F}}(\vec{\lambda})$ is clearly at a local maximum. At a nearby point $\vec{\lambda}' = \vec{\lambda} + \Delta\vec{\lambda}$, the template $h(\vec{\lambda}')$ is no longer perfectly matched to the signal, and thus $2\bar{\mathcal{F}}(\vec{\lambda}')$ will decrease relative to $2\bar{\mathcal{F}}(\vec{\lambda})$. If the difference $\Delta\vec{\lambda}$ is small, the difference between $2\bar{\mathcal{F}}(\vec{\lambda})$ and $2\bar{\mathcal{F}}(\vec{\lambda}')$ can be written as a second-order Taylor expansion:

$$2\bar{\mathcal{F}}(\vec{\lambda}') \approx 2\bar{\mathcal{F}}(\vec{\lambda}) + \frac{1}{2}\Delta\vec{\lambda} \cdot \mathbf{H}\Delta\vec{\lambda}. \quad (7.9)$$

Because $2\bar{\mathcal{F}}(\vec{\lambda})$ is at a local maximum, the first-order term, i.e. the gradient of $2\bar{\mathcal{F}}$ at $\vec{\lambda}$, is zero, and the second-order term contains a negative-definite Hessian matrix $\mathbf{H}(\vec{\lambda})$, with entries

$$[\mathbf{H}(\vec{\lambda})]_{ij} = \frac{\partial^2(2\bar{\mathcal{F}})}{\partial\lambda_i\partial\lambda_j}\bigg|_{\vec{\lambda}}. \quad (7.10)$$

Re-arranging equation 7.9, we define the *mismatch*

$$\mu(\vec{\lambda}, \vec{\lambda}') = \frac{2\bar{\mathcal{F}}(\vec{\lambda}) - 2\bar{\mathcal{F}}(\vec{\lambda}')}{2\bar{\mathcal{F}}(\vec{\lambda})} \quad (7.11)$$

$$= (\vec{\lambda}' - \vec{\lambda}) \cdot \mathbf{M}(\vec{\lambda})(\vec{\lambda}' - \vec{\lambda}) \quad (7.12)$$

to be the loss of signal-to-noise ratio, i.e. $2\bar{\mathcal{F}}$, due to the mismatch between the Doppler parameters of the signal, $\vec{\lambda}$, and the Doppler parameters of the template, $\vec{\lambda}'$. The mismatch is calculated from the difference in parameters, $\Delta\vec{\lambda} = \vec{\lambda}' - \vec{\lambda}$, and from the positive-definite matrix

$$\mathbf{M}(\vec{\lambda}) = -\frac{\mathbf{H}(\vec{\lambda})}{4\bar{\mathcal{F}}(\vec{\lambda})}. \quad (7.13)$$

Equation 7.12 is simply the expression for the distance μ between two points $\vec{\lambda}$, $\vec{\lambda}'$, given by the metric \mathbf{M} .

To construct the template bank, we require that the mismatch is bounded by a maximum mismatch

$$\mu_{\max} \geq \mu(\vec{\lambda}, \vec{\lambda}'), \quad (7.14)$$

to limit the loss in signal-to-noise ratio to any potential signal. The region defined by equation 7.14 is an n -dimensional ellipse, centred on $\vec{\lambda}$, with axes

$$\vec{a}_i = \hat{v}_i \sqrt{\frac{\mu_{\max}}{\nu_i}}, \quad (7.15)$$

where \hat{v}_i are the unit vectors along the eigenvectors of \mathbf{M} , and ν_i are the corresponding eigenvalues. In order to enforce equation 7.14 across the entire parameter space, one would choose a set of points \mathcal{P} , such that every point in the parameter space falls within at least one of the mismatch ellipses surrounding one of the points in \mathcal{P} ; this is precisely the definition of a sphere covering.

In section 7.3.2 we demonstrate that the mismatch ellipses, under a suitable coordinate transformation, are mapped onto the covering spheres. This completes the correspondence between the template bank and the sphere covering we use to construct it. For this construction to work, however, there is one further requirement; the value of \mathbf{M} must be *flat*, i.e. it must be independent of all parameter that are being searched over (i.e. $\vec{\lambda}$ in the above derivation). This is generally not the case for the full metric of the \mathcal{F} statistic, although it is true for particular subspaces of the full space of Doppler parameters.

7.2.1 The metric of the \mathcal{F} statistic

The derivation of the \mathcal{F} statistic metric presented above is a simplification of the full metric derived in Prix (2007a). In particular, the full metric of the \mathcal{F} statistic depends on the amplitude parameters \vec{A} , which leads to a family of \mathcal{F} statistic metrics parameterised by \vec{A} . The \vec{A} are, in general, unknown, and the variation of the metric family over all possible \vec{A} implies a range of mismatches. The full metric can therefore be used to compute the extremum of the range of mismatches, due to the unknown \vec{A} ; in practise, however, it is more useful to marginalise over the \vec{A} to obtain an average metric.

The metric of the \mathcal{F} statistic typically scales with some power of T_{span} , the total time span of the data (see section 5.2.3). When T_{span} is greater than several days, the metric of the \mathcal{F} statistic simplifies to a *phase* metric, where the amplitude modulation of the signal is ignored, and the overall mismatch

between signal and template is due only to the mismatch in their respective phases. The phase metric is given by

$$[\mathbf{M}^\Phi]_{ij} = \left\langle \frac{\partial\Phi(t)}{\lambda_i} \frac{\partial\Phi(t)}{\lambda_j} \right\rangle - \left\langle \frac{\partial\Phi(t)}{\lambda_i} \right\rangle \left\langle \frac{\partial\Phi(t)}{\lambda_j} \right\rangle, \quad (7.16)$$

where the phase $\Phi(t)$ is given by equation 5.12. The first term is equal to the Fisher information matrix; the second term arises from maximising the metric over the initial phase ϕ_0 . The phase metric has been computed for the single-detector \mathcal{F} statistic by Whitbeck (2006), and since Prix (2007a) found that the full \mathcal{F} statistic metric does not scale with the number of detectors,² the same metric can be used for the multi-detector \mathcal{F} statistic. Prix (2007a) also defines an *orbital* metric, which further neglects the phase modulation due to the sidereal motion of the detector, which gives similar performance.

The phase metric is not flat (see above) over the space of Doppler parameters $\alpha, \delta, f, \dot{f}, \ddot{f}$, etc. If the sky position is known, however, the phase metric over the remaining parameters f, \dot{f}, \ddot{f} , etc. is flat, and is given by (Wette *et al.* 2008)

$$[\mathbf{M}^{f,\dot{f},\ddot{f},\dots}]_{ij} = \frac{4\pi^2 T_{\text{span}}^{i+j+2} (i+1)(j+1)}{(i+2)!(j+2)!(i+j+3)}, \quad (7.17)$$

where the indices i, j take the values 0 for f , 1 for \dot{f} , etc. Thus, the template bank generation algorithm presented in this chapter may be used to search for gravitational waves from a periodic source with known sky position but unknown frequency evolution; such a search, directed at the supernova remnant Cassiopeia A, is presented in Chapters 8 and 9. Equation 7.17 is sometimes referred to as the *spindown-only* metric.

Finally, we note that the metric, because it is the result of a truncated Taylor expansion, is only a local approximation to the mismatch, and becomes less accurate as the difference between the parameters of signal and template increase. Fortunately, the tendency is for the metric to overestimate the expected mismatch (Prix 2007a). We verify this behaviour, for the spindown-only metric, in section 7.4.1.

7.3 A template bank generation algorithm

This section presents an algorithm for generating the minimal bank of templates to cover a given n -dimensional parameter space, subject to a maximum

² The sensitivity of an \mathcal{F} statistic search does, however, scale with the harmonic sum of the noise floors of the detectors; see section 8.5.

allowed mismatch (see also Prix 2007b). The algorithm uses a lattice to position points in the template parameter space. The order in which the lattice points are generated guarantees that every part of the parameter space will be visited (section 7.3.1). The transformation from the covering spheres of the lattice to the mismatch ellipses given by the parameter space metric is presented in section 7.3.2. In section 7.3.3 we show how to calculate the number of templates required for a complete coverage of the parameter space.

An algorithm based on the sphere coverings of lattices is not the only possible approach to the problem of template placement. A very different, but simple and effective approach, described in Messenger *et al.* (2009), is to simply place templates randomly in the parameter space. While complete coverage is not assured, the mean fractional loss in coverage can be controlled, through the number of generated templates, and the statistical distribution of the worst possible mismatch can be determined (Messenger *et al.* 2009, Manca & Vallisneri 2010). Control over the fractional loss only applies, however, within the bulk of the parameter space; at its edges, the coverage worsens considerably, and additional templates are required (Manca & Vallisneri 2010). In higher dimensions, starting at $n \sim 6$, and particularly for $n \gtrsim 10$, random template placement requires many fewer templates (by a factor of ~ 6 in $n = 19$ dimensions) than a lattice-based algorithm to cover the bulk of a parameter space, at the cost of a small ($\sim 1\%$) loss in coverage; this trade-off may well be appealing for high-dimensional, computationally expensive searches. Stochastic template placement (e.g. Harry *et al.* 2009) is similar to random templates, except that that algorithm additionally removes templates which, by chance, excessively overlap each other. Messenger *et al.* (2009) also presented a relaxed lattice-based placement strategy, where the mismatch μ is simply beyond the desired maximum μ_{\max} , again reducing the number of required templates at the cost of some loss in coverage due to $\mu > \mu_{\max}$. Manca & Vallisneri (2010) have presented a template placement algorithm based on a triangulation of the local (possibly curved) geometry of the parameter space; this approach does not require explicit knowledge of the parameter space metric.

7.3.1 Iteration over the parameter space

The parameter space points

$$\vec{p} = \mathbf{G}\vec{k} \quad (7.18)$$

are calculated using the generator of the lattice \mathbf{G} , and an n -dimensional vector of integers k . To generate all the points needed to cover the parameter space, an efficient method is needed of iterating over all possible integers k ,

so that every part of the space is visited, regardless of its shape. Such a method is presented in this section.

We require that the bounds on the parameters $\vec{p} = (p_1, \dots, p_n)$ follow the following pattern. The first parameter, p_1 , is bounded by two constants, p_1^{\min} and p_1^{\max} . The second parameter, p_2 , is bounded by p_2^{\min} and p_2^{\max} , which may be functions only of p_1 . The third parameter, p_3 , is bounded by p_3^{\min} and p_3^{\max} , which may be functions only of the preceding two parameters, p_1 and p_2 . This pattern is continued, until p_n is reached. In short, the bounds on any parameter may be functions only of the previous parameters, or constant if the parameter is p_1 . Thus, the parameter space is defined by the following region:

$$\begin{aligned} p_1^{\min} &\leq p_1 \leq p_1^{\max}, \\ p_2^{\min}(p_1) &\leq p_2 \leq p_2^{\max}(p_1), \\ p_3^{\min}(p_1, p_2) &\leq p_3 \leq p_3^{\max}(p_1, p_2), \end{aligned} \tag{7.19}$$

...

$$p_n^{\min}(p_1, p_2, \dots, p_{n-1}) \leq p_n \leq p_n^{\max}(p_1, p_2, \dots, p_{n-1}).$$

To generate the parameter space points, via equation 7.18, we need to transform the bounds on \vec{p} , given by equation 7.19, into the equivalent bounds on \vec{k} . It would be convenient if the nested structure of the bounds on \vec{p} , evident in equation 7.19, was also true of the bounds on \vec{k} , i.e.

$$\begin{aligned} k_1^{\min} &\leq k_1 \leq k_1^{\max}, \\ k_2^{\min}(k_1) &\leq k_2 \leq k_2^{\max}(k_1), \\ k_3^{\min}(k_1, k_2) &\leq k_3 \leq k_3^{\max}(k_1, k_2), \end{aligned} \tag{7.20}$$

...

$$k_n^{\min}(k_1, k_2, \dots, k_{n-1}) \leq k_n \leq k_n^{\max}(k_1, k_2, \dots, k_{n-1}).$$

If so, the parameter space could be easily iterated over using a sequence of nested loops. The outermost loop (over k_1) would take constant bounds (k_1^{\min} , k_1^{\max}); the first inner loop (over k_2) would calculate the value of its bounds ($k_2^{\min}(k_1)$, $k_2^{\max}(k_1)$) using the current value of k_1 ; the next inner loop (over k_3) would calculate the value of its bounds ($k_3^{\min}(k_1, k_2)$, $k_3^{\max}(k_1, k_2)$) using the current values of k_1 and k_2 ; and so the pattern would continue. For equation 7.20 to be true, however, the generator matrix \mathbf{G} must have a particular form. To see why, we transform the lower bound on p_i , $p_i^{\min}(p_{i-1}, p_{i-2}, \dots)$,

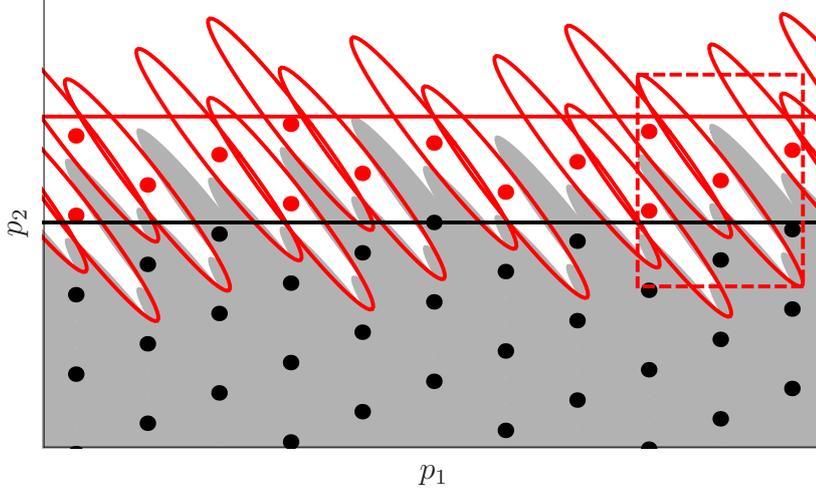


Figure 7.2: Edges effects at the boundary of a two-dimensional parameter space in p_1 and p_2 . Black template points are laid within the parameter space (the region below the black line, p_2^{\max}); the grey area indicates where the space is covered by the mismatch ellipses of the black templates. The red line, $p_2^{\max} + \partial p_2$, indicates an extension of the extent parameter space by half the height of the bounding box, $d_2^{\text{b.b.}}$, of a mismatch ellipse (red, dashed). Additional template points (plotted in red, with their mismatch ellipses) are laid within this boundary region. A hexagonal lattice is used; the generator matrix used is that of Figure 7.3d.

into the equivalent lower bound on k_i , using equation 7.18:

$$k_i^{\min}(k_{i-1}, k_{i-2}, \dots) = \left\lfloor \frac{1}{G_{ii}} \left\{ p_i^{\min} \left(\sum_{j=1}^n G_{(i-1)j} k_j, \sum_{j=1}^n G_{(i-2)j} k_j, \dots \right) - \sum_{\substack{j=1 \\ j \neq i}}^n G_{ij} k_j \right\} \right\rfloor, \quad (7.21)$$

where $\lfloor \cdot \rfloor$ is the floor function; the upper bound on k_i follows in a similar fashion. For k_i^{\min} to be a function only of k_{i-1} , k_{i-2} , etc., as implied by the left-hand side of equation 7.21, we require \mathbf{G} to be a *lower triangular* matrix, i.e.

$$G_{ij} = 0 \text{ for all } j > i. \quad (7.22)$$

In section 7.3, we show that a generator matrix in lower triangular form may be constructed for any lattice; equation 7.22 does not therefore limit the applicability of the algorithm.

The iteration procedure, while guaranteeing coverage of the bulk of the parameter space, will however fail to adequately cover the edges of the parameter space. This is illustrated in Figure 7.2. The geometry of how the ellipses fit together leads to gaps below the boundary of the parameter space. This is easily rectified, however, by adding a boundary layer around the edges of the parameter space, such that parameter space bounds become

$$- \partial p_i + p_i^{\min}(p_{i-1}, p_{i-1}, \dots) \leq p_i \leq p_i^{\max}(p_{i-1}, p_{i-1}, \dots) + \partial p_i, \quad (7.23)$$

where ∂p_i denotes³ the width of the boundary layer appropriate for the parameter p_i . An appropriate width for the boundary layer is given by the bounding box of the mismatch ellipse, which we define to be the smallest n -dimensional rectangle, with sides parallel to the directions of the coordinates p_i , that completely encloses the mismatch ellipse. It is sufficient to choose a boundary width ∂p_i that is half the extent of the bounding box, d_i^{bb} , in the p_i coordinate direction. This implies that a template should be included in the bank whenever any part of its mismatch ellipse is within the *original* bounds on the parameter space, given by equation 7.19.

An expression for the size of the bounding box in terms of the metric can be derived as follows. Consider an n -dimensional ellipse, centred at the origin, and given in terms of the coordinates \vec{x} and the positive-definite matrix \mathbf{M} by

$$\vec{x} \cdot \mathbf{M} \vec{x} = \mu. \quad (7.24)$$

We consider the intersection of the ellipse with a plane perpendicular to the direction of one of the x_i coordinates; without loss of generality, we choose the x_1 coordinate. We partition \mathbf{M} and \vec{x} into

$$\mathbf{M} = \begin{pmatrix} a & \vec{b}^{\text{T}} \\ \vec{b} & \mathbf{N} \end{pmatrix}, \quad \vec{x} = \begin{pmatrix} x_1 \\ \vec{y} \end{pmatrix}, \quad (7.25)$$

where a and x_1 are scalars, \vec{b} and \vec{y} are $n-1$ vectors, \vec{b}^{T} denotes the row vector form of \vec{b} , and \mathbf{N} is a $(n-1) \times (n-1)$ positive-definite matrix. Substituting equations 7.25 into equation 7.24 gives, after some rearrangement,

$$(\vec{y} - \vec{y}_0) \cdot \mathbf{N} (\vec{y} - \vec{y}_0) = \nu \quad (7.26)$$

where

$$\vec{y}_0 = -\mathbf{N}^{-1} \vec{b} x_1, \quad (7.27)$$

$$\nu = \mu - x_1^2 (a - \vec{b} \cdot \mathbf{N}^{-1} \vec{b}). \quad (7.28)$$

³ Here, the symbol ∂ does not mean partial differentiation, but is used in the spirit of its meaning in topology, where it denotes the boundary of a set.

With x_1 constant, equation 7.26 describes the ellipse defined by the intersection of equation 7.24 with the constant- x_1 plane. We note that when

$$x_1 = \pm \sqrt{\mu(a - \vec{b} \cdot \mathbf{N}^{-1} \vec{b})^{-1}}, \quad (7.29)$$

then $\nu = 0$, by equation 7.28, and the intersection ellipse shrinks to a point. This indicates that $|x_1|$ is at the furthest extent of the original ellipse (equation 7.24); if $|x_1|$ is increased further, $\nu < 0$ and equation 7.26 has no solution, i.e. there is no longer any intersection.

To simplify equation 7.29, we partition the inverse of \mathbf{M} in the same fashion as in equation 7.25, and equate

$$\mathbf{M}\mathbf{M}^{-1} = \begin{pmatrix} a & \vec{b}^T \\ \vec{b} & \mathbf{N} \end{pmatrix} \begin{pmatrix} c & \vec{d}^T \\ \vec{d} & \mathbf{P} \end{pmatrix} = \begin{pmatrix} 1 & \vec{0}^T \\ \vec{0} & \mathbf{1}_{n-1} \end{pmatrix} = \mathbf{1}_n, \quad (7.30)$$

where $\mathbf{1}$ is the identity matrix. This results in four equations, two of which,

$$ac + \vec{b} \cdot \vec{d} = 1, \quad \vec{b}c + \mathbf{N}\vec{d} = \vec{0}, \quad (7.31)$$

can be rearranged to give⁴

$$(a - \vec{b} \cdot \mathbf{N}^{-1} \vec{b})^{-1} = c = [\mathbf{M}^{-1}]_{11}. \quad (7.32)$$

Thus, $x_1 = \pm \sqrt{\mu[\mathbf{M}^{-1}]_{11}}$, and it follows that the extent of the bounding box in the x_i coordinate direction is given by

$$d_i^{\text{bb}} = 2\sqrt{\mu[\mathbf{M}^{-1}]_{ii}}. \quad (7.33)$$

7.3.2 The lattice generator

In this section, we construct a generator matrix for the template bank lattice. We start with the generator of a lattice with a optimally thin sphere covering, and transform it, using equation 7.3, into an equivalent lattice whose covering spheres are the mismatch ellipses of the metric. The generator will also be in lower triangular form, to satisfy the requirements of the previous section. To illustrate the transformation, we again use the hexagonal lattice, A_2^* , as an example. Prix (2007b) derives a similar procedure, but does not explicitly require that the generator be lower triangular.

⁴ Equation 7.32 gives the Schur complement c of the block \mathbf{N} of the matrix \mathbf{M} .

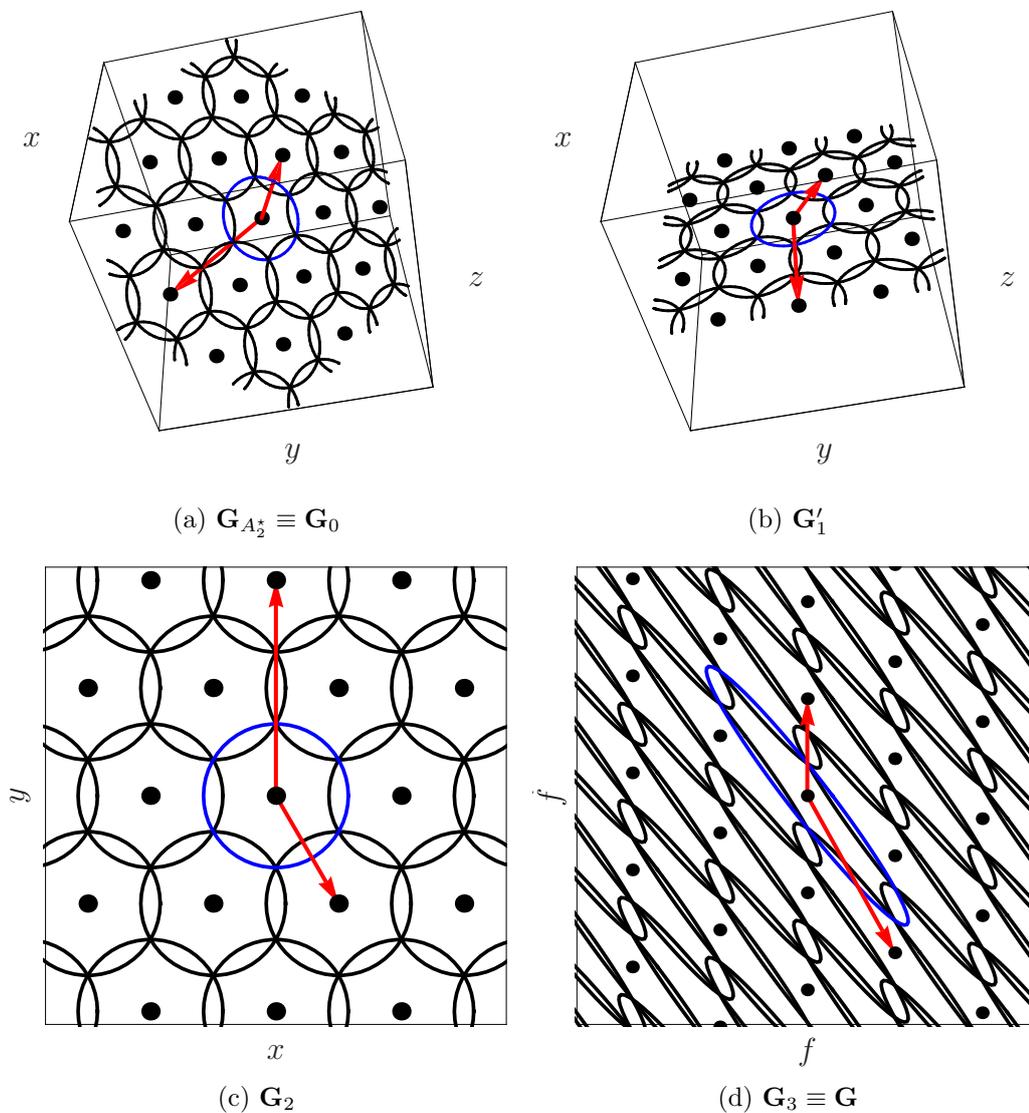


Figure 7.3: Transformation of the the original hexagonal lattice generator $\mathbf{G}_{A_2^*}$ (a), through intermediate stages \mathbf{G}'_1 (b) and \mathbf{G}_2 (c), to the final generator for a 2-dimensional template bank \mathbf{G} (d). One covering sphere/mismatch ellipse is plotted in blue for illustration. The red arrows represent the basis vectors given by the columns of the generator matrix. See the text of section 7.3.2 for details.

The generator of the A_n^* lattices is most conveniently expressed by the $(n+1) \times n$ matrix (Conway & Sloane 1988b)

$$\mathbf{G}_{A_n^*} = \begin{pmatrix} 1 & \cdots & 1 & \frac{-n}{n+1} \\ -1 & \cdots & 0 & \frac{1}{n+1} \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & -1 & \frac{1}{n+1} \\ 0 & \cdots & 0 & \frac{1}{n+1} \end{pmatrix}. \quad (7.34)$$

In particular, the 3×2 matrix

$$\mathbf{G}_{A_2^*} = \begin{pmatrix} 1 & -\frac{2}{3} \\ -1 & \frac{1}{3} \\ 0 & \frac{1}{3} \end{pmatrix} \equiv \mathbf{G}_0 \quad (7.35)$$

generates the A_2^* lattice over a 2-dimensional plane embedded in 3 dimensions, as shown in Figure 7.3a.

The first step is to reduce \mathbf{G}_0 to a 2×2 matrix \mathbf{G}_1 , thereby eliminating the extra dimension. Prix (2007b) suggests obtaining \mathbf{G}_1 by finding the components of the lattice vectors (the columns of \mathbf{G}_0) with respect to a basis which spans the 2-dimensional plane of the lattice. An alternative is to rotate the lattice plane to be perpendicular to one of the coordinate directions, e.g. the z axis in figures 7.3a–7.3b. The value of the z coordinate will then always be zero, and thus the z dimension may be discarded. The desired rotation is given by a QR decomposition (Golub & Loan 1983):

$$\mathbf{G}_0 = \mathbf{Q}\mathbf{R}, \quad (7.36)$$

where \mathbf{Q} is a 3×3 orthogonal matrix and \mathbf{R} is a 3×2 upper triangular matrix. If we take

$$\mathbf{G}'_1 = \mathbf{R} = \mathbf{Q}^T \mathbf{G}_0, \quad (7.37)$$

then \mathbf{G}'_1 is equivalent to \mathbf{G}_0 , as equation 7.37 satisfies equation 7.3. The generator

$$\mathbf{G}'_1 = \mathbf{Q}^T \mathbf{G}_0 = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & 1 \\ -\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & 1 \\ 0 & \sqrt{\frac{2}{3}} & 1 \end{pmatrix}^T \mathbf{G}_0 = \begin{pmatrix} \sqrt{2} & -\frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{6}} \\ 0 & 0 \end{pmatrix} \quad (7.38)$$

is plotted in Figure 7.3b. The plane of the lattice has no extent in the z direction, and thus the last row of \mathbf{G}'_1 is zero. We remove this row to obtain

$$\mathbf{G}_1 = \mathbf{J}\mathbf{G}'_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \sqrt{2} & -\frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{6}} \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} \sqrt{2} & -\frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{6}} \end{pmatrix}. \quad (7.39)$$

Next, we transform \mathbf{G}_1 to a lower triangular matrix

$$\mathbf{G}_2 = \mathbf{K}\mathbf{G}_1\mathbf{L} = \begin{pmatrix} \frac{1}{\sqrt{6}} & 0 \\ -\frac{1}{\sqrt{2}} & \sqrt{2} \end{pmatrix}, \quad (7.40)$$

where

$$\mathbf{K} = \begin{pmatrix} 0 & \text{sign}([\mathbf{G}_1]_{22}) \\ \text{sign}([\mathbf{G}_1]_{11}) & 0 \end{pmatrix}, \quad \mathbf{L} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (7.41)$$

are anti-diagonal matrices. The sign expressions in \mathbf{K} ensure that the diagonal elements of \mathbf{G}_2 are always positive (although this is not required for the example hexagonal lattice considered here). This is convenient for a practical implementation (see section 7.4) as it ensures that the parameters p_i are always incremented in a consistent direction: from the lower bound p_i^{\min} to the upper bound p_i^{\max} . We note that \mathbf{G}_2 is equivalent to \mathbf{G}_1 by equation 7.3.

Finally, we perform a coordinate transformation so that the covering spheres of \mathbf{G}_2 are mapped to the mismatch ellipses of the metric \mathbf{M} , as required for the template bank. Let $\vec{q} = \mathbf{G}_2\vec{k}$ be a lattice point generated by \mathbf{G}_2 . If we take any point \vec{q}' on the surface of the covering sphere surrounding \vec{q} , of radius R , then the square of the distance between \vec{q} and \vec{q}' is, by definition,

$$(\vec{q} - \vec{q}') \cdot (\vec{q} - \vec{q}') = R^2. \quad (7.42)$$

Now consider the points $\vec{p} = \mathbf{T}\vec{q}$ and $\vec{p}' = \mathbf{T}\vec{q}'$, where the transformation \mathbf{T} is defined by

$$\mathbf{T} = \frac{\sqrt{\mu_{\max}}}{R} \mathbf{D}, \quad (7.43)$$

and the matrix \mathbf{D} has the property that its columns are orthonormal with respect to the metric, i.e.

$$\mathbf{D}^T \mathbf{M} \mathbf{D} = \mathbf{1}. \quad (7.44)$$

The square of the *metric distance* between \vec{p} and \vec{p}' is

$$(\vec{p} - \vec{p}') \cdot \mathbf{M}(\vec{p} - \vec{p}') = \frac{\mu_{\max}}{R^2} \left[(\vec{q} - \vec{q}') \cdot \mathbf{D}^T \mathbf{M} \mathbf{D} (\vec{q} - \vec{q}') \right] = \mu_{\max}; \quad (7.45)$$

in other words, \vec{p}' is on the surface of the metric mismatch ellipse surrounding \vec{p} . Thus, the transformation \mathbf{T} maps the covering spheres of the lattice to the mismatch ellipses of the template bank, as required. The generator of the template bank points is

$$\mathbf{G}_3 = \mathbf{T}\mathbf{G}_2 \equiv \mathbf{G}. \quad (7.46)$$

We note that equation 7.46 does not need to satisfy equation 7.3; we take the view that \mathbf{T} is a transformation acting on the space containing the lattice (by changing the coordinate), and is not acting on the lattice itself (within some fixed space). Therefore, the lattice generated by \mathbf{G}_3 has the same properties of the lattice generated by \mathbf{G}_2 , in particular whether the coverage is optimal (see section 7.1.2).

For the example hexagonal lattice, and using the spindown-only metric (equation 7.17, with an arbitrary T_{span}),

$$\mathbf{T} = \frac{\sqrt{\mu_{\text{max}}}}{\sqrt{2}/3} \begin{pmatrix} \frac{8}{\sqrt{3}\pi} & 0 \\ -\frac{10}{\sqrt{3}\pi} & \frac{2\sqrt{5}}{3\pi} \end{pmatrix}; \quad (7.47)$$

the generator

$$\mathbf{G}_3 = \mathbf{T}\mathbf{G}_2 = \sqrt{\mu_{\text{max}}} \begin{pmatrix} \frac{4}{\pi} & 0 \\ -\frac{5+\sqrt{5}}{\pi} & \frac{2\sqrt{5}}{\pi} \end{pmatrix} \quad (7.48)$$

is plotted, for an arbitrary μ_{max} , in Figure 7.3d.

It remains to check that the transformation \mathbf{T} is lower triangular, so that it preserves the lower triangular form of \mathbf{G}_2 . The matrix \mathbf{D} may be found using Gram-Schmidt orthogonalisation (Golub & Loan 1983), where vectors \vec{u}_i are considered orthogonal if $\vec{u}_i \cdot \mathbf{M}\vec{u}_j = \delta_{ij}$, where δ_{ij} is the Kronecker delta. The order of operations can be arranged so that \mathbf{D} is lower triangular; thus the final generator \mathbf{G} is lower triangular, and the parameter space bounds on the integers \vec{k} are of the form given by equation 7.20.

In summary, the $n \times n$ generator matrix \mathbf{G} of the template bank for an n -dimensional parameter space, given the metric \mathbf{M} , and starting from any $m \times n$ lattice generator $\mathbf{G}_{\mathcal{L}}$, is given by

$$\mathbf{G} = \frac{\sqrt{\mu_{\text{max}}}}{R} \mathbf{D}\mathbf{K}\mathbf{J}\mathbf{Q}^T \mathbf{G}_{\mathcal{L}} \mathbf{L}, \quad (7.49a)$$

where

$$\mathbf{G}_{\mathcal{L}} = \mathbf{Q}\mathbf{R}, \quad (7.49b)$$

$$\mathbf{Q} = m \times m, \text{ orthogonal}, \quad (7.49c)$$

$$\mathbf{R} = m \times n, \text{ upper triangular}, \quad (7.49d)$$

$$\mathbf{J} = n \times m, [\mathbf{J}]_{ij} = \delta_{ij}, \quad (7.49e)$$

$$\mathbf{K} = n \times n, [\mathbf{K}]_{ij} = \delta_{(n-i+1)j} \text{sign}([\mathbf{R}]_{jj}) \quad (7.49f)$$

$$\mathbf{L} = n \times n, [\mathbf{L}]_{ij} = \delta_{(n-i+1)j} \quad (7.49g)$$

$$\mathbf{D} = n \times n, \text{ lower triangular}, \mathbf{D}^T \mathbf{M} \mathbf{D} = \mathbf{1}, \quad (7.49h)$$

$$R = \theta^{1/n} (\det \mathbf{G}_{\mathcal{L}}^T \mathbf{G}_{\mathcal{L}})^{1/(2n)}. \quad (7.49i)$$

7.3.3 Number of templates required for coverage

The number of templates the algorithm requires to cover the parameter space can be estimated by dividing the volume of the parameter space by the volume covered by a single template point (see Prix 2007b). This approach is accurate if the extent of the parameter space is always much larger than a single template, so the the parameter space edges are not important. If, however, the edges of the parameter space do become important, in particular if the parameter space becomes smaller than the extent of a single template, the calculation will assume that only fractions of a template are needed, and the sum of those fractions will underestimate the total number of whole templates actually required. Fortunately, the measures taken in section 7.3.1 to ensure proper coverage of the edges can also be used here to arrive at an accurate estimate of the number of templates.

The parameter space of the gravitational wave search for Cassiopeia A (Chapters 8 and 9) is a particular case where the edge effects of the parameter space are significant. The application of the template bank algorithm to this parameter space is discussed in section 8.4; in particular, we show that the estimate for the number of templates presented here accurately predicts the number of templates generated by the algorithm.

Given the parameter space of equation 7.19, with the additional coverage of the edges given by equation 7.23, the coordinate volume of the parameter space \mathcal{V} is found by integrating over the space, starting with p_n and ending with p_1 :

$$\mathcal{V} = \int_{p_1^{\min} - \delta p_1}^{p_1^{\max} + \delta p_1} dp_1 \int_{p_2^{\min}(p_1) - \delta p_2}^{p_2^{\max}(p_1) + \delta p_2} dp_2 \int_{p_3^{\min}(p_1, p_2) - \delta p_3}^{p_3^{\max}(p_1, p_2) + \delta p_3} dp_3 \cdots \int_{p_n^{\min}(p_1, \dots, p_{n-1}) - \delta p_n}^{p_n^{\max}(p_1, \dots, p_{n-1}) + \delta p_n} dp_n \quad (7.50)$$

The volume of a single template, $\sqrt{\det \mathbf{G}^T \mathbf{G}}$, is simply the volume of the fundamental region (equation 7.2) of the lattice generated by \mathbf{G} . It follows from equations 7.49 that the determinant of $\mathbf{G}^T \mathbf{G}$ is

$$\det \mathbf{G}^T \mathbf{G} = \left(\frac{\sqrt{\mu_{\max}}}{R} \right)^{2n} (\det \mathbf{D})^2 (\det \mathbf{G}_{\mathcal{L}}^T \mathbf{G}_{\mathcal{L}}). \quad (7.51)$$

Deducing the determinant of \mathbf{D} from equation 7.44, and substituting equation 7.6 for the normalised thickness of the lattice, this simplifies to

$$\det \mathbf{G}^T \mathbf{G} = \frac{\mu_{\max}^n}{\theta^2 \det \mathbf{M}}. \quad (7.52)$$

The estimate of the number of templates required for coverage is therefore

$$N = \frac{\mathcal{V}}{\sqrt{\det \mathbf{G}^T \mathbf{G}}} = \theta \mu_{\max}^{-n/2} \sqrt{\det \mathbf{M}} \mathcal{V}. \quad (7.53)$$

Aside from the inclusion of the parameter space edges in \mathcal{V} , this expression agrees with Prix (2007b).

7.4 Implementation

The template bank algorithm, presented above, was implemented in the C programming language, and integrated in the `ComputeFStatistic_v2` implementation of the \mathcal{F} statistic (see section 8.3). The implementation is named `FlatLatticeTiling`, and is available as part of LAL/LALApps (see bibliography), a collection of libraries and applications for gravitational wave data analysis.

The implementation is broken into four parts:

- `FlatLatticeTiling.c`, and associated header `FlatLatticeTiling.h`, are located in the LAL `tools` package. It contains the implementation of the template bank.
- `FlatLatticeTilingPulsar.c`, and associated header `FlatLatticeTilingPulsar.h`, are located in the LAL `pulsar` package. It provides a few functions specific to periodic gravitational wave searches.
- `testFlatLatticeTilingPulsar.c` is located in LALApps under `src/pulsar/templateBanks`. It is a stand-alone testing and debugging application.
- `DopplerFullScan.c`, and associated headers `DopplerFullScan.h` and `DopplerScan.h`, are located in the LAL `support` package. They were written by Reinhard Prix to connect a number of template bank implementations to `ComputeFStatistic_v2` through a common interface. Code was added to include the `FlatLatticeTiling` template bank.

7.4.1 Testing the implementation

In this section, we verify that the `FlatLatticeTiling` implementation of the template bank algorithm provides complete coverage over three example parameter spaces. We also confirm that the spindown-only metric provides a satisfactory, if conservative, estimator of the mismatch of the \mathcal{F} statistic (see also Prix 2007a). In section 8.4.1, we also show that the implementation generates the correct number of templates, as predicted by equation 7.53.

The example parameters spaces are subspaces of the parameter space of the Cassiopeia A search; this parameter space, as discussed in section 8.4, is dominated by edge effects from the second spindown (\ddot{f}) dimension. For this

reason, it is a suitable parameter space with which to test the algorithm's ability to account for the parameter space edges. We choose three frequency slices through the parameter space, with starting frequencies of 100, 200, and 300 Hz, and bandwidths of 10^{-4} Hz. The remaining parameters of the space are as described in section 8.4. The `FlatLatticeTiling` implementation generates $\sim 1.8 \times 10^6$, 3.7×10^6 , and 5.8×10^6 templates respectively to cover the three subspaces.

To test whether the template bank covers a particular slice, we perform the following procedure. We generate *injection points* by selecting random locations within the parameter space under test; approximately 10 injection points are generated for each template in the bank. For each injection point, we compute the distance, using the spindown-only metric (equation 7.17), from it to each template in the bank, and record the minimum distance, which we denote μ . This distance must always be less than the maximum mismatch μ_{\max} supplied to the algorithm. Figure 7.4a shows histograms of μ/μ_{\max} for the three subspaces. In each case, no injection points were found where $\mu > \mu_{\max}$, confirming that the template bank completely covered each subspace. The histograms are also identical, indicating the coverage does not change with the frequency-dependent size of the parameter spaces. The shape of the histograms reflects the geometry of the A_3^* lattice used to generate the template bank, and agree with the results of equivalent simulations by Messenger *et al.* (2009).

The above test assumed that the spindown-only metric \mathbf{M} is a reliable estimator of the mismatch behaviour of the \mathcal{F} statistic; we now confirm that this assumption is valid. We use two LALApps applications; `ComputeFStatistic_v2`, the implementation of the \mathcal{F} statistic, and `Makefake-data_v4`, which generates SFTs containing simulated periodic gravitational wave signals and, optionally, Gaussian noise.

We first use `Makefakedata_v4` to generate noise-free SFTs with a time span of $T_{\text{span}} = 12$ days. The SFTs contain a single simulated periodic signal with a fixed h_0 , and the remaining parameters chosen at random: values for $\cos \iota$, ψ , and ϕ_0 are chosen from their maximum possible ranges; a value for f is chosen from a range of 10^{-5} Hz, starting at 100 Hz; a value for \dot{f} is chosen from the range -1×10^{-8} to -5×10^{-9} Hz s $^{-1}$; and a value for \ddot{f} is chosen from the range -1×10^{-18} to -5×10^{-19} Hz s $^{-2}$. We then use `ComputeFStatistic_v2` to search for the signal at its known Doppler parameters $\lambda_{\text{signal}} = (f, \dot{f}, \ddot{f})$, thereby obtaining the maxima of the \mathcal{F} statistic, which we denote $2\mathcal{F}_{\max}$. Finally we use `ComputeFStatistic_v2` to perform a search for the signal, over the same ranges of f , \dot{f} , and \ddot{f} from which the random parameters of the signal were chosen. The search uses `FlatLatticeTiling`, with a maximum mismatch of $\mu_{\max} = 0.2$. We find the template with the

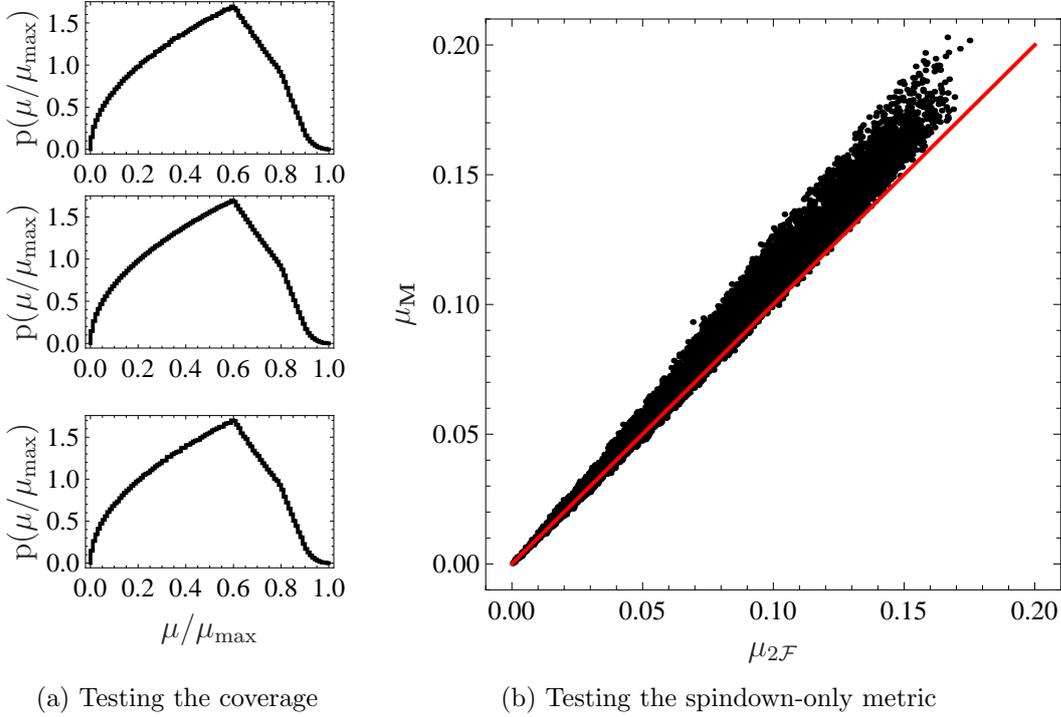


Figure 7.4: Testing the FlatLatticeTiling implementation of the template bank algorithm. (a) Testing the coverage: histograms of μ/μ_{\max} for the three subspaces of the Cassiopeia A search parameter space, with starting frequencies of (top to bottom) 100, 200, and 300 Hz. (b) Testing the spindown-only metric: the mismatch according to the metric, $\mu_{\mathbf{M}}$, versus the mismatch according to the \mathcal{F} statistic, $\mu_{2\mathcal{F}}$.

maximum value of $2\mathcal{F}$, which we denote $2\mathcal{F}_{\text{search}}$, and its Doppler parameters, which we denote $\vec{\lambda}_{\text{search}}$. From this we record the mismatch according to the \mathcal{F} statistic:

$$\mu_{2\mathcal{F}} = \frac{2\mathcal{F}_{\max} - 2\mathcal{F}_{\text{search}}}{2\mathcal{F}_{\max}}. \quad (7.54)$$

We also record, using the Doppler parameters of the signal and of its closest match, the mismatch according to the metric:

$$\mu_{\mathbf{M}} = (\vec{\lambda} - \vec{\lambda}_0) \cdot \mathbf{M}(\vec{\lambda} - \vec{\lambda}_0). \quad (7.55)$$

Figure 7.4b plots $\mu_{2\mathcal{F}}$, $\mu_{\mathbf{M}}$; each point represents one of the ~ 7400 repetitions of the above procedure. We expect $\mu_{2\mathcal{F}} \lesssim \mu_{\mathbf{M}}$, if the spindown metric is a conservative estimator of the \mathcal{F} statistic mismatch, and this is indeed the case. We note that $\sim 10\%$ of the points do fall below the red line

$\mu_{2\mathcal{F}} = \mu_{\mathbf{M}}$, but in these cases the deviation is small: the average value of $|\mu_{2\mathcal{F}} - \mu_{\mathbf{M}}| \lesssim 10^{-4}$. The slight curve of the black points away from the red line is due to the quadratic nature of the metric, which at high μ begins to significantly over-estimate the \mathcal{F} statistic mismatch.

7.5 Summary

This chapter presented an efficient template bank generation algorithm for coherent periodic gravitational wave searches. We introduced optimal lattices and sphere coverings, which are used to minimise the number of templates required to cover the template parameter space of the search. A metric on the parameter space is used to find the correct template spacing, such that any potential gravitational wave signal will only suffer a prescribed maximum loss in signal-to-noise ratio when matched to its nearest template. We presented each step in the template bank algorithm, showing how to efficiently iterate over the parameter space (and how to correctly account for its edges), how to construct the lattice generator used to position the templates, and how to estimate the number of templates required for complete coverage. Finally, we presented an implementation of the algorithm and validated its performance. This algorithm will be used in the search for periodic gravitational waves from Cassiopeia A presented in the next two chapters.

Chapter 8

A search for gravitational waves from Cassiopeia A

This chapter, and Chapter 9, present a search for gravitational waves targeted at the young Galactic supernova remnant Cassiopeia A (Cas A). This search is the first search by the LIGO Scientific Collaboration (LSC) to target a non-pulsing isolated neutron star. We discuss the motivation for a search targeting Cas A in section 8.1. We then derive an indirect upper limit on the strength of gravitational waves from Cas A, based on energy conservation, in section 8.2. Sections 8.3 and 8.4 present the proposed search, including the choice of analysis method, the time span of the data set, and the parameter space of frequencies and spindowns to search. Section 8.5 calculates the sensitivity of the search, and section 8.6 calculates its computational cost and the upper limits it is expected to attain. Finally, section 8.7 computes the expected distribution of the largest value of the detection statistic that will be found by the search, assuming that there is no gravitational wave signal.

See section 1.1 for author contributions and publications, a disclaimer, and acknowledgements relevant to this chapter.

8.1 The central compact object

Cas A is a core-collapse supernova remnant (Krause *et al.* 2008). It is $3.4_{-0.1}^{+0.3}$ kpc distant (Reed *et al.* 1995), and is estimated to have occurred in the year 1681 ± 19 (Fesen *et al.* 2006a). The supernova was possibly observed by the first Astronomer Royal, John Flamsteed (Hughes 1980). It is the second youngest supernova remnant in the Galaxy, after G1.9+0.3 (Reynolds *et al.* 2008), and is one of the strongest radio sources in the sky

(Baade & Minkowski 1954, McLaughlin *et al.* 2001). At the centre of the remnant is a bright x-ray point source, first identified in first-light images taken by the Chandra x-ray satellite (Tananbaum 1999), and retro-actively identified in archival data from the Röntgen (ROSAT) and Einstein satellites (Pavlov *et al.* 2000). The point source, designated CXOU J232327.8+584842, is almost certainly associated with the remnant (Pavlov *et al.* 2000, Kaplan *et al.* 2001). Its right ascension $\alpha = 23\text{h } 23\text{m } (27.943 \pm 0.05)\text{s}$ and its declination $\delta = 58^\circ 48' (42.51 \pm 0.4)''$ (Fesen *et al.* 2006b). The point source is included as a member of the class of central compact objects (CCOs), of which it is the youngest (De Luca 2008).

The nature of the Cas A CCO remains uncertain, although a number of possibilities have been ruled out. A black hole is excluded (Pavlov & Luna 2009), due to the lack of observed x-ray variability which would arise from accretion (Pavlov *et al.* 2000). This leaves an accreting or isolated neutron star as the most likely possibilities (e.g. Murray *et al.* 2002, Pavlov & Luna 2009). The transverse velocity of the CCO, assuming an initial position coincident with the centre of expansion, is $\sim 350 \text{ km s}^{-1}$ (Fesen *et al.* 2006a). Accretion by the CCO, due to its motion, of remnant SNR material or the inter-stellar medium is, however, inconsistent with the observed x-ray flux (Pavlov *et al.* 2000, Fesen *et al.* 2006b). Optical and near-infrared observations by Ryan *et al.* (2001), Kaplan *et al.* (2001), Fesen *et al.* (2006b), and Wang *et al.* (2007) found no counterpart to the CCO. Upper limits in these bands constrain the absolute magnitude of any binary companion to less than any main-sequence star or known x-ray binary companion (Fesen *et al.* 2006b). A residual accretion disk, formed from fallback material from the SNR, was not detected by Wang *et al.* (2007), but a cool passive disk is still a possibility.

McLaughlin *et al.* (2001) searched for radio pulsations from the CCO with frequencies between $\sim 0.1 \text{ Hz}$ and $\sim 2 \text{ kHz}$ (depending on the dispersion measure), and found no credible signal. This is evidence against the CCO being an active pulsar, although it may simply be beamed away from Earth. The upper limits on radio pulsations derived by McLaughlin *et al.* are lower than for any known pulsar younger than 10^4 yr . The absence of a pulsar wind nebula (Hwang *et al.* 2004) is also evidence against an active pulsar. X-ray pulsations from the CCO have been searched for by Chakrabarty *et al.*, Murray *et al.* (2002), Mereghetti *et al.* (2002), and most recently by Pavlov & Luna (2009); together these searches have covered frequencies between $30 \mu\text{Hz}$ and 500 Hz . No pulsations were observed¹; a weakly-significant

¹ According to Pavlov & Luna (2009): "A 300 ks observation with the Chandra HRC-S detector, aimed at the search for the CCO period, has been carried out recently (2009

candidate at ~ 83 Hz (~ 12 ms), reported by Murray *et al.*, was not confirmed by Ransom (2002).

A power-law fit to the CCO's x-ray spectrum is considerably steeper than expected for a classical young pulsar; a blackbody fit yields a temperature too high, and a radius too low, to be consistent with x-ray emission from the whole surface of a cooling neutron star with a uniform temperature (Pavlov *et al.* 2000, Chakrabarty *et al.* 2001). Better fits are provided by more sophisticated neutron star atmosphere models (Pavlov & Luna 2009), but are still inconsistent with the parameters of a standard neutron star. Pavlov *et al.* and Chakrabarty *et al.* suggested that the CCO might possess a very strong magnetic field ($\sim 10^{14}$ – 10^{15} G), which would produce a non-uniform temperature distribution. Pavlov *et al.* proposed a two-component spectral model of a neutron star with hot polar caps; a problem is that the model predicts (as-yet unobserved) x-ray pulsations at the rotation frequency.

Two possible models for the Cas A CCO are magnetars and anti-magnetars. Magnetars possess very strong magnetic fields and rotate slowly at ~ 0.1 – 0.5 Hz, while anti-magnetars are weakly magnetised ($\lesssim 10^{11}$ G) and rotate relatively quickly at ~ 10 Hz (Pavlov & Luna 2009). To date, there is no observational evidence which conclusively favours either model (Pavlov *et al.* 2000, Gotthelf & Halpern 2008, Pavlov & Luna 2009). Magnetars include the anomalous x-ray pulsars (AXPs) and soft γ -ray repeaters (SGRs), both of which emit occasional energetic outbursts (Fesen *et al.* 2006b). Anti-magnetars are those CCOs which do not show any pulsar- or magnetar-like activity (Gotthelf & Halpern 2008, Pavlov & Luna 2009). Using infrared echoes from the inter-stellar medium, Krause *et al.* (2005) claimed to have detected a ~ 60 -year-old high-energy flare from the CCO, similar to those seen from SGRs; the claim was, however, not supported by subsequent observations (Kim *et al.* 2008, Dwek & Arendt 2008). The lack of any observed bursts or flares from the Cas A CCO is evidence against the magnetar model; although the CCO may be at an early, weakly-magnetised stage of development, when outbursts does not occur, this is a difficult scenario to distinguish observationally (Pavlov & Luna 2009). The Cas A CCO's low x-ray luminosity is also inconsistent with the spindown luminosity expected if the CCO were a young magnetar (Gotthelf & Halpern 2008, Pavlov & Luna 2009).

8.1.1 Motivation for a gravitational wave search

From the perspective of detecting gravitational waves from the CCO, an anti-magnetar model is more promising, as it implies a more rapid rotation rate.

March; PI: D. Chakrabarty), but no results have been published yet.”

Three of the known CCOs have confirmed spin frequencies (2.3–9.5 Hz); a fourth measured period is possibly due to binary orbital motion (De Luca 2008, Gotthelf & Halpern 2009). The upper limits on their spindowns are low, implying near-constant rotation rates over the stars’ lifetimes; they also suggest weak magnetic fields consistent with the magnetar model (Pavlov & Luna 2009). Non-axisymmetric gravitational radiation at twice the known CCO spin frequencies, i.e. $\lesssim 20$ Hz, would be undetectable by LIGO, due to the steep seismic noise floor below ~ 50 Hz. Nevertheless, given the lack of a definite model for the CCOs, it is premature to suggest that the small sample of known CCOs are representative of the class in general. For example, the Cas A CCO is much younger than the other known CCOs (De Luca 2008), and therefore it may be true that Cas A exhibits different physics, such as r -modes melting the crust (Owen *et al.* 1998, Lindblom *et al.* 2000, Wu *et al.* 2001).

In short, the wealth of multi-wavelength observations and theoretical modelling of the Cas A CCO suggests the possibility that it might be a gravitational wave source. If so, it may indeed be detectable even before advanced interferometers come online, as we show in the next section. Such a detection would provide the rotational frequency and spindown of the CCO, resolving its ambiguous classification (Pavlov *et al.* 2000, Chakrabarty *et al.* 2001, Gotthelf & Halpern 2008, Pavlov & Luna 2009). There is, therefore, a wealth of information which may be potentially gained by the gravitational wave search presented below.

8.2 Indirect upper limits

It is possible to determine an *a priori*, indirect upper limit on periodic gravitational waves from Cas A² (Wette *et al.* 2008). We first assume that the gravitational wave luminosity (i.e. the energy radiated in gravitational waves per unit time) is bounded by the loss of rotational kinetic energy:

$$L_{\text{gw}} = \frac{32G}{5c^5} I_{zz}^2 \epsilon^2 (\pi f)^6 \lesssim -\frac{d}{dt} \left(\frac{1}{2} \pi^2 I_{zz} f^2 \right) = -\dot{E}_{\text{rot}}, \quad (8.1)$$

where L_{gw} is the gravitational wave luminosity (e.g. Zimmermann & Szednits 1979), E_{rot} is the rotational kinetic energy, f is the gravitational wave frequency (which is, assuming non-axisymmetric emission, twice the rotation frequency), I_{zz} is the principal moment of inertia, and ϵ is the equatorial el-

² Hereafter, for brevity, we mean “Cas A” to refer to “the CCO in Cas A”.

lipticity. Equation 8.1 can be rearranged to give an upper limit on

$$\epsilon \lesssim \sqrt{\frac{5c^5}{32\pi^4 GI_{zz}} \frac{-\dot{f}}{f^5}}, \quad (8.2)$$

and on the strain amplitude

$$h_0 \lesssim \frac{1}{D} \sqrt{\frac{5GI_{zz}}{2c^3} \frac{-\dot{f}}{f}} \quad (8.3)$$

by substituting equation 2.4; here, D is the distance to the source. Equations 8.2 and 8.3 are known as *spindown upper limits*. They are used to derive upper limits on known pulsars (e.g. Abbott *et al.* 2010), where the frequency and spindown are known from electromagnetic observations of the pulsar.

For Cas A, the frequency and spindown are unknown, so the spindown upper limits cannot be used directly. If we assume, however, that the CCO was born spinning at a substantially higher rate than at present, and that its spindown since then has been primarily due to quadrupolar gravitational wave emission, we can relate the unknown frequency and spindown to the known age τ of the CCO as follows (Ostriker & Gunn 1969, Palomba 2000, Abbott *et al.* 2007b):

$$\tau \approx \frac{1}{n-1} \left(\frac{f}{-\dot{f}} \right). \quad (8.4)$$

The quantity

$$n = f\ddot{f}/\dot{f}^2 \quad (8.5)$$

is known as the braking index. A neutron star spins down according to $\dot{f} \propto f^n$, where the constant of proportionality and n depend on the mechanism by which energy is lost; for example: electromagnetic (dipolar) radiation ($n = 3$), gravitational (quadrupolar) radiation ($n = 5$), or a saturated r -mode ($n = 7$). Substituting equation 8.4 into equations 8.2–8.3 gives the *age-based upper limits*

$$\epsilon \lesssim \sqrt{\frac{5c^5}{32\pi^4 GI_{zz} f^4 (n-1)\tau}}, \quad (8.6)$$

$$h_0 \lesssim \frac{1}{D} \sqrt{\frac{5GI_{zz}}{2c^3 (n-1)\tau}}. \quad (8.7)$$

While the age-based upper limits are similar to the spindown upper limits, they rely on less information, and are therefore less solid. Like the spindown

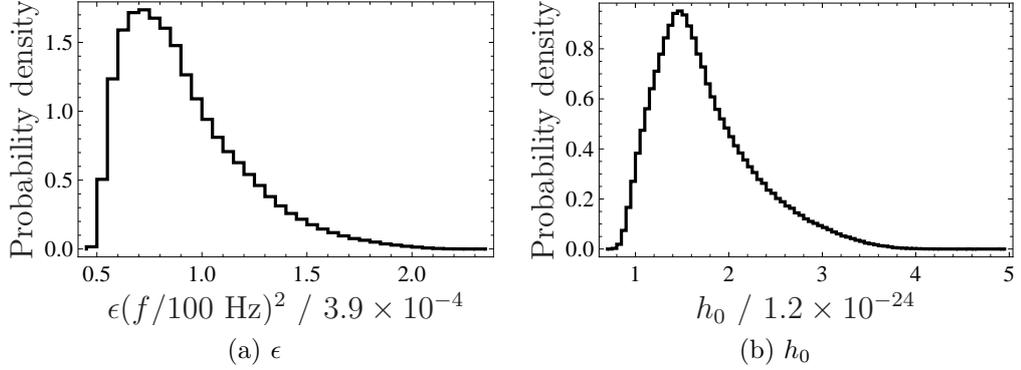


Figure 8.1: Histograms of the probability distributions of the upper limits on (a) the ellipticity ϵ , assuming $f = 100$ Hz, and (b) the strain h_0 . The ϵ and h_0 are plotted in units of their respective fiducial values for Cas A.

upper limits, they are most useful as a relative indicator of which astronomical objects would be interesting to target with a gravitational wave search, as opposed to a rigorous estimator of the *expected* gravitational wave strength.

Using $n = 5$, the known age and distance, and a fiducial value for I_{zz} , we arrive at the following age-based upper limits on Cas A:

$$\epsilon \lesssim 3.9 \times 10^{-4} \left(\frac{100 \text{ Hz}}{f} \right)^2 \sqrt{\frac{4}{n-1} \frac{300 \text{ yr}}{\tau} \frac{10^{45} \text{ g cm}^2}{I_{zz}}}, \quad (8.8)$$

$$h_0 \lesssim 1.2 \times 10^{-24} \left(\frac{3.4 \text{ kpc}}{D} \right) \sqrt{\frac{4}{n-1} \frac{300 \text{ yr}}{\tau} \frac{I_{zz}}{10^{45} \text{ g cm}^2}}. \quad (8.9)$$

Note that the fiducial upper limit on h_0 is independent of frequency. We adopt a fiducial value for the age τ at the lower end of the estimated range (Fesen *et al.* 2006a); this choice gives more conservative bounds on the search parameter space (section 8.4), as well as a higher upper limit. The upper limit on ϵ is at the speculative end of the range of theoretical predictions (see section 2.4).

8.2.1 Uncertainty in the upper limits

There are a number of sources of uncertainty in the age-based upper limits on Cas A (equations 8.8 and 8.9), aside from the validity of the assumptions used in their derivation. In addition to the measured uncertainties in the age (Fesen *et al.* 2006a) and distance (Reed *et al.* 1995), the principal moment of inertia could conceivably be three times the fiducial value assumed here

(see discussion in Abbott *et al.* 2007e). Since the CCO is observed in x-rays, the true braking index will at least include electromagnetic radiation, and possibly other mechanisms. Furthermore, the relative weighting of the mechanisms spinning down the star, and thus the braking index, are likely to change over its lifetime. Therefore, instead of the assumed $n = 4$, the appropriate quantity for the upper limits is the braking index averaged over the lifetime of the star, which is unknown.

Here, we present a rough estimate of the effect of the above uncertainties in τ , D , I_{zz} , and n on the value of the indirect upper limit; we do not attempt a more sophisticated treatment such as, for example, modelling different spindown histories of the CCO. We evaluate equations 8.6–8.7 using random values for the parameters τ , D , I_{zz} , and n drawn from the following distributions: for τ , a Gaussian distribution with a mean of 328 yr and a standard deviation of 19 yr; for D , a Gaussian distribution with a mean of 3.4 kpc and a standard deviation of 0.2 kpc (i.e. the average of the upper and lower errors in Reed *et al.* 1995); for I_{zz} , a uniform distribution over the range $1\text{--}3 \times 10^{45}$ g cm²; and for n , a uniform distribution over the range 2–7 (see section 8.4).

Figure 8.1 shows histograms, computed using 10^6 trials, of the probability distributions on the upper limits on ϵ (assuming $f = 100$ Hz) and h_0 . Due to the uncertainties in τ , D , I_{zz} , and n , the upper limit on ϵ is $\sim 70\%$ likely to be less than its fiducial value of 3.9×10^{-4} . On the other hand, the upper limit on h_0 is only $\sim 3\%$ likely to be less than its fiducial value of 1.2×10^{-24} . The fiducial upper limit on Cas A is comparable to the spindown upper limit on the Crab pulsar of $h_0 = 1.4 \times 10^{-24}$, which has been beaten by a search of LIGO S5 data (Abbott *et al.* 2008c).

8.3 Choice of analysis method and time span of data set

Both coherent (e.g. Abbott *et al.* 2007b) and semi-coherent (e.g. Abbott *et al.* 2008a) analysis methods have been used to search for periodic gravitational waves. Semi-coherent methods are in practice more sensitive than coherent methods for many searches of large data sets, when computational limits are taken into account (see section 5.2.3). For Cas A, however, the integration time required to beat the indirect limit is short enough that a simpler and more sensitive coherent method can be used for reasonable computational cost (Wette *et al.* 2008). In section 8.6, we show that a \mathcal{F} statistic search of data, at the design sensitivity of LIGO, can beat the indirect limit if the

span of the data set is

$$T_{\text{span}} \approx 12 \text{ days}, \quad (8.10)$$

and that the computational cost of such a search is feasible.

The implementation of the \mathcal{F} statistic used by the search is named `ComputeFStatistic_v2`, and is available as part of the LALApps (see bibliography) software package.³ The implementation (Prix 2008) takes SFTs as input, and computes the multi-detector \mathcal{F} statistic of Cutler & Schutz (2005) in the frequency domain. `ComputeFStatistic_v2` has previously been used in a wide-parameter search for the Crab pulsar (Abbott *et al.* 2008c), and is also being used in Einstein@Home searches. It is the successor to the single-detector \mathcal{F} statistic implementation used in Abbott *et al.* (2007b, 2009d).

8.4 The frequency and spindown parameter space

The parameters of a periodic gravitational wave search include the sky position and frequency evolution of the source (section 5.2.2). The sky position of Cas A is known to sufficiently high precision that it need not be searched over. On the other hand, because Cas A is younger than objects considered in previous LSC multi-template searches, a second spindown derivative is required (Wette *et al.* 2008). Over the chosen length of the Cas A data set (12 days), the maximum shift in frequency, over the range of second spindowns presented below (equation 8.11c) is $\sim 1.3 \times 10^{-5}$ Hz, which is larger than the frequency extent of the mismatch ellipse (equation 7.33) of $\sim 4.8 \times 10^{-6}$. Thus, if the second spindown was omitted, a gravitational wave signal from Cas A could potentially be lost due to decoherence between signal and template. For this reason, previous and contemporary all-sky searches for periodic gravitational waves using PowerFlux (Abbott *et al.* 2009b) and Einstein@Home (Abbott *et al.* 2009d) were not sensitive to the complete range of frequency evolution parameters possible for Cas A. Indeed, Cas A is young enough that the range of first spindowns of these searches may not cover all possibilities; the parameter space of (Abbott *et al.* 2009d) is restricted to search for objects with spindown ages greater than 8000 years.

The parameter space of the Cas A search thus covers frequency f , first spindown \dot{f} , and second spindown \ddot{f} . A search of this type of parameter space is often referred to as a *directed* search. Previous directed searches by the LSC have targeted Scorpius X-1 (see section 5.1.4); the search for Cas A is the first to target a non-pulsing neutron star. Figure 8.2 visualises the

³ The version of the code used here is tagged with the identifier “S5CasASearch”.

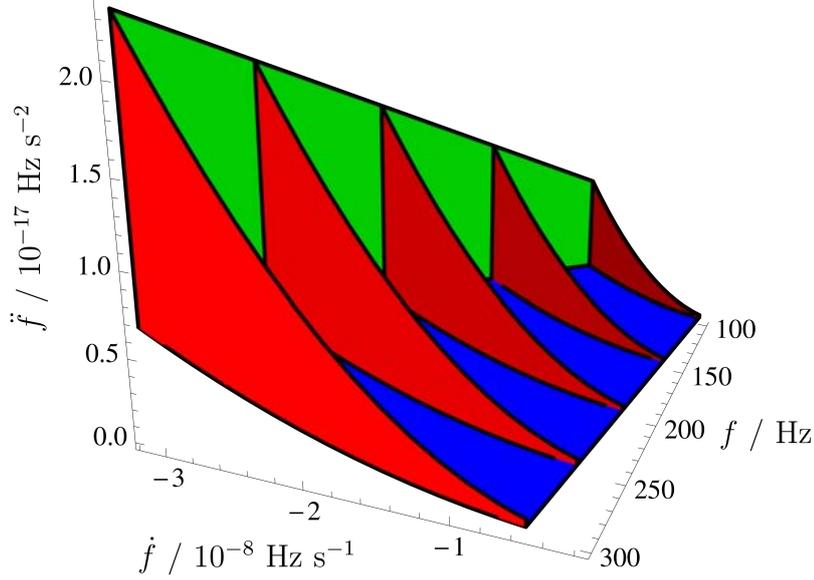


Figure 8.2: Visualisation of the Cas A search parameter space, defined by equations 8.11. The red surfaces show the bounds on \dot{f} and \ddot{f} , at constant f from (dark to light) 100 to 300 Hz, in 50 Hz intervals. The green surface shows the lower bound on \dot{f} as a function of f , while the blue surface shows the lower bound of \ddot{f} as a function of f and \dot{f} . (The upper bounds on \dot{f} and \ddot{f} are not shown.)

bounds on the parameter space, which are chosen as follows (Wette *et al.* 2008). The frequency band is chosen to be

$$100 \text{ Hz} \leq f \leq 300 \text{ Hz} \quad (8.11a)$$

which surrounds the frequencies where the LIGO interferometers are most sensitive. It roughly corresponds to the band over which the search is expected to beat the indirect upper limit. The bounds on the spindowns are chosen to be

$$\frac{-f}{(\min n - 1)\tau} \leq \dot{f} \leq \frac{-f}{(\max n - 1)\tau}, \quad (8.11b)$$

$$\frac{(\min n)\dot{f}^2}{f} \leq \ddot{f} \leq \frac{(\max n)\dot{f}^2}{f}, \quad (8.11c)$$

where τ is the age of Cas A. The bounds on the first spindown come from rearranging equation 8.4, while the bounds on the second spindown come from

re-arranging equation 8.5; in both cases the lower and upper bounds result from assuming a range of possible braking indices. The choice of $\tau = 300$ yr (section 8.2) results in conservative bounds on the first spindown.

The range of braking indices is chosen to be $2 \leq n \leq 7$. This range covers all known pulsars except the Vela pulsar, which is visibly interacting with its wind nebula, and thus does not follow the spindown-only evolution assumed here. We note that the braking index in equation 8.11b is a lifetime average (as in the upper limits), while in equation 8.11c it is its present-day value. We therefore allow the lifetime-average and present-day braking indices to take independent values, which are then searched for over the ranges of \dot{f} and \ddot{f} respectively. Finally, we note that the bounds fit the form of equation 7.19, as required by the template bank algorithm of Chapter 7.

8.4.1 Application of the template bank algorithm

The search for Cas A uses the template bank algorithm of Chapter 7 to cover the parameter space of equations 8.11. The three-dimensional A_3^* (also known as body-centred cubic) lattice is used; the maximum mismatch $\mu_{\max} = 20\%$. Figures 8.3 and 8.4 illustrate the tiling produced by the algorithm. They show cross-sections through the tiling over a subspace of the Cas A parameter space, where (for the sake of illustration) f and \dot{f} take restricted ranges, but \ddot{f} takes its full range.

It is important to emphasise that the ellipses plotted in these figures are not the (3-dimensional) mismatch ellipses of the tiling, but instead are the intersection of the 3-dimensional ellipses with the cross-section plane, which yield 2-dimensional ellipses. This is the reason why, for example, the black ellipses in Figure 8.3, which originate from template points where f is exactly in the cross-section plane, need not completely cover the space. The intervening gaps are covered by the grey ellipses, which originate from template points outside of the cross-section plane, but whose 3-dimensional ellipses nevertheless intersect it. Conversely, the reason why some of the black ellipses in Figure 8.3 do not intersect the tiled region is because the full three-dimensional ellipse will intersect the parameter space in some other $f = \text{constant}$ cross-section. These extra templates are needed to completely cover the edges of the parameter space (see section 7.3.1), particularly in the \ddot{f} dimension, where the extent of the parameter space is clearly much smaller than the size of the metric ellipses. Indeed, the parameter space is so much smaller than the metric ellipses in \ddot{f} that it seems possible that this dimension could be flattened to a single point; we delay further discussion until Chapter 10.

The properties of the tiling in Figure 8.3 are similar to the tiling in Fig-

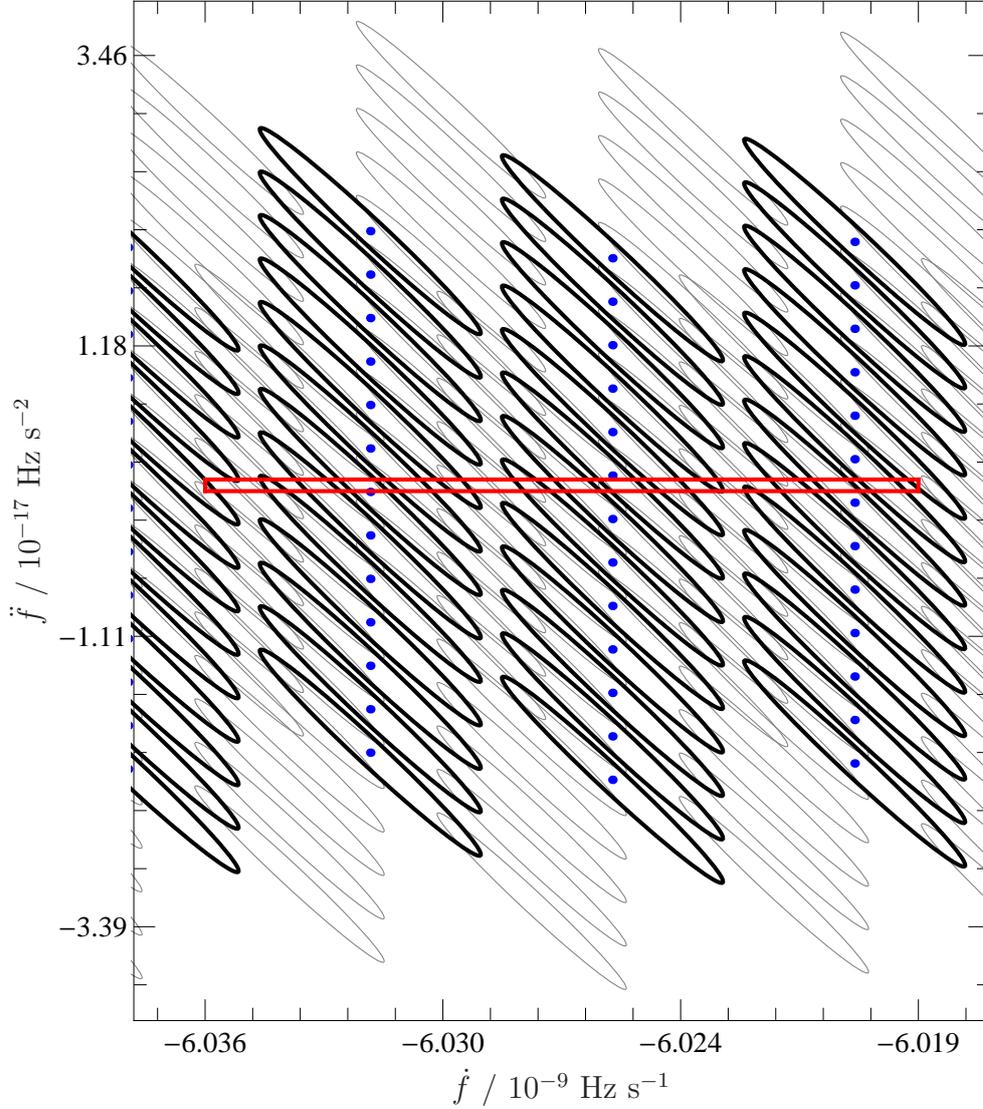


Figure 8.3: \dot{f} - \ddot{f} cross-section through the template tiling of the Cas A parameter space, at $f = 200$ Hz. The plotted ellipses are the intersections of the 3-dimensional metric ellipse with the cross-section plane. The range of \dot{f} is restricted to the braking indices $4.50 \leq n \leq 4.51$. The red box indicates the tiled range of \dot{f} and \ddot{f} . Intersection ellipses from templates where $f = 200$ Hz are plotted in black, with the template points plotted in blue; all other intersection ellipses are plotted in grey.

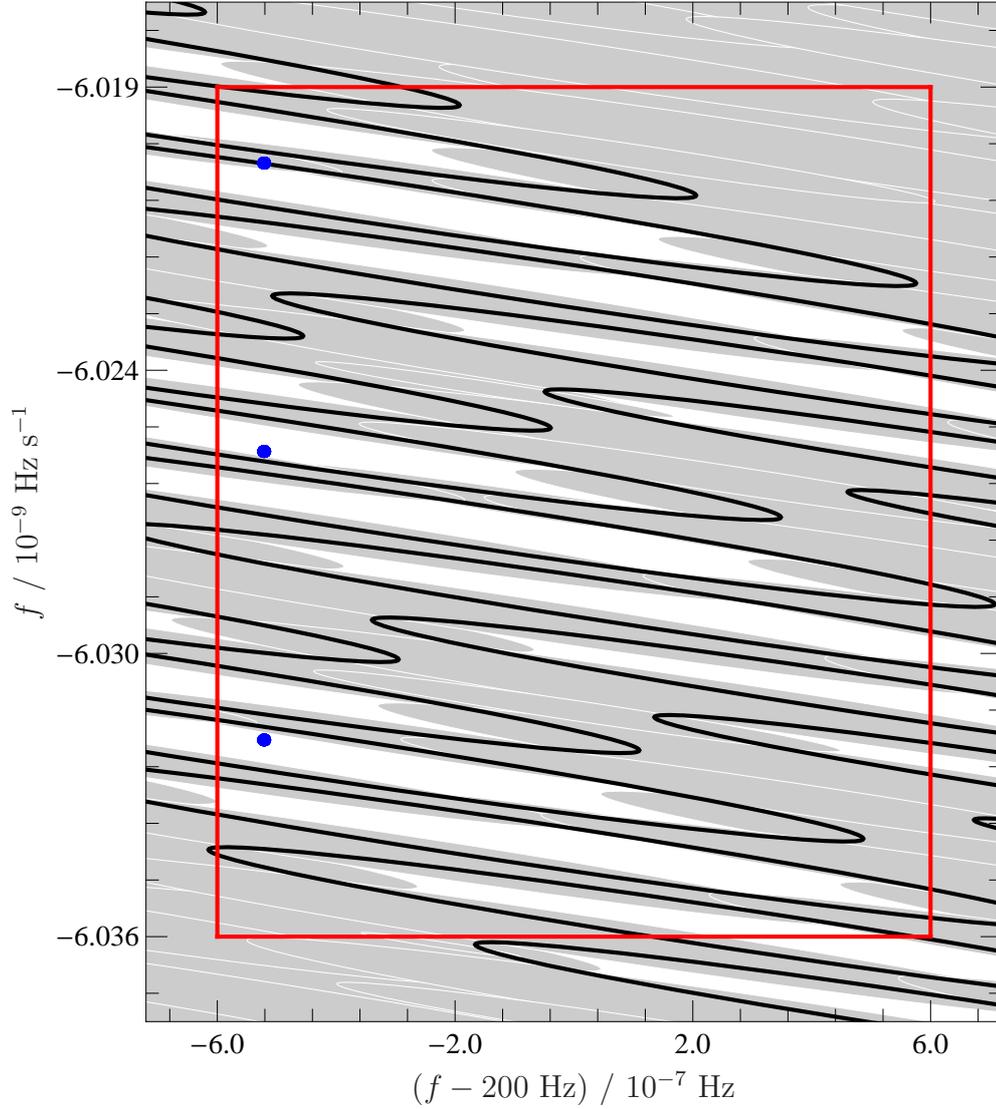


Figure 8.4: f - \dot{f} cross-section through the template tiling of the Cas A parameter space, at $\ddot{f} = 8.38 \times 10^{-19} \text{ Hz s}^{-2}$, the mid-point of the \ddot{f} parameter space range. As in Figure 8.3, the plotted ellipses are the intersections of the 3-dimensional metric ellipse with the cross-section plane. The range of \dot{f} is restricted to the braking indices $4.50 \leq n \leq 4.51$, and the range of f is restricted to $(200 \pm 6 \times 10^{-7}) \text{ Hz}$. The red box indicates the tiled range of f and \dot{f} . Intersection ellipses from templates where (f, \dot{f}) is inside the red box are plotted in black, with the template points plotted in blue; all other intersection ellipses are shaded in grey with a white border. The three blue template points in this figure correspond to the three lines of template points in Figure 8.3.

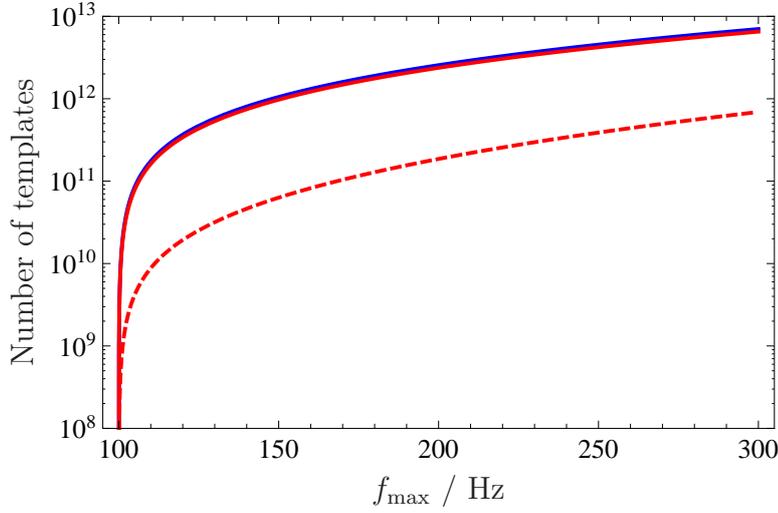


Figure 8.5: Cumulative number of templates required to cover the Cas A search parameter space, as a function of the maximum frequency searched f_{\max} . The number of templates used in the search (blue) is plotted against the estimate of equation 7.53 (red), including the parameter space edges (solid), and excluding the parameter space edges (dashed). The blue line is partially obscured by the solid red line.

ure 8.4; the f - \dot{f} region is covered either by templates laid within the region (black ellipses) or templates laid outside the region by whose 3-dimensional mismatch ellipses nevertheless intersect in a different $\ddot{f} = \text{constant}$ plane.

We use the equations in section 7.3.3 to estimate the number of templates required for the search. We use the normalised thickness of an A_3^* or body-centred cubic lattice (Table 7.1), and the spindown-only metric of the \mathcal{F} statistic (equation 7.17). Importantly, we also take account of the number of templates required to cover the edges of the parameter space. Figure 8.5 shows the number of templates required to cover the parameter space, plotted cumulatively as a function of the maximum frequency searched. The number of templates at $f_{\max} = 300$ Hz is the number required to cover the Cas A parameter space. The number of templates generated by the template bank algorithm is accurately estimated when the edges of the parameter space are taken into account, whereas ignoring the parameter spaces edges results in an estimate an order of magnitude too small, as calculated by Wette *et al.* (2008).

This indicates that the majority of the templates are used to cover the parameter edges. We deduce from Figure 8.3 that the small extent of the parameter space in the \ddot{f} direction is responsible for this. To confirm, we

expand equation 7.53, which gives the total number of templates in terms of the edge widths ∂f , $\partial \dot{f}$, and $\partial \ddot{f}$, which are given by extent of the bounding box of the mismatch ellipses (see section 7.3.1). The largest 3 terms read:

$$N \approx \underbrace{5.823 \times 10^{12}}_{\partial \ddot{f}} + \underbrace{6.91 \times 10^{11}}_1 + \underbrace{7 \times 10^9}_{\partial \dot{f} \partial \ddot{f}} + \dots, \quad (8.12)$$

where the under-braces indicate the origin of each term. The largest term arises from $\partial \ddot{f}$, and is thus the number of templates required to covering the \ddot{f} boundary, i.e. it would be zero if the \ddot{f} boundary were ignored. It is an order of magnitude larger than the next term, which is a constant, and is therefore the number of templates required to cover the bulk of the parameter space. This agrees with the order-of-magnitude difference between including and neglecting the boundary seen in Figure 8.5.

8.5 Sensitivity of the search

The sensitivity of a periodic gravitational wave search is an estimate of the upper limits that can be set by the search, in the absence of a detection. Upper limits are always given with respect to a certain confidence; for periodic gravitational wave searches, a 95% confidence is typically chosen. The meaning of an 95% confidence upper limit, denoted $h_0^{95\%}$, is that if a large population of signals with amplitude $h_0^{95\%}$ were injected into the search pipeline, we would expect 95% of them to be detected. The sensitivity is also set by the false alarm rate (see below).

For a coherent search, the estimated upper limit at the level of confidence C is given by

$$h_0^C = \Theta \sqrt{\frac{S_n}{T_{\text{obs}}}}, \quad (8.13)$$

where S_n is the one-sided power spectral densities of the noise in each detector, and T_{obs} is the amount of data available from each detector. For a multiple-detector search, the ratio S_n/T_{obs} is harmonically summed over the detector network. Strictly speaking, the ratio S_n/T_{obs} is also multiplied by a factor accounting for the orientation of each detector; in practice, however, this can usually be safely neglected (see Fig. 4 of Jaranowski *et al.* 1998). The sensitivity factor, denoted Θ , is determined by the following two equations:

$$p_\alpha = 1 - \left[\text{cdf}(\chi_4^2; 2\mathcal{F}_\alpha) \right]^N, \quad (8.14)$$

$$1 - C = p_\beta = \left\langle \text{cdf}(\chi_4^2(\Theta^2 R^2); 2\mathcal{F}_\alpha) \right\rangle, \quad (8.15)$$

where $p_\alpha = 1\%$ is the false alarm rate, and $p_\beta = 5\%$ is the false dismissal rate (see section 5.2.2), which is the complement of the confidence. Other parameters are explained below.

A central χ^2 distribution with 4 degrees of freedom is denoted χ_4^2 ; a non-central χ^2 distribution with 4 degrees of freedom and non-centrality parameter ρ^2 is denoted $\chi_4^2(\rho^2)$, where cdf denotes the cumulative distribution function (see section 5.2.2). Using equation 5.51, the optimal signal-to-noise ratio is

$$\rho^2 \approx \left(h_0^C \sqrt{\frac{T_{\text{obs}}}{S_n}} \right)^2 \left[\sum_{n=-2}^2 \left(\frac{\bar{h}[1, n]}{h_0^C} \right)^2 + \left(\frac{\bar{h}[-1, n]}{h_0^C} \right)^2 \right] \quad (8.16)$$

$$= \Theta^2 R^2; \quad (8.17)$$

the function R is independent of h_0^C , T_{obs} , and S_n . For multiple detectors, whose noise densities S_n can be regarded as equal, R^2 is averaged over the detector network.

In equation 8.14, N denotes the number of searched templates from which the upper limit is derived. Upper limits are typically set by partitioning the parameter space into small ($\lesssim 1$ Hz) frequency bands, over which the noise floor of the detectors can be regarded as constant; a single h_0^C is then calculated for each band. Therefore, N is the number of templates⁴ in a particular upper limit band, and therefore may be a function of frequency, depending on the parameter space.

Finally, the brackets $\langle \rangle$ in equation 8.15 imply marginalisation over the unknown parameters in R^2 : the right ascension α , declination δ , inclination angle ι , and polarisation angle ψ . For an object with a known sky position (e.g. known pulsars, and Cas A), the marginalisation then takes the form

$$\langle f \rangle = \left[\frac{1}{2} \int_{-1}^1 d(\cos \iota) \right] \left[\frac{2}{\pi} \int_{-\pi/4}^{\pi/4} d\psi \right] f(\cos \iota, \psi). \quad (8.18)$$

The average sensitivity factor of a single-template search for a known pulsar is $\Theta \approx 11.4$ (e.g. Abbott *et al.* 2007b, Prix 2009). In contrast to the method presented here, this factor is derived by taking a fixed non-centrality parameter to be the average; Monte Carlo simulations were used to confirm that this factor is approximately correct. The sensitivity factor is, as we shall see, significantly larger for multi-templates searches.

To determine the sensitivity, we solve equation 8.14 to obtain $2\mathcal{F}_\alpha$, substitute into equation 8.15, and solve for Θ . We set $p_\alpha = 1\%$, $p_\beta = 5\%$, and

⁴ A small caveat: we have assumed that the templates are statistically independent. See the extensive discussion in section 8.7.1.

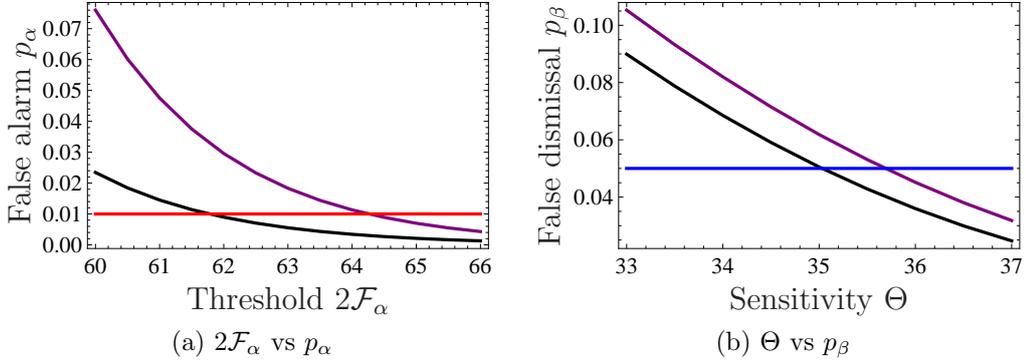


Figure 8.6: (a) False alarm rate p_α versus detection threshold $2\mathcal{F}_\alpha$. The target $p_\alpha = 1\%$ is plotted in red. The black and purple curves plot p_α for the lowest and highest values of N respectively. (b) False dismissal rate p_β versus sensitivity Θ . The target $p_\beta = 5\%$ is plotted in blue. The black and grey curves plot p_β for the lowest and highest values of $2\mathcal{F}_\alpha$ respectively, and correspond to the like-coloured curves in figure (a).

adopt the range $8.1 \times 10^9 < N < 2.8 \times 10^{10}$, which encompasses the actual number of templates used in each upper limit band (see section 9.5). The marginalisation of the cumulative distribution function is more easily computed⁵ using a series expansion of the non-central χ^2 distribution (Johnson & Kotz 1970):

$$\left\langle \text{cdf}(\chi_4^2(\rho^2); 2\mathcal{F}_\alpha) \right\rangle = \sum_{i=0}^{\infty} \left\langle e^{-\frac{1}{2}\rho^2} \frac{(\frac{1}{2}\rho^2)^i}{i!} \right\rangle \text{cdf}(\chi_{4+2i}^2; 2\mathcal{F}_\alpha). \quad (8.19)$$

The series is a weighted sum of central χ^2 distributions with $4 + 2i$ degrees of freedom, where the weights are Poisson distributions of i with mean $\rho^2/2$.

Figure 8.6a plots p_α against $2\mathcal{F}_\alpha$, using equation 8.14; the desired detection threshold takes the range $61.8 \lesssim 2\mathcal{F}_\alpha \lesssim 64.3$. Substituting into equation 8.15, Figure 8.6b plots p_β against Θ . Thus, we see that the sensitivity factor for the Cas A search is ~ 35 , in agreement with Wette *et al.* (2008).

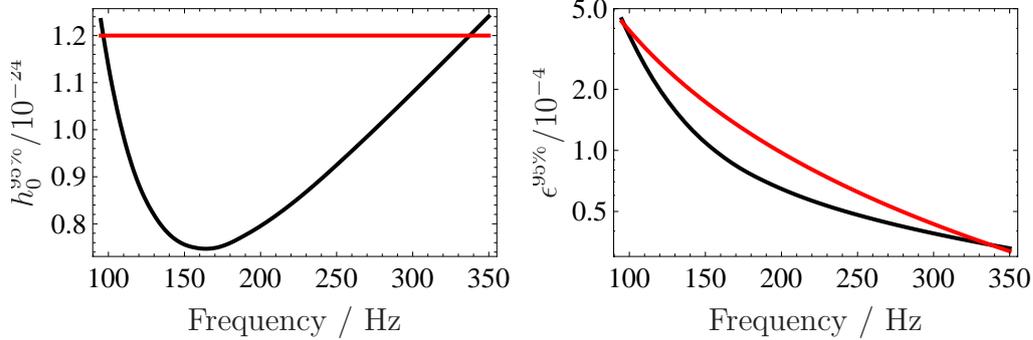


Figure 8.7: Expected upper limits obtainable by the Cas A search. Plotted in black are estimated 95% confidence upper limits on h_0 (left) and ϵ (right); plotted in red are the indirect upper limits of equations 8.9 and 8.8.

8.6 Computational cost and expected upper limits

Finally, we estimate the computational cost of a coherent search of 12 days of LIGO data (section 8.3) over the parameter space of section 8.4, and the sensitivity (section 8.5) we can expect from such a search.

The initial estimate of the computational cost (Wette *et al.* 2008) was 20 days on 200 nodes of the `ac` cluster⁶ of the Australian NCI National Facility (see bibliography). The search was, instead, run on the ATLAS cluster.⁷ We measure the performance of `ComputeFStatistic_v2` on ATLAS to be $\sim 2.4 \times 10^{-7}$ seconds per template per SFT. The estimated number of templates, $N \approx 6.5 \times 10^{12}$, is given by equation 8.12. Finally, the data set selected for the Cas A search (see section 9.1.1) contains 934 SFTs, implying a duty cycle $T_{\text{obs}}/T_{\text{span}} \approx 80\%$. Multiplying these quantities together the total computational cost of the Cas A search:

$$3.3 \text{ days} \left(\frac{f_{\text{max}}}{300 \text{ Hz}} \right)^3 \left(\frac{300 \text{ years}}{\tau} \right)^3 \left(\frac{T_{\text{span}}}{12 \text{ days}} \right)^7 \left(\frac{5000}{\text{nodes}} \right)^3, \quad (8.20)$$

where the scalings are from Wette *et al.* (2008). The fiducial number of nodes was the approximate capacity of ATLAS at the time the search was run.

⁵ For example, Mathematica has difficulty integrating equation 8.15 directly, but is able to numerically integrate each term in equation 8.19.

⁶ Such is the relentless march of technology that this cluster has since been decommissioned!

⁷At the Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Hanover, Germany.

Figure 8.7 plots the 95% confidence upper limits achievable by the search, using equations 8.13 and 8.10. For S_n , we use the initial LIGO design sensitivity curve (LIGO Laboratory 2009b), divided by $\sqrt{2}$ for two identical detectors. The search is clearly expected to beat the indirect upper limit on h_0 over 100-300 Hz, and by up to $\sim 40\%$ at $f \approx 160$ Hz. The estimated upper limits plotted here are slightly better than plotted in Wette *et al.* (2008), which assumed a more conservative duty cycle of 70%. The estimated upper limits intersect the indirect upper limit at $f \approx 97$ and 340 Hz, which is approximately the band chosen for the Cas A search (Wette *et al.* 2008).

8.7 The expected largest $2\mathcal{F}$

Finally, we consider the potential outcome of the Cas A search, before presenting the results in Chapter 9. We expect the search to be sensitive to only one (potential) gravitational wave source, i.e. Cas A. Therefore, we are interested in only one template, from the $\sim 7 \times 10^{12}$ templates searched: that with the largest value of $2\mathcal{F}$. If we adopt the hypothesis that no gravitational wave signal will be detected, we can derive (and do so in this section) the distribution of expected values for the largest value of $2\mathcal{F}$. If the largest $2\mathcal{F}$ found by the search is significantly larger than its expected value, it would be a first step towards claiming a detection.

In the absence of any signal, and assuming Gaussian noise, the probability distribution of a single value of $2\mathcal{F}$ is a χ^2 distribution with 4 degrees of freedom (section 5.2.2). The probability distribution of the largest of N values of $2\mathcal{F}$ is derived as follows. Let $i = 1, \dots, N$ index the N values of $2\mathcal{F}$. Assume that the j th value is the largest, and denote its value $2\mathcal{F}^*$; the probability of its occurrence is $p(\chi_4^2; 2\mathcal{F}^*)$. For $2\mathcal{F}^*$ to be the largest, the $N - 1$ remaining $2\mathcal{F}$ values must be less than $2\mathcal{F}^*$; the probability of this occurring is

$$p \left[\bigcap_{\substack{i=1 \\ i \neq j}}^N (2\mathcal{F}_i \leq 2\mathcal{F}^*) \right] = \prod_{\substack{i=1 \\ i \neq j}}^N p(2\mathcal{F}_i \leq 2\mathcal{F}^*) \quad (8.21)$$

$$= \left[\text{cdf}(\chi_4^2; 2\mathcal{F}^*) \right]^{N-1}. \quad (8.22)$$

Finally, any of the N values of $2\mathcal{F}$ may have been the largest $2\mathcal{F}$, so we multiply by N possible combinations. In short, the probability of $2\mathcal{F}^*$ being the largest of a collection of N values of $2\mathcal{F}$ is

$$p(N; 2\mathcal{F}^*) = N p(\chi_4^2; 2\mathcal{F}^*) \left[\text{cdf}(\chi_4^2; 2\mathcal{F}^*) \right]^{N-1}. \quad (8.23)$$

We note, however, the following important subtlety. Equation 8.21 assumes, by definition, that each of the N values of $2\mathcal{F}$ are statistically independent, i.e. the probability of the i value of $2\mathcal{F}$ being less than $2\mathcal{F}^*$ depends only on the values of $2\mathcal{F}_i$ and $2\mathcal{F}^*$, and not on the values of i and j (the index of the $2\mathcal{F}^*$ value). This is not necessarily true. To see why, we recall that the value of $2\mathcal{F}$, in the absence of a signal, is given by the sum of the squares of 4 Gaussian-distributed random variables (Jaranowski *et al.* 1998). For each value of $2\mathcal{F}$, the value of the 4 Gaussian random variables are computed from the SFT data; there are, however, many more values of $2\mathcal{F}$ than there are SFT data samples. For example, at 100 Hz, $\sim 10^8$ templates are required to cover a $\sim 2 \times 10^{-2}$ Hz band (section 9.2), whereas the number of spectral bins $\tilde{x}(f)$ in this band is approximately 2×10^{-2} Hz $\times T_{\text{SFT}} N_{\text{SFTs}} \approx 3 \times 10^4$ bins. Assuming that the spectral bins are all independent random samples of the detector noise, there are nevertheless an insufficient number of them from which to generate $\sim 4 \times 10^8$ independent random values. Therefore, the values of the Gaussian variables will, to some degree, be statistically correlated, and it follows that the N values of $2\mathcal{F}$ will also be correlated. We note that this statistical correlation between $2\mathcal{F}$ values is separate from the mismatch between the corresponding templates in parameter space; as there is no signal present, it is unlikely that the parameter space metric remains an accurate descriptor of this mismatch between templates.

The approach we adopt to this problem is to assume that the $2\mathcal{F}$ values are statistically independent, so that equation 8.23 holds, but that the total number of templates N is substituted with the total number of equivalent statistically independent templates $\bar{N} \leq N$. Note that \bar{N} is *not* the size of the subset of $2\mathcal{F}$ values which are statistically independent, implying a complementary subset of $N - \bar{N}$ $2\mathcal{F}$ values which are statistically dependent. Instead, *all* N values of $2\mathcal{F}$ are mutually statistically dependent, and \bar{N} is simply the number with which to replace N in order for equation 8.23 to be valid.

The degree of mutual statistical dependence between the N values of $2\mathcal{F}$ is a mathematical problem which we do not attempt to investigate further. Instead, we present below an empirical method to determine \bar{N} . We demonstrate that, for the purpose of determining the expected distribution of the largest $2\mathcal{F}$ for the entire search, the N values of $2\mathcal{F}$ may be assumed to be statistically independent.

8.7.1 Number of statistically independent templates

In section 9.2, the search frequency band was partitioned into smaller job frequency bands, such that each job band is covered by approximately the

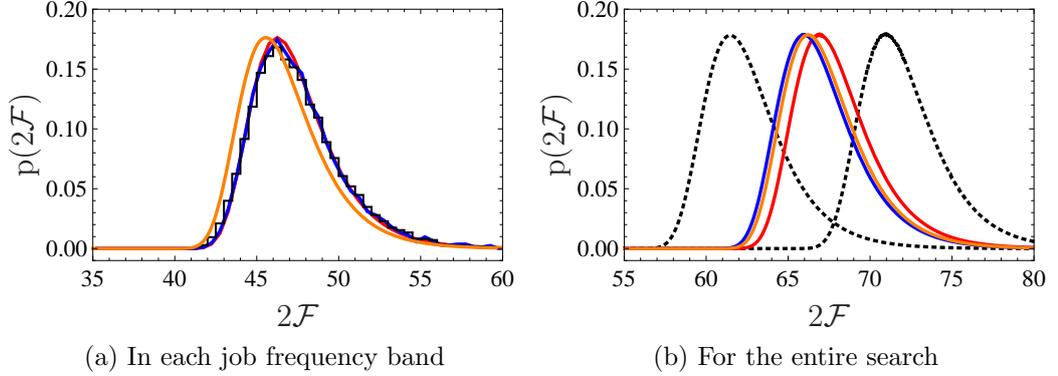


Figure 8.8: (a) Distribution of the largest $2\mathcal{F}$ in each search job frequency band. The histogram of $2\mathcal{F}$ values, in black, is closely fitted by $p(\bar{N}_{\text{job}} \approx 1.39N_{\text{job}}; 2\mathcal{F}^*)$ in red, and by $P(\bar{N}_{\text{job}} \approx 0.88N_{\text{job}}; 2\mathcal{F}^*)$ in blue. For comparison, $p(N_{\text{job}}; 2\mathcal{F}^*)$ is plotted in orange. (b) Distribution of the largest $2\mathcal{F}$ for the entire search. Plotted are $p(\bar{N} \approx 1.39N; 2\mathcal{F}^*)$ in red, $P(\bar{N} \approx 0.88N; 2\mathcal{F}^*)$ in blue, and $p(N; 2\mathcal{F}^*)$ in orange. For illustration, $p(0.1N; 2\mathcal{F}^*)$, left, and $p(10N; 2\mathcal{F}^*)$, right, are plotted in black (dotted).

same number of templates. As a consequence, the largest $2\mathcal{F}$ in each job band should be drawn from the same distribution: $p(\bar{N}_{\text{job}}; 2\mathcal{F}^*)$, given by equation 8.23, where \bar{N}_{job} is the number of statistically independent templates per job band. Therefore, we can determine \bar{N}_{job} by fitting $p(\bar{N}_{\text{job}}; 2\mathcal{F}^*)$ to a histogram of the largest $2\mathcal{F}$ values found by the 21439 jobs. This gives the ratio $\bar{N}_{\text{job}}/N_{\text{job}} = \bar{N}/N$, and so \bar{N} is determined.

The fit to determine \bar{N}_{job} is shown in Figure 8.8a. The histogram of the largest $2\mathcal{F}$ values is plotted in black, and the fitted curve $p(\bar{N}_{\text{job}}; 2\mathcal{F}^*)$ is overplotted in red. The fitted ratio $\bar{N}_{\text{job}}/N_{\text{job}} \approx 1.39$ is clearly wrong, as we expect $\bar{N}_{\text{job}} \leq N_{\text{job}}$, with equality holding if all $2\mathcal{F}$ values were indeed statistically independent. In Figure 8.8a we also plot $p(N_{\text{job}}; 2\mathcal{F}^*)$ (i.e. $\bar{N}_{\text{job}}/N_{\text{job}} = 1$); the histogram is shifted to the right relative to this curve, indicating that the $2\mathcal{F}$ values are being positively biased.

The reason for this bias is that `ComputeFStatistic_v2` is not an exact implementation of the \mathcal{F} statistic. Small modifications to the algorithm, designed to decrease the computational cost (Prix 2008), introduce small differences between the distribution of $2\mathcal{F}$, as computed by `ComputeFStatistic_v2`, and a χ^2 distribution with 4 degrees of freedom. There are also additional biases in the noise estimation (Prix 2006), which have not been taken into account here. Figure 8.9a shows the fractional difference between theoretical distribution $p(\chi_4^2; 2\mathcal{F})$, and a histogram $P(2\mathcal{F})$ of all $2\mathcal{F}$ values

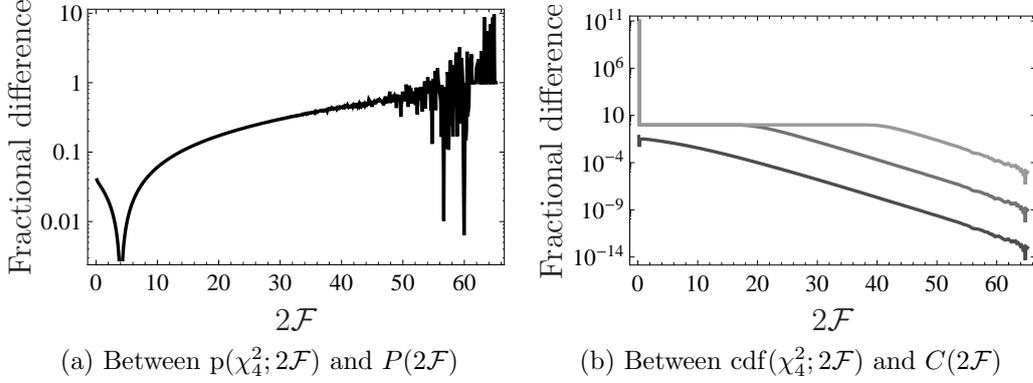


Figure 8.9: Fractional difference between the theoretical distribution of the \mathcal{F} statistic and the distribution of the \mathcal{F} statistic as implemented by `ComputeFStatistic_v2`. The latter is derived from a histogram of all $2\mathcal{F}$ values produced by the Cas A search, after the removal of outliers as described in section 9.3. (a) Fractional difference $|p(\chi_4^2; 2\mathcal{F}) - P(2\mathcal{F})|/p(\chi_4^2; 2\mathcal{F})$ between the theoretical distribution, $p(\chi_4^2; 2\mathcal{F})$, and the `ComputeFStatistic_v2` distribution, $P(2\mathcal{F})$. (b) Fractional difference $|\text{cdf}(\chi_4^2; 2\mathcal{F})^N - C(2\mathcal{F})^N|/\text{cdf}(\chi_4^2; 2\mathcal{F})^N$ between the theoretical cumulative distribution, $\text{cdf}(\chi_4^2; 2\mathcal{F})$, and the `ComputeFStatistic_v2` cumulative distribution, $C(2\mathcal{F})$, for (dark to light grey) $N = 1, 10^4$, and 10^8 .

produced by the Cas A search. All outliers identified in the post-processing described in section 9.3 are removed prior to computing $P(2\mathcal{F})$, which thus represents the distribution of the \mathcal{F} statistic, as implemented by `ComputeFStatistic_v2`, in the absence of any non-Gaussian features. The fractional difference increases from a minimum at $2\mathcal{F} = 4$, where both distributions are approximately at their maxima; at higher $2\mathcal{F}$, the histogram contains increasingly fewer samples and the difference increases, becoming noisy at $2\mathcal{F} \gtrsim 45$. While the fractional difference between the equivalent cumulative distributions, $\text{cdf}(\chi_4^2; 2\mathcal{F})$ and $C(2\mathcal{F}) = \int_0^{2\mathcal{F}} P(2\mathcal{F}')d(2\mathcal{F}')$, is also small, this difference is rapidly magnified when both distributions are raised to large powers, as shown in Figure 8.9b. Thus, the computation of equation 8.23 magnifies small differences between the theoretical and practical distributions of the \mathcal{F} statistic.

To account for these differences, we also fit a modified version of equation 8.23 to the histogram of the largest $2\mathcal{F}$ values from each job band (Figure 8.8a), where the histogram $P(2\mathcal{F})$ is substituted for the theoretical \mathcal{F} statistic distribution:

$$P(N; 2\mathcal{F}^*) = NP(2\mathcal{F}^*)C(2\mathcal{F}^*)^{N-1} \quad (8.24)$$

The fitted curve $P(\bar{N}_{\text{job}}; 2\mathcal{F}^*)$ is plotted in blue in Figure 8.8a; the fitted ratio $\bar{N}_{\text{job}}/N_{\text{job}} \approx 0.88$. Thus, we empirically estimate that the equivalent statistically independent templates is $\bar{N} = 0.88N \approx 6.2 \times 10^{12}$.

Figure 8.8b plots the expected distribution of the largest $2\mathcal{F}$ value found by the Cas A search. The three curves, which represent different choices for \bar{N} of $\sim 0.88N$, N , and $\sim 1.39N$, are nonetheless very similar; they predict that the largest $2\mathcal{F}$ value found by the Cas A search, in the absence of a signal, is mostly likely to lie in the range $62 \lesssim 2\mathcal{F}^* \lesssim 75$.

The similarity of the distributions can be explained by the following observations: at high $2\mathcal{F}$ values, $\text{cdf}(\chi_4^2; 2\mathcal{F})^N$ and $C(2\mathcal{F})^N$ remain relatively similar (see the light grey curve in Figure 8.9b), and the position of the distributions is only weakly dependent on N . To increase the position of the maximum of $p(\chi_4^2; 2\mathcal{F})$ from ~ 61 to ~ 71 , N must be increased by 2 orders of magnitude (see Figure 8.8b). We conclude, therefore, that for the purpose of estimating the largest $2\mathcal{F}$ value likely to be found by the Cas A search, templates can be regarded as statistically independent.

8.8 Summary

In this chapter, we presented the motivation for a periodic gravitational wave search targeting Cas A. We first reviewed astronomical observations of Cas A conducted since the discovery of its CCO, and suggested that the CCO might be a gravitational wave source. We derived an indirect upper limit, based on energy conservation, on the strength of gravitational waves from Cas A. We then presented the proposed search, including the choice of analysis method, the time span of the data set, and the frequency and spindown parameter space to be searched. We determined the computational cost of the search, demonstrated that the search will beat the indirect upper limit, and predicted the largest value of $2\mathcal{F}$ that will be found by the search in the absence of a detection. Chapter 9 will present the results of this search.

Chapter 9

Upper limits on gravitational waves from Cassiopeia A

In Chapter 8, we motivated a search for gravitational waves targeting the supernova remnant Cassiopeia A (Cas A), and demonstrated that such a search would beat the indirect limit on the strength of gravitational waves based on energy conservation. In this chapter, we present the results of this search. We start by describing the LIGO S5 science run, and the selection from it of the data set used in the search (section 9.1). We then describe the search pipeline (section 9.2). We show how the results of the search may become contaminated by instrumental noise artifacts, present a method which identifies where this occurs, and apply it to the search (section 9.3). From the remaining results, we find the largest value of the detection statistic, and discuss whether they indicate that we have detected a gravitational wave signal (section 9.4); we conclude that they do not. Finally, we present upper limits derived from the search on gravitational waves from Cas A (section 9.5).

See section 1.1 for author contributions and publications, a disclaimer, and acknowledgements relevant to this chapter.

9.1 The LIGO S5 science run

The S5 science run is LIGO's fifth and most recently completed science run. It commenced at UTC 16:00 November 4, 2005 (GPS¹ 815155213) at Hanford, and exactly 10 days later at Livingston; it ended at UTC 00:00 October

¹ The Global Positioning System (GPS) is used to time-stamp LIGO data (Abbott *et al.* 2009e). GPS times give the number of seconds (excluding UTC corrections) since UTC 00:00:00 January 6, 1980.

1, 2007 (GPS 875232014; Abbott *et al.* 2009b). The S5 run collected over a year of science-quality data coincident between i) all three detectors, and ii) the two detector sites (Lazzarini 2007). The duty cycle of triple-coincident science data rose from 48% to 67%, and averaged 54%, over the course of the run. Interruptions by a wide variety of environmental and anthropogenic disturbances, as well as scheduled breaks for maintenance and commissioning of equipment, account for the downtime. During S5, the detectors were operating at very near their design sensitivities; the strain noise of the two 4-km detectors was, on average, less than $3 \times 10^{-23} \text{ Hz}^{-1/2}$ at their most sensitive frequencies (around 140 Hz) and less than $5 \times 10^{-23} \text{ Hz}^{-1/2}$ over 100–300 Hz. The S5 run successfully met the goal of the initial phase of LIGO: to collect a year of science data at design sensitivity.

Acquired science data passes through a number of post-acquisition procedures. Times when the data may be unsuitable for analysis, for a number of reasons (Riles *et al.* 2009, Zweizig *et al.* 2009), are marked with data quality flags, which are then used to exclude undesired data segments. The remaining data is calibrated (Abbott *et al.* 2009e) to produce a non-contiguous time series of gravitational wave strain, $h(t)$. Finally, the time series are broken into segments of length $T_{\text{SFT}} = 30$ min which are used to create SFTs.

9.1.1 Selection of data for the search

The Cas A search uses S5 data from the two 4-km arm-length interferometers at Livingston and Hanford, denoted L1 and H1 respectively. During S5, LIGO also operated a third, 2-km arm-length interferometer at Hanford, denoted H2, and co-located with H1; its data was not analysed. Data from H2 has roughly twice the strain noise density of data from L1 and H1, due the halved arm length, but its analysis would still carry the same computational cost. The SFTs used in the search are selected from those generated from data acquired before GPS 861000000 (approx. UTC April 19, 2007); the first SFTs generated from data acquired after this time were subsequently found to have used incorrect data segments (Betzwieser *et al.* 2008), and had not been re-generated at the time the search was conducted. The S5 run collected approximately 9525 hours of H1 and 7959 hours of L1 science data before GPS 861000000 (Riles *et al.* 2009). The data quality flags applied to this data are shown in Table 9.1. From the remaining data, 17309 H1 and 13784 L1 SFTs² were generated. The fraction of available science data not included in the SFTs, either due to data quality vetoes or to losses during

² The number of H1 and L1 S5 SFTs before GPS 861000000 was determined from the SFTs on the ATLAS cluster on September 4, 2009, using the `ligo_data_find` software.

Table 9.1: Data quality flags applied to the L1 and H1 science data, up to GPS 861000000, before generation of the S5 SFTs (Landry & Mendell 2007, Betzwieser *et al.* 2008). Descriptions are derived from Riles *et al.* (2009) and Zweizig *et al.* (2009).

Name	Description
AS_TRIGGER	Saturation of output photodiodes
CALIB_BAD_COEFFS_60	Bad calibration coefficients
CALIB_DROPOUT_1SAMPLE	Calibration waveform generator dropouts
CALIB_DROPOUT_1SEC	...
CALIB_DROPOUT_AWG_STUCK	...
CALIB_GLITCH_ZG	...
INVALID_DARMERR	Data acquisition system errors
MISSING_RDS_C03_L2 ^a	Short gaps due to $h(t)$ finite impulse filters
OUT_OF_LOCK	Corrections to end-of-lock times
PRE_LOCKLOSS_10_SEC	Noisy data 10 seconds prior to lock-loss
PRE_LOCKLOSS_30_SEC	Noisy data 30 seconds prior to lock-loss
Wind_Over_30MPH ^b	High winds measured by weather stations

^a This flag appears not to have been applied to the L1 data before GPS 846374414 (Landry & Mendell 2007). Nevertheless, the only flagged segment in this interval (GPS 838066816–838066848; Mendell 2008) spans only 32 seconds, and does not appear in the list of L1 segments used to generate SFTs (Landry & Mendell 2007).

^b This flag is defined for H1 only.

the 30-minute segmentation, totals $\sim 9\%$ for H1 and $\sim 13\%$ for L1.

The time span of the Cas A search data set is $T_{\text{span}} = 12$ days (section 8.3). To choose the start time t_0 of the data, we consider all intervals $[t_0, t_0 + T_{\text{span}}]$ contained within the time span of the available SFTs, and compute a figure of merit $F(t_0)$ for each interval. The figure of merit is chosen to favour intervals containing as many SFTs as possible, and SFTs with the lowest strain noise possible; it is given by

$$F(t_0) = \frac{T_{\text{SFT}}}{N_{\text{IFO}} T_{\text{span}}} \sum_{\substack{\text{SFTs in} \\ [t_0, t_0 + T_{\text{span}}]}} \frac{\min_{\text{all SFTs}} S_{\text{SFT}}}{S_{\text{SFT}}}, \quad (9.1)$$

where $N_{\text{IFO}} = 2$ is the number of interferometers and, for each SFT,

$$\frac{1}{S_{\text{SFT}}} = \sum_{f=100 \text{ Hz}}^{300 \text{ Hz}} \frac{1}{S_n(f)}, \quad (9.2)$$

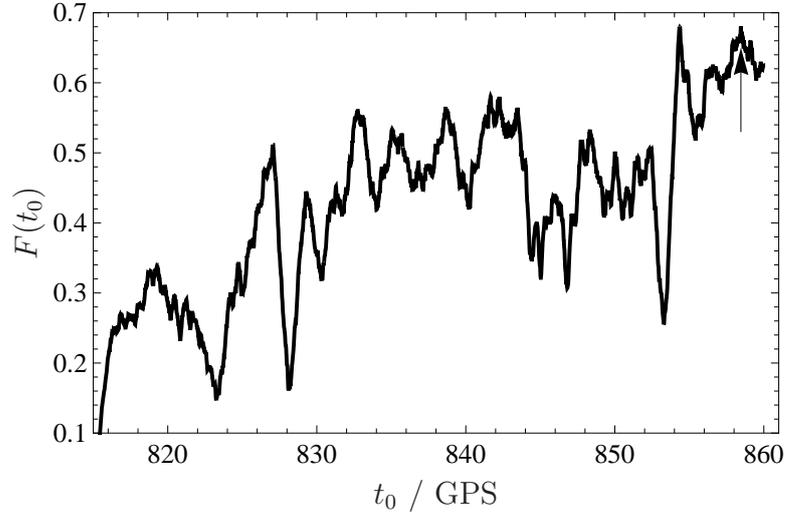


Figure 9.1: Figure of merit $F(t_0)$ for all 12-day intervals containing SFTs before GPS 861000000, as a function of the interval start time t_0 . The arrow indicates the start time of the interval chosen for the Cas A search. GPS times of tick marks are abbreviated to their first 3 digits.

where $S_n(f)$ is the noise power spectral density of each SFT. The normalisation of equation 9.1 is chosen so that $0 \leq F \leq 1$.

Figure 9.1 plots the figure of merit $F(t_0)$ for all possible 12-day intervals within the time span of the available SFTs. The maximum value of F is achieved by the interval beginning at $t_0 = \text{GPS } 858459411$ (UTC 20:56:37 March 20, 2007), and ending at GPS 859495818 (UTC 20:50:04 April 1 2007). This interval is chosen as the Cas A search data set. It contains 934 SFTs in total, 445 from L1 and 489 from H1. The average amount of data from each detector is ~ 9.7 days, giving a duty cycle of $\sim 81\%$ over the 12-day interval. Figure 9.2 shows spectrograms of the SFTs. The lowest strain noise level is at ~ 150 Hz; at the extremes of the search band, the strain noise is a factor of ~ 2 – 3 higher. The occasional red/green horizontal stripe denotes SFTs of data acquired when the detectors were unusually noisy. The vertical red/green bands correspond to lines (i.e. narrow-band, constant-frequency, non-Gaussian noise) in the data; the three most prominent lines, at 120, 180, and 240 Hz, are harmonics of the 60 Hz frequency of the U.S. A.C. mains power. There also appears to be a wandering noise feature in the H1 data between 180 and 230 Hz; this feature is unlikely to impact upon the search, as it changes frequency too rapidly to overlap with a search template for any significant length of time.

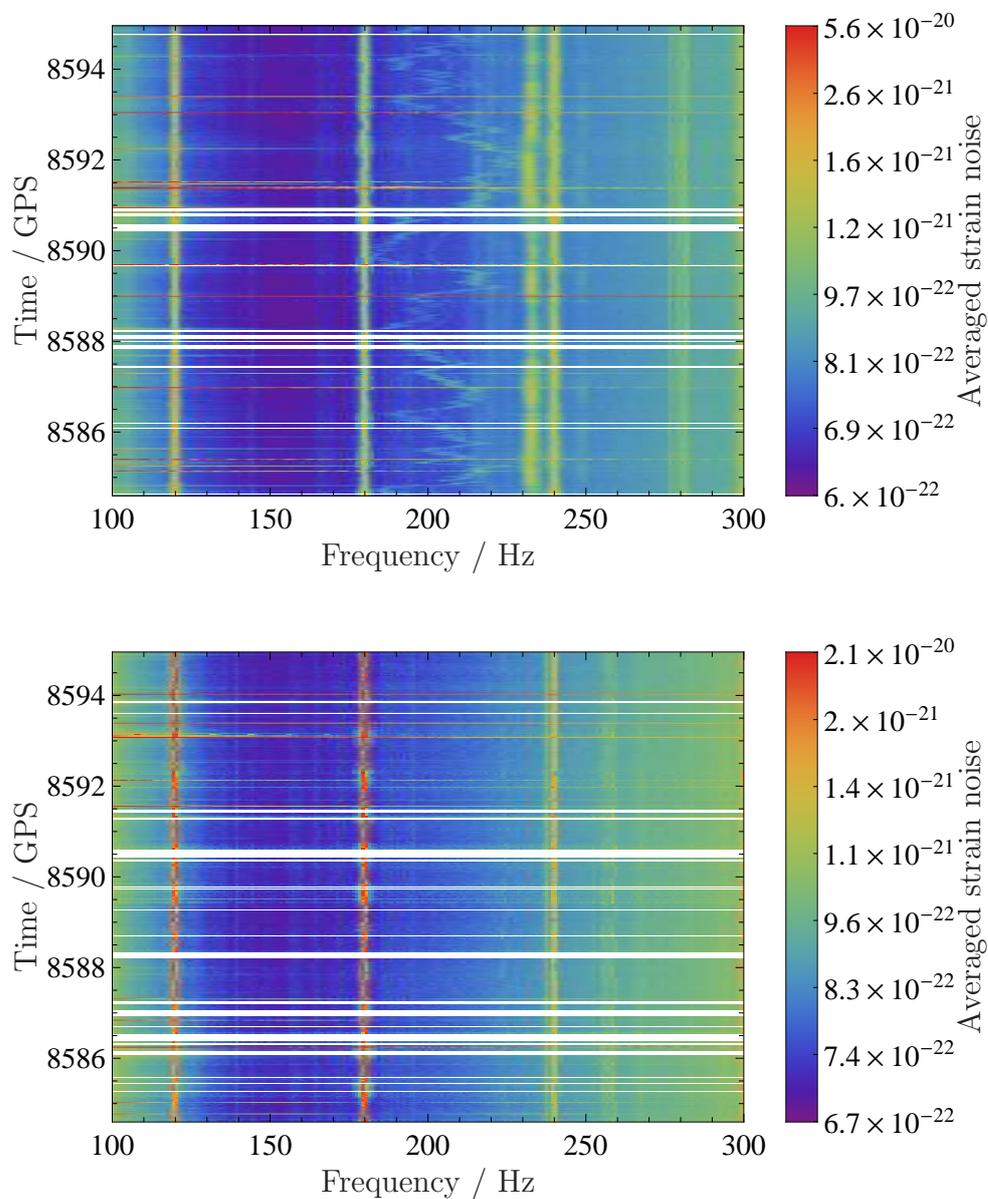


Figure 9.2: Spectrograms of the H1 (top) and L1 (bottom) SFTs comprising the Cas A search data set. The vertical axes give the first 4 digits of the GPS starting times of each SFTs; white gaps in the spectrograms denotes times where no SFTs exist. Note that the colour scales for each spectrogram are different from each other, and are non-linear.

9.2 The search pipeline

We search the data set, selected in section 9.1.1, for periodic gravitational waves signals at the sky position of Cas A, and over the frequency and spin-down parameter space of section 8.4. To distribute the search over the ATLAS cluster, we partition the full search frequency band (100–300 Hz) into 21439 *job frequency bands*. The widths of the job bands are chosen so that roughly the same number of templates is required to cover each band; that number is $N_{\text{job}} = (3.27 \pm 0.02) \times 10^8$ templates. Because the size of the spin-down parameter space increases with frequency, the widths of the job bands decrease from $\sim 2.0 \times 10^{-2}$ Hz (at 100 Hz) to $\sim 1.8 \times 10^{-3}$ Hz (at 300 Hz).

For each job band, a separate Condor (see bibliography) job is submitted to the ATLAS cluster. Each job calls `ComputeFStatistic_v2` (section 8.3), which uses the algorithm of Chapter 7 to generate a bank of templates over the given job band, and the full range of first and second spindowns. The maximum mismatch μ_{max} is set to 20%. `ComputeFStatistic_v2` then computes the value of $2\mathcal{F}$ for each template in the bank; to limit the size of the output files, only the top 0.01% templates with the largest values of $2\mathcal{F}$ are recorded.

Because the jobs are given the same number of templates to search, they take approximately the same time (~ 20 hours) to complete. Thus, the total computational cost of the search is 21439×20 hours ≈ 3.5 days, in good agreement with the prediction of equation 8.20.

9.3 Vetoing of spuriously large $2\mathcal{F}$

Figure 9.3 shows the raw output of the Cas A search: the largest value of $2\mathcal{F}$ found in each of the job frequency bands, prior to the post-processing described in this section. The majority of $2\mathcal{F}$ values lie in the range $40 \lesssim 2\mathcal{F} \lesssim 60$; there are, however, a number of $2\mathcal{F}$ significantly above the background level. Since we expect the search to be sensitive to at most one gravitational wave source, i.e. Cas A, it is unlikely that these outliers are gravitational wave signals. Indeed, a careful inspection of the spectrograms in Figure 9.2 reveals that many of the $2\mathcal{F}$ outliers appear to be at the same frequencies as visible line features. For example, the $2\mathcal{F} \sim 90$ outlier is close to the 120 Hz mains power harmonic, and the $2\mathcal{F} \sim 95$ –110 outliers appear to coincide with a faint line at ~ 145 Hz.

In section 9.3.1, we describe how it is possible for the computation of $2\mathcal{F}$ to be contaminated by lines. In section 9.3.2, we describe a procedure which identifies narrow line features which may result in unreasonably large

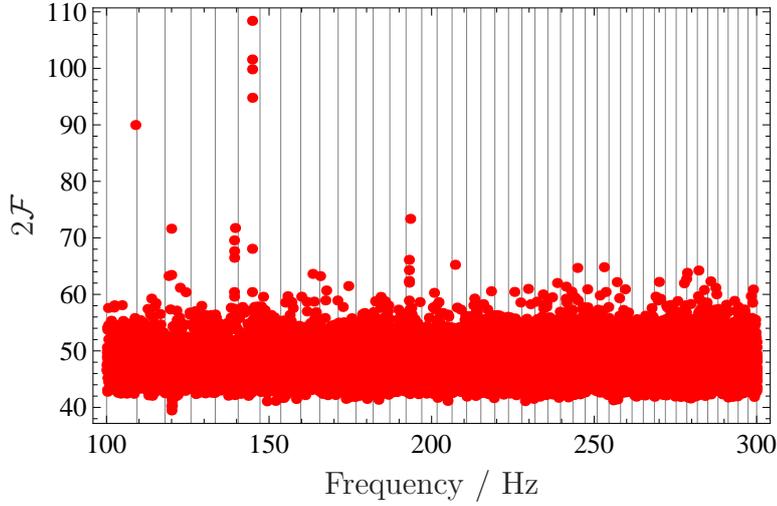


Figure 9.3: Largest $2\mathcal{F}$ found in each search job frequency band, before the post-processing described in section 9.3. The thin vertical lines mark the divisions between groups of 500 search job frequency bands.

values of $2\mathcal{F}$. The application of this procedure to the Cas A search results is described in section 9.3.3

9.3.1 Contamination of the \mathcal{F} statistic by lines

We start with the \mathcal{F} statistic as expressed by equation 5.42, using the A , B , C , and D coefficients of equation 5.43. Ignoring the difference between the h_i components of the signal waveform, $A \sim B \sim C \propto (h||h)$, while $D \propto (h||h)^2$. Noting that, from equation 5.33, $(x||y) \approx \tilde{x}(f)\tilde{y}(f)^*/S_n(f)$ over a narrow frequency band centred on f , we find that

$$\begin{aligned}
 2\mathcal{F} &\propto \frac{1}{(h||h)^2} \left\{ (h||h)(x||h)^2 \right\} \\
 &= \frac{\left[\tilde{x}(f)\tilde{h}(f)^*/S_n(f) \right]^2}{\left[\tilde{h}(f)\tilde{h}(f)^*/S_n(f) \right]} \\
 &= \frac{|\tilde{x}(f)|^2}{S_n(f)}.
 \end{aligned} \tag{9.3}$$

In other words, the value of the \mathcal{F} statistic is, very roughly, proportional to the SFT power, $|\tilde{x}(f)|^2$, divided by the noise power spectral density, $S_n(f)$, at the same frequency. Any outlying non-Gaussian noise in $\tilde{x}(f)$, such as

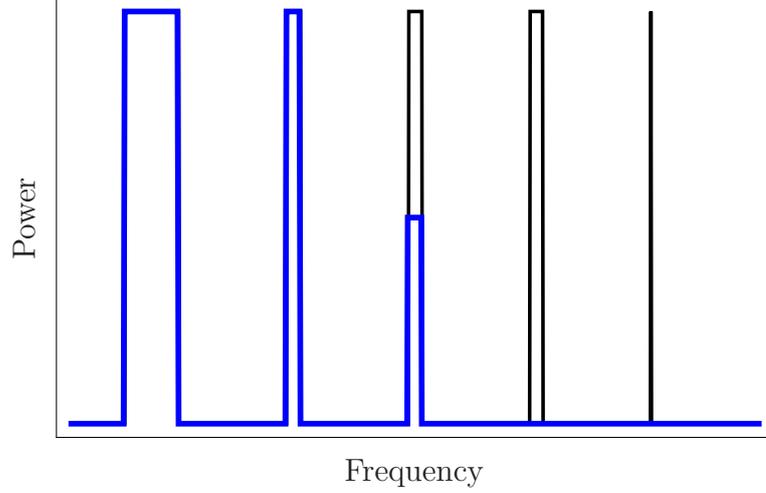


Figure 9.4: The action of a spectral running median (blue), with a window size of 50 bins, on a discrete spectrum (black) containing lines with widths of (left to right) 100, 26, 25, 24, and 1 bin.

an instrumental line, will therefore result in a spuriously large value of $2\mathcal{F}$, unless it is removed by an accurate estimation of the noise floor $S_n(f)$.

A robust method of estimating $S_n(f)$, used by `ComputeFStatistic.v2`, is the spectra running median (Mohanty 2003, Abbott *et al.* 2007b). The noise spectral density $S_n(f)$ is given by the median value of the periodogram

$$s_n(f) = \frac{1}{N_{\text{SFTs}}} \sum_{\text{SFTs}} \frac{2|\tilde{x}(f)|^2}{T_{\text{SFT}}} \quad (9.4)$$

within a frequency window of N bins centred on f :

$$S_n(f) = \text{median} \left\{ s_n(f') \left| \frac{|f' - f|}{T_{\text{SFT}}} \leq \frac{N}{2} \right. \right\}. \quad (9.5)$$

The use of the median makes the method robust against outlying non-Gaussian noise; because using the median to estimate the mean introduces a systematic bias, $S_n(f)$ is also divided by a normalisation factor. A typical window size, used in this analysis and elsewhere (Abbott *et al.* 2007b), is $N = 50$, which for 30-minute SFTs gives a bandwidth of ~ 27.8 mHz.

The running median will track any spectral features in the SFT power $|\tilde{x}(f)|^2$, provided that such features are wider than half the running median window. Spectral features smaller than this size will be removed from $S_n(f)$. This behaviour is illustrated in Figure 9.4. For the large 100-bin feature, $2\mathcal{F}$

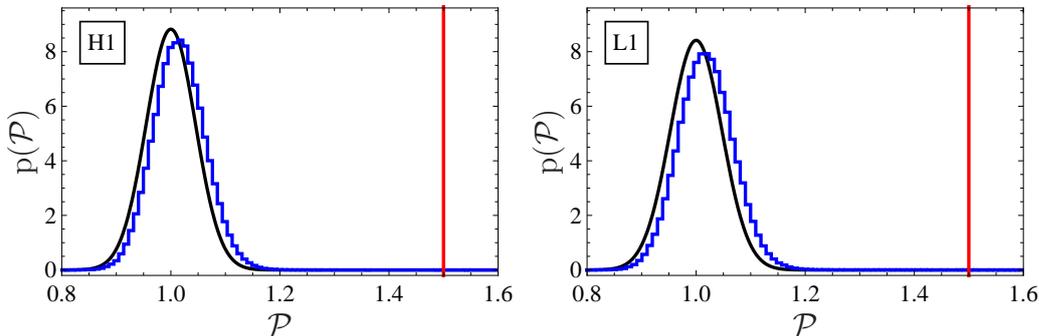


Figure 9.5: Histograms (blue) of the StackSlide power $\mathcal{P}(f)$, $100 \text{ Hz} \leq f \leq 300 \text{ Hz}$, of the Cas A search SFTs. Gaussian distributions with unit means and standard deviations, calculated from the number of SFTs, of ~ 0.045 for H1 (left), and ~ 0.047 for L1 (right) are plotted in black. The veto threshold \mathcal{P}_{thr} is plotted in red.

would remain unchanged as the feature remains in both $|\tilde{x}(f)|^2$ and $S_n(f)$; the 1-bin narrow line, however, remains in $|\tilde{x}(f)|^2$ but is removed by the running median. As a result, at the frequency of the narrow line feature, $|\tilde{x}(f)|^2$ increases while $S_n(f)$ remains constant, and hence $2\mathcal{F}$ increases despite no signal being present.

9.3.2 Identification of narrow line features

We first compute the power in the search SFTs, normalise by the running median, and averaged over time:

$$\mathcal{P}(f) = \frac{2}{N_{\text{SFTs}} T_{\text{SFT}}} \sum_{\text{SFTs}} \frac{|\tilde{x}(f)|^2}{S_n(f)}. \quad (9.6)$$

This quantity is simply the StackSlide power (Abbott *et al.* 2007b); thus, we essentially perform a StackSlide search for constant-frequency, un-Doppler modulated signals. For SFTs containing only Gaussian noise, $\mathcal{P}(f)$ is Gaussian distributed, with unit mean and standard deviation $N_{\text{SFTs}}^{-1/2}$.

We identify frequency bands where \mathcal{P} exceeds a threshold $\mathcal{P}_{\text{thr}} = 1.5$ in either H1 or L1. The threshold is chosen to be far above the expected mean of \mathcal{P} , so we can confidently assume that it will identify only non-Gaussian features in the SFTs. To confirm, Figure 9.5 plots histograms of \mathcal{P} , computed from the search SFTs over the search frequency band. The distributions of \mathcal{P} are close to their expected Gaussian distributions, with slight positive bias (see section 8.7.1), and the threshold is well above the mean.

Table 9.2: Vetoed frequency bands identified by the threshold on $\mathcal{P}(f)$. The first column gives the figure detailing each band. The remaining columns give: the lower (f_L) and upper (f_U) frequencies defining the band; the maximum values of \mathcal{P} and $2\mathcal{F}$ found in each band; and a description of the instrumental line or other phenomenon that triggered the veto.

Fig.	f_L / Hz	f_U / Hz	\mathcal{P}_{\max}	$2\mathcal{F}_{\max}$	Description
9.7a	108.842	108.877	1.8	89.97	Pulsar hardware injection no. 3
9.14a	119.858	119.896	5.3	71.62	Associated with 60 Hz harmonic
9.14b	127.983	128.018	4.4	55.67	16 Hz harmonic
9.15	139.167	139.283	3.9	69.56	Definite L1-only line
9.8a	139.493	139.528	9.4	71.75	Probable L1-only line
9.16	144.696	144.807	4.6	108.40	Definite L1-only line
9.14c	179.794	179.829	1.6	51.36	Associated with 60 Hz harmonic
9.17	185.575	185.684	1.9	58.78	Definite L1-only line
9.18	192.959	193.052	2.2	66.12	Definite L1-only line
9.7b	193.373	193.408	2.2	73.37	Pulsar hardware injection no. 8
9.8b	209.248	209.283	2.1	53.92	Probable L1-only line

We refer to the frequency bands, identified as above, as *veto bands*. The veto bands are centred on the SFT frequency bin where \mathcal{P} exceeds the threshold, and the widths of the bands are equal to the running median window. Veto bands which overlap are merged together. Templates whose instantaneous frequency $f(t)$ (including spindown and Doppler modulation; see section 5.2.1), at any time during the time span of the search SFTs, falls within any of the veto bands are excluded from the search results.

The above procedure does not necessarily identify all narrow line features in the Cas A search SFTs, nor is it guaranteed to identify all instances where line features contaminate the \mathcal{F} statistic. Provided that the procedure identified and removes all instances of contaminated values of $2\mathcal{F}$, so that the loudest remaining value of $2\mathcal{F} = 2\mathcal{F}^*$ shows no sign of contamination, then all remaining values of $2\mathcal{F} < 2\mathcal{F}^*$ are irrelevant, since contamination can only increase the value of $2\mathcal{F}$.

9.3.3 Application to the search results

Table 9.2 shows the eleven veto bands identified by the procedure described above. In each figure references in the table are two *veto band plots*, one each for H1 and L1 (as indicated in their upper left hand corner), which show why the veto was triggered, the extent of the veto band, and the templates

removed by the post-processing as a result of the veto: see the top row of plots in Figures 9.7a and 9.7b, and all the plots in Figure 9.14. The veto band plots, which are identical in form, are described below.

The veto band plots show the StackSlide power $\mathcal{P}(f)$ in blue, and the noise power spectral density $S_n(f)$ in purple; the scales of both these quantities are given on the right hand axis, that of $\mathcal{P}(f)$ by the un-bracketed numbers, and that of $S_n(f)$ by the bracketed numbers. The vertical red lines delimit the extent of the veto band. The StackSlide power threshold \mathcal{P}_{thr} is plotted in blue, dashed, where appropriate. It can be seen from these plots that, for each veto band, $\mathcal{P}(f)$ exceeds \mathcal{P}_{thr} in either H1 or L1, as required. The largest value of $2\mathcal{F}$ found in each for the search job frequency bands are plotted as horizontal red or black bars; the scale for these is given on the left hand axis. The bars are identical in both the H1 and L1 plots. The horizontal extent of the bars gives the total range of frequencies covered by the template, corresponding to the value of $2\mathcal{F}$, during the Cas A search data set; the sidereal Doppler modulation is included. Red bars indicates templates which overlap the veto band; black bars indicate templates which never intersect it. For each search job band, the post-processing eliminates all templates in the job band which intersect the veto band, in order of decreasing value of $2\mathcal{F}$, until a non-intersecting template is found. In this case, the template with the largest $2\mathcal{F}$ value in the job band prior to vetoing is plotted as a red bar, the non-intersecting template (which is the template with the largest $2\mathcal{F}$ value after vetoing) is plotted as a black bar, and the two bars are linked with a shaded grey quadrilateral. If no non-intersecting template is found, the entire job band is removed from the search; its loudest $2\mathcal{F}$ template is plotted as a single red bar. Single black bars are the largest $2\mathcal{F}$ templates of job bands unaffected by the post-processing.

Seven of the eleven bands are triggered by instrumental lines that were either well known prior to the search, or else readily identified by follow-up investigations. They are briefly discussed in Appendix 9.A.1. The remaining four veto bands are discussed below.

During the S5 run, ten simulated periodic gravitational wave signals were injected into the LIGO interferometers at the hardware level, by physically manipulating the position of the end mirrors (Abbott *et al.* 2009d). The injections were turned alternately on and off approximately every two weeks (see Figure 9.6). Four of the injections, nos. 0, 3, 6, and 8 (Riles *et al.* 2008), have frequencies within the band searched for Cas A. None of the injections are at the sky position of Cas A. It is well established, however, that the sky position and frequency parameters are globally correlated (Prix & Itoh 2005, Pletsch 2008), and so it is possible for a strong signal, injected at a given sky position, to be visible at a different sky position, with an accompanying

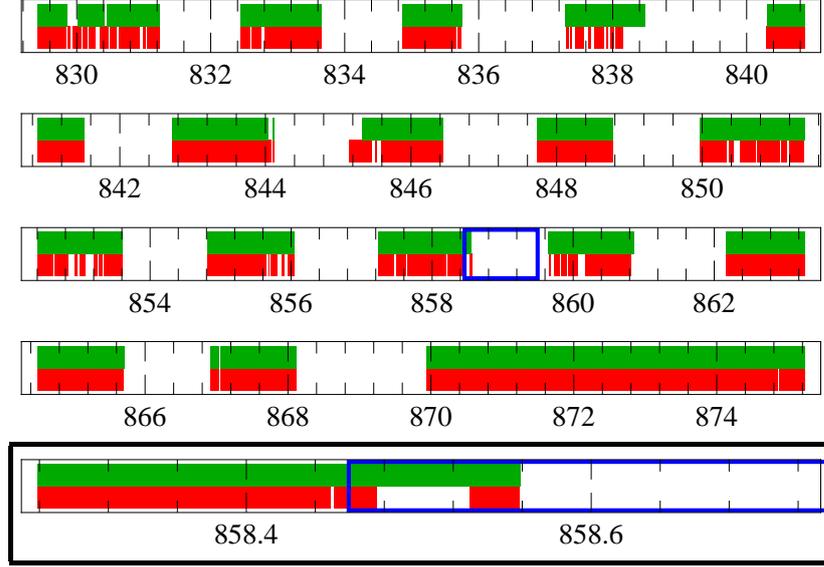
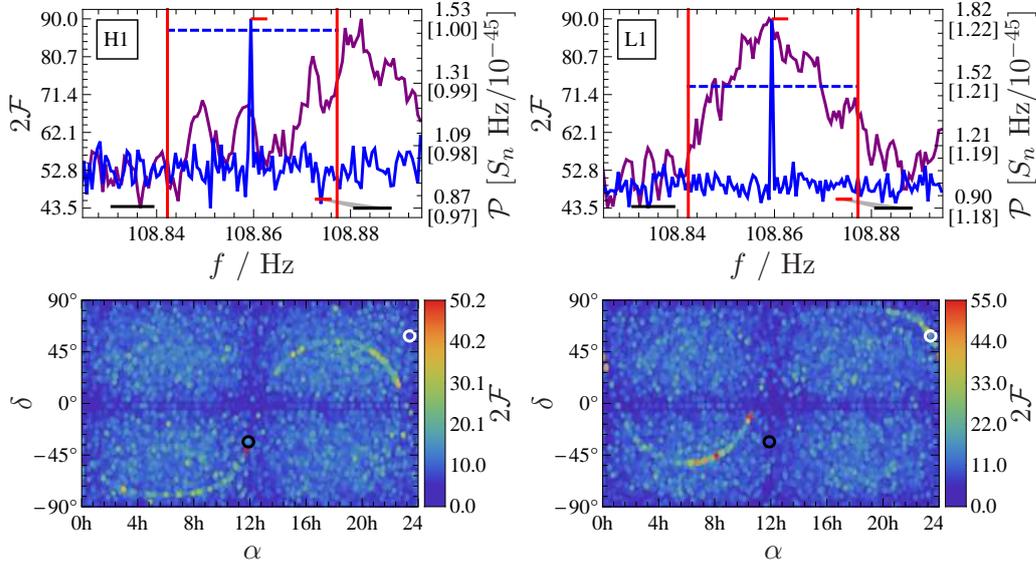


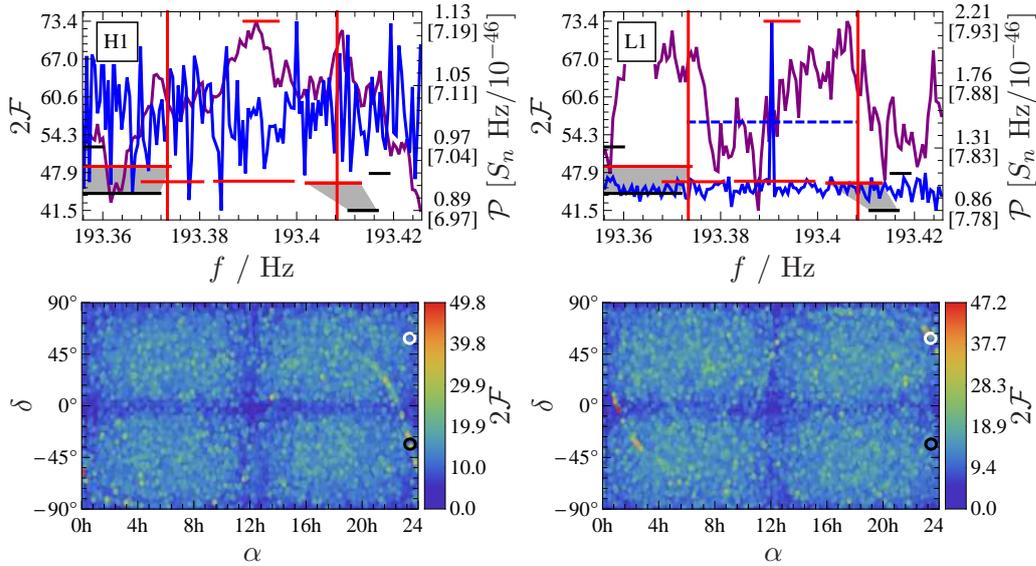
Figure 9.6: Times during which pulsar hardware injections were injected into L1 (green) and H1 (red) during the S5 “try3” epoch (Riles *et al.* 2008). The time span of the Cas A search data set is indicated in blue. Tick marks are the first 3 digits of GPS times. (Framed) A close-up of the injection times at the beginning of the Cas A search data set.

Table 9.3: Doppler parameters of the two pulsar hardware injections strong enough to be seen by the Cas A search, despite the limited overlap between the search data set and the injection times. *Nominal* indicates the nominal parameters of each injection (Riles *et al.* 2008), *Found* indicates the parameters of the injection found by the search, and the start and end times of the search data set are $t_0 = \text{GPS } 858459411$ and $t_1 = \text{GPS } 859495818$.

	α	δ	$\frac{f(t_0)}{\text{Hz}}$	$\frac{f(t_1)}{\text{Hz}}$	$\frac{\dot{f}(t_0)}{\text{Hz s}^{-1}}$	$\frac{\ddot{f}(t_0)}{\text{Hz s}^{-1}}$
Pulsar hardware injection no. 3						
Nominal	3.113	-0.584	108.857	108.857	-1.5×10^{-17}	0
Found	6.124	1.026	108.862	108.860	-2.5×10^{-9}	-6.2×10^{-19}
Pulsar hardware injection no. 8						
Nominal	6.133	-0.583	193.385	193.376	-8.7×10^{-9}	0
Found	6.124	1.026	193.396	193.389	-6.5×10^{-9}	1.5×10^{-17}



(a) Pulsar hardware injection no. 3



(b) Pulsar hardware injection no. 8

Figure 9.7: Veto bands triggered by pulsar hardware injections (a) no. 3, and (b) no. 8. In Figures (a) and (b): (top) Veto band plots, see section 9.3.3 for a description; (bottom) Sky map of $2\mathcal{F}$ from follow-up searches at the nominal injection frequency and spindown of each injection (left), and the frequency and spindown at which each injection was found by the search (right). The black circle is sky position of the respective injection; the white circle is the sky position of Cas A.

shift in frequency.

The Cas A search data set was not chosen with regard to when the hardware injections were active, and as a result only $\sim 4.2\%$ of the selected H1 data and $\sim 9.2\%$ of the selected L1 data contains the hardware injections. Since the expectation value of $2\mathcal{F}$ is linear in $\rho^2 \propto T_{\text{obs}}$, we expect the strength of the injections to be reduced by a factor of $\gtrsim 10$. Taking this reduction into account, injections no. 0 and 6 have respective strengths $h_0 \sim 2.5 \times 10^{-26}$ and $\sim 6.9 \times 10^{-26}$, and are thus too weak to be detected by the search performed here (see Figure 8.7). Injections no. 3 and 8, on the other hand, have strengths $h_0 \sim 1.63 \times 10^{-24}$ and $\sim 1.59 \times 10^{-24}$ respectively, and thus should be detectable; this is indeed the case.

Table 9.3 shows the nominal parameters of pulsar hardware injections no. 3 and 8 (Riles *et al.* 2008), the parameters at which the injections were found by the search. The frequencies at which the injections are found are slightly higher than the nominal frequencies. The very small \dot{f} of injection no. 3 is poorly recovered, but the nominal and recovered \dot{f} of injection no. 8 agree to within $\sim 25\%$. The hardware signals were injected with zero \ddot{f} ; the range of recovered \ddot{f} is consistent with the extent of the metric mismatch ellipses for the time span of the Cas A data set.

Indeed, the two injections are strong enough to have been detected by the StackSlide power $\mathcal{P}(f)$, and thus were flagged during the post-processing. Figure 9.7 shows veto band plots (top row of plots in Figures 9.7a and 9.7b) for the two hardware injections. Injection no. 3 is a visible spike $\mathcal{P}(f)$ in both H1 and L1, and exceeds the threshold \mathcal{P}_{thr} . Injection no. 8, however, is only seen in L1, although again $\mathcal{P}(f)$ exceeds \mathcal{P}_{thr} . It is not entirely clear why injection no. 8 is not seen in H1, but the following explanation seems possible. First, note that injection no. 8 possesses a high spindown, so that it moves in frequency by half a SFT frequency bin width in $(T_{\text{SFT}}/2)^{-1}/8.7 \times 10^{-9} \approx 8$ hours. In the framed close-up in Figure 9.6, we see that the total overlap of the Cas A data set with the H1 injection times consists of two segments separated by ~ 11 hours; in this intervening time, injection no. 8 will have moved more than half an SFT bin. Since $\mathcal{P}(f)$ is calculated from a single SFT bin $|\tilde{x}(f)|^2$, the injection signal in the two segments will no longer be summed together, and instead will effectively be split into two signals with reduced signal-to-noise ratios. In addition, as the total H1 injection is spread over less than a day, the modulation of the beam patterns F_+ and F_\times (see section 5.2.1) become important. It may be the case that the H1 injection is active at times when the interferometer is unfavourably oriented to the sky position of the signal, which implies that the beam patterns are small and, as a result, the signal-to-noise is reduced during these times (see equation 5.20). We do not, however, investigate this further.

To confirm that the spikes found in $\mathcal{P}(f)$ are indeed the pulsar hardware injections, we perform two searches, for each injection, of the Cas A data set over right ascension α and declination δ , using `ComputeFStatistic_v2`. We use the default sky position template bank generator of `ComputeFStatistic_v2` (i.e. not the template bank algorithm used in the search and presented in Chapter 7), with a large mismatch of 7.0 to reduce the number of data points for plotting purposes. The first search is performed at the nominal frequency and spindown parameters of each injection; the second search is at the frequency and spindowns at which each injection was found by the Cas A search.

The results are shown in the bottom rows of plots in Figures 9.7a and 9.7b; plotted are sky maps of $2\mathcal{F}$ as a function of α and δ . Circles in the sky (Prix & Itoh 2005) are distinguishable in all four sky maps. In the leftmost sky maps, which were produced at the nominal injection parameters, the circles pass through the sky position of the respective injection (black circle), showing that this signal is very probably the hardware injection. In the rightmost sky maps, which were conducted at the parameters of the injections, the circles pass through the Cas A sky position, explaining why these signals were seen by the search. Note that, in Figure 9.7a, the circles in the leftmost sky maps pass reasonably close to the Cas A sky position, and conversely the circles in the rightmost sky maps pass reasonably close to the sky position of the injections; the same is likely true in Figure 9.7b but is more difficult to see due to the fainter circles.

The displacement of the circles is consistent with the signals in the left and right sky maps being one and the same, and the movement being due to the different frequencies at which the signal is searched for. For each injection, and between the leftmost and rightmost sky maps, the circles moves upward in the sky map, which corresponds to a movement in the direction of the approximate orbital velocity vector of the Earth at the observation time; see Figure 1 of Prix & Itoh (2005). The positive sign of the accompanying change in frequencies (see Table 9.3) is consistent with the behaviour found by Prix & Itoh (see their Figure 2). In summary, we are confident that the veto bands of Figures 9.7a and 9.7b contain the pulsar hardware injections no. 3 and 8 respectively, and therefore we exclude the $2\mathcal{F}$ values in these bands (the loudest of which was $2\mathcal{F} \sim 90$) from the remainder of the search.

We are now left with nine veto bands. Three veto bands are readily identified with well-known instrumental lines (see Appendix 9.A.1). The six remaining veto bands were investigated by Nelson Christensen and colleagues, as part of an ongoing effort to identify and understand long-term instrumental coherence in the S5 data (Christensen *et al.* 2008). The coher-

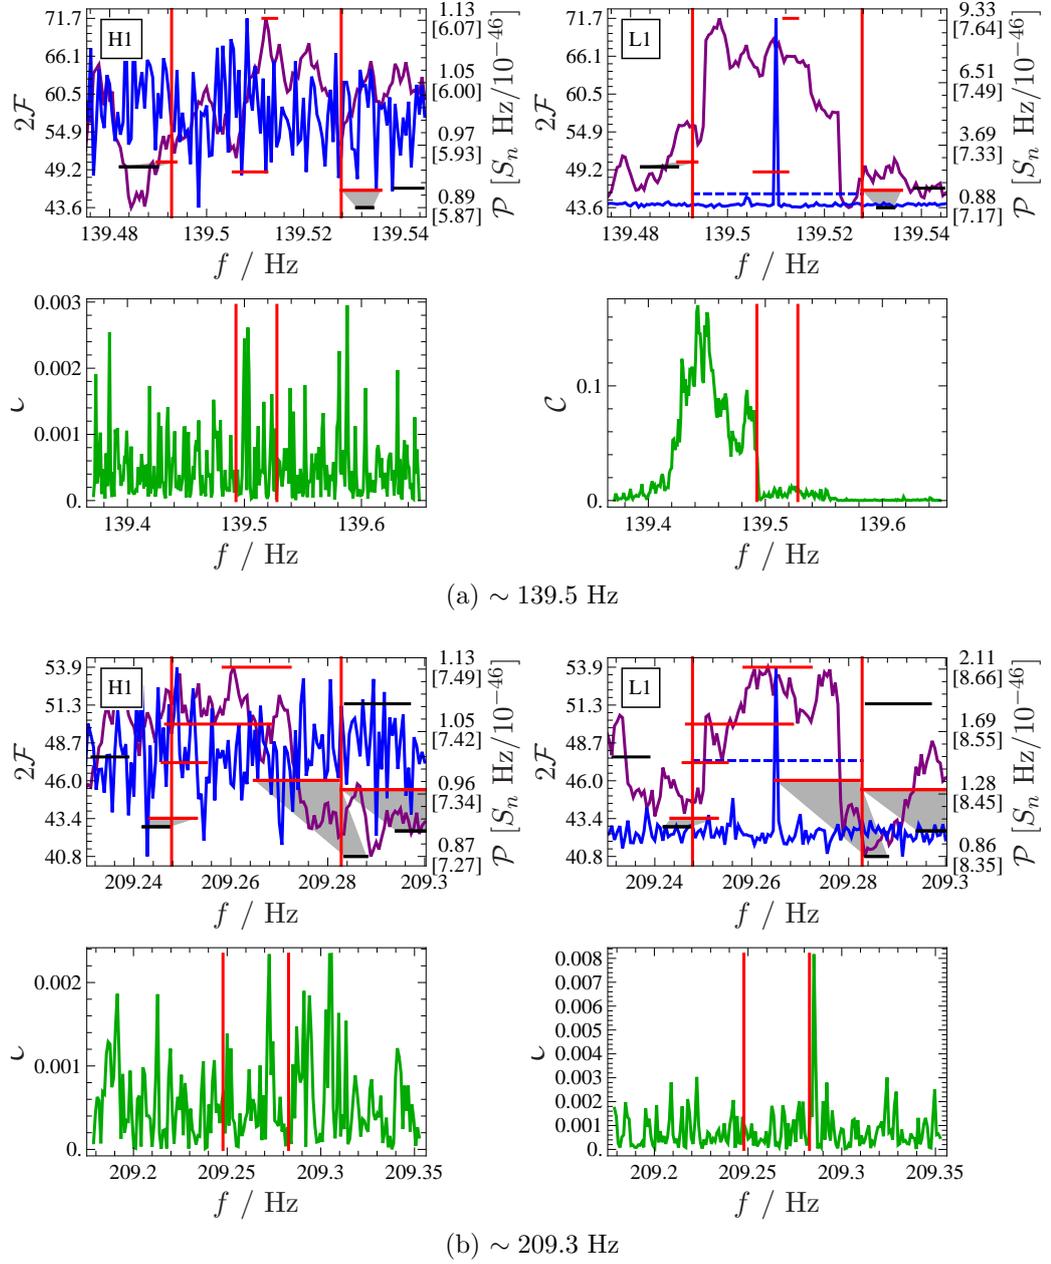


Figure 9.8: Veto bands triggered by probable instrumental lines in L1, at (a) ~ 139.5 Hz, and (b) ~ 209.3 Hz. In Figures (a) and (b): (top) Veto band plots, see section 9.3.3 for a description. In Figure (a): Coherence between the channels L1:LSC-DARM_ERR and L0:PEM-RADIO_LVEA during (bottom left) March 2007, and (bottom right) May 2007. In Figure (b): Coherence between the channels L1:LSC-DARM_ERR and L0:PEM-LVEA_MAGZ during (bottom left) March 2007, and (bottom right) October 2006.

ence

$$\mathcal{C}[x, y](f) = \left\langle \frac{|\tilde{x}(f)^* \tilde{y}(f)|^2}{|\tilde{x}(f)|^2 |\tilde{y}(f)|^2} \right\rangle \quad (9.7)$$

is computed between two data channels x and y , where $\langle \rangle$ implies averaging of many suitably-sized blocks of data. One channel is always `L1:LSC-DARM_ERR`, which is the un-calibrated gravitational wave channel of the L1 detector, and the other channel is chosen to be a physical environmental monitoring (PEM) channel. PEM channels that were found to be of interest are (LSC 2009):

- `L0:PEM-BSC4_MIC`, a microphone on Beam Splitter Chamber #4;
- `L0:PEM-EX_MAGZ`, the Z axis of a magnetometer at the end station of the X arm;
- `L0:PEM-ISCT4_ACCZ`, the Z axis of an accelerometer on Instrument Sensing and Control Table #4;
- `L0:PEM-LVEA_MAGZ`, the Z axis of a magnetometer in the Laser Vacuum Equipment Area; and
- `L0:PEM-RADIO_LVEA` a wire in the Laser Vacuum Equipment Area used to look for radio signals.

Of the six veto bands, four are definitely identified with instrumental noise (see Appendix 9.A.1). The case is not, however, as strong for the remaining two veto bands (shown in Figure 9.8) as there is no significant coherence between `L1:LSC-DARM_ERR` and any PEM channels during the month containing the Cas A search data set (March 2007). Coherence is seen, however, during other months between `L1:LSC-DARM_ERR` and a number of PEM channels; the coherence between `L1:LSC-DARM_ERR` and representative channels are plotted in Figure 9.8. A transitory or wandering noise source may be responsible. Nonetheless, there are additional reasons for excluding these two veto bands. In both bands, $\mathcal{P}(f)$ exceeds the threshold \mathcal{P}_{thr} only in L1. In particular, the $\mathcal{P}(f)$ spike at ~ 139.5 Hz exceeds the threshold by a factor of ~ 6 , or ~ 60 standard deviations of \mathcal{P} ; it is very unlikely that such a strong signal, if it were a gravitational wave, would be visible so clearly only in one detector. Finally, we note that the largest values of $2\mathcal{F}$ in each veto band, ~ 71 (Figure 9.8a) and ~ 54 (Figure 9.8b) are not particularly significant candidates, considering the distribution of the largest value of $2\mathcal{F}$ for the entire search (see Figure 9.11). Therefore, we also exclude the $2\mathcal{F}$ values in these bands from the remainder of the search.

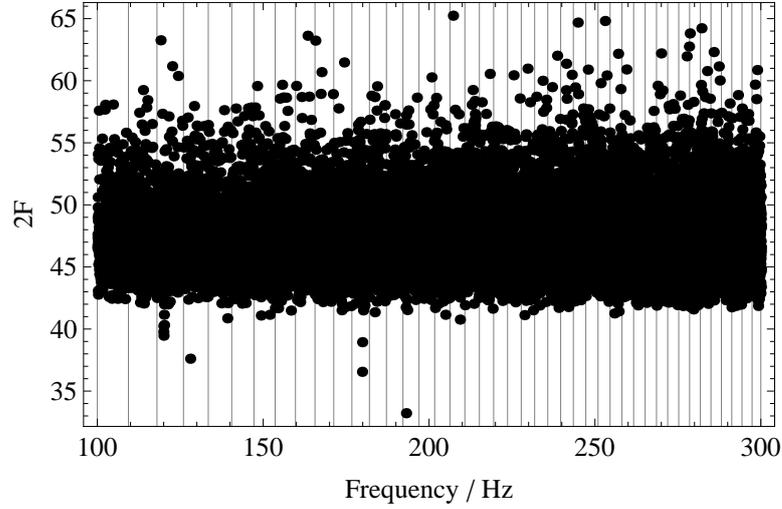


Figure 9.9: Largest $2\mathcal{F}$ found in each search job frequency band, after the post-processing described in section 9.3. The thin vertical lines mark the divisions between groups of 500 search job frequency bands.

9.4 The largest $2\mathcal{F}$

Figure 9.9 shows the largest value of $2\mathcal{F}$ found in each of the Cas A search job frequency bands, after the removal of spuriously large $2\mathcal{F}$ values detailed in section 9.3. In contrast to Figure 9.3, the post-processed results show no obvious outliers. The largest non-vetoed value is $2\mathcal{F}^* \approx 65.2$. Figure 9.10 show plots, in the style of the veto band plots in section 9.3.3, of the job band containing $2\mathcal{F}^*$. No significant spikes are seen in the StackSlide power $\mathcal{P}(f)$, and the range of $\mathcal{P}(f)$ is consistent with its expected Gaussian noise

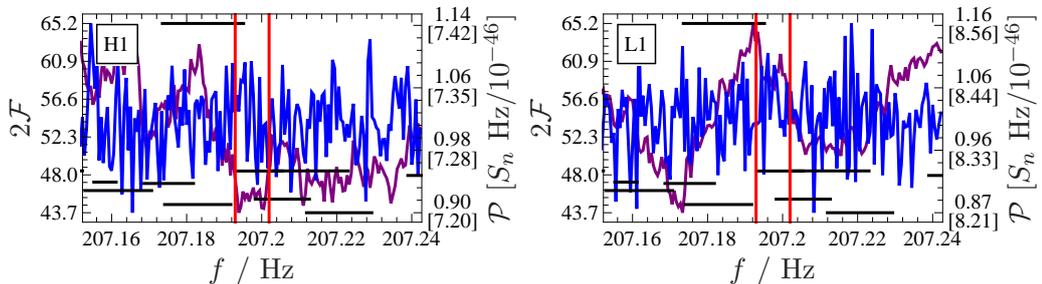


Figure 9.10: The search job band containing the largest non-vetoed value of $2\mathcal{F}$. The form of the plots are as described in section 9.3.3, except that the vertical red lines now indicate the extent of the search job band. The value of the largest $2\mathcal{F}$ is ~ 65.2 .

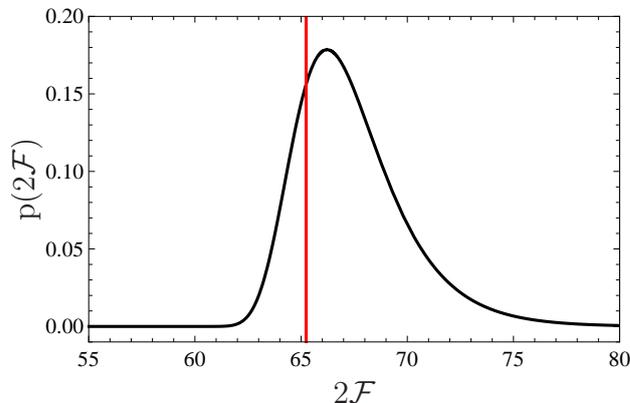


Figure 9.11: The largest value of $2\mathcal{F}$ found by the Cas A search, $2\mathcal{F}^*$ (red), and its expected distribution (black), assuming no signal and independent templates.

distribution (see Figures 9.5). We conclude that $2\mathcal{F}^*$ is not due to an instrumental line or other non-Gaussian noise, and is therefore the best candidate for a detection of gravitational waves from Cas A produced by this search.

Figure 9.11 plots $2\mathcal{F}^*$ against its expected distribution, assuming that no signal is found. As discussed in section 8.7, we also assume that the searched templates are statistically independent, and so the distribution in Figure 9.11 is calculated using the full number of templates. It is evident that the $2\mathcal{F}^*$ found by this search is consistent with the hypothesis that no signal has been detected; its value would need to be much greater than the spread of the distribution in order for the possibility of a detection to be seriously considered.

We conclude, therefore, that the results of the search presented here show no evidence of a detected periodic gravitational wave signal from the central compact object in Cas A.

9.5 Upper limits

Having concluded that we have not detected a gravitational waves from Cas A, we proceed to use the results of the search to set 95% confidence upper limits on the strength of gravitational waves from Cas A. The full search frequency band (100–300 Hz) is partitioned into 400 *upper limit bands*, each of width 0.5 Hz; upper limits are set individually in each upper limit bands. The width of the bands is chosen to be small enough so that the noise floor of the detectors is roughly constant over the band, and large enough to limit the

number of bands for computational reasons. No upper limits are set within 2 Hz either side of harmonics of 60 Hz (i.e. within 118–122 Hz, 238–242 Hz, and 298–300 Hz), due to strong non-Gaussian contamination from the 60 Hz mains power harmonics.

In previous periodic gravitational wave searches (e.g. Abbott *et al.* 2007b, 2008c), the 95% confidence upper limit of the gravitational wave strain, $h_0^{95\%}$, in each upper limit band was determined using software injections, as follows. First, `Makefakedata_v4` (see section 7.4.1) is used to generate SFTs, with time stamps matching the time stamps of the searched SFTs, and frequencies limited to the upper limit band under consideration. The generated SFTs contain simulated Gaussian noise, of the same amplitude as the noise in the searched SFTs, and a simulated periodic gravitational wave signal with fixed amplitude, h_0 (whose value is initially guessed), and randomly chosen inclination angles $\cos \iota$, polarisation angles ψ , and initial phases ϕ_0 . The signal is injected into the generated SFTs at a random frequency, and with random first (and second, if applicable) spindowns drawn from the original parameter space of the search. The generated SFTs are then searched, using the same software (e.g. `ComputeFStatistic_v2`) and parameter space as the original search, in order to find the injected signal. For the sake of computational efficiency, only a small subset of the original parameter space is searched, surrounding the (known) frequency and spindown of the injected signal. Care must be taken, however, to ensure that there remains a random mismatch between the injected signal and the searched template closest to it. The maximum value of $2\mathcal{F}$ found by the search for the injection is compared to $2\mathcal{F}_{\text{ul}}$, the largest value of $2\mathcal{F}$ found in the upper limit band under consideration. If $2\mathcal{F} > 2\mathcal{F}_{\text{ul}}$, the injection is considered to have been detected. The above procedure is performed many times. The number of detected injections, as a fraction of the total, is the confidence $C(h_0)$ at h_0 . This confidence is then compared to the desired confidence $C = 95\%$. If $C(h_0) < C$, the injected amplitude h_0 is too small; if $C(h_0) > C$, h_0 is too large. In either case, the value of h_0 is adjusted, and the entire procedure repeated, until finally the $C(h_0) = C$, whereupon $h_0 = h_0^{95\%}$ is the desired 95% confidence upper limit on strain.

9.5.1 Analytic estimation of upper limits

To expedite the derivation of the upper limits, we instead use an analytic model of the software injection procedure, described above, to quickly arrive at an estimate of $h_0^{95\%}$. The strain upper limit on h_0 at confidence C is

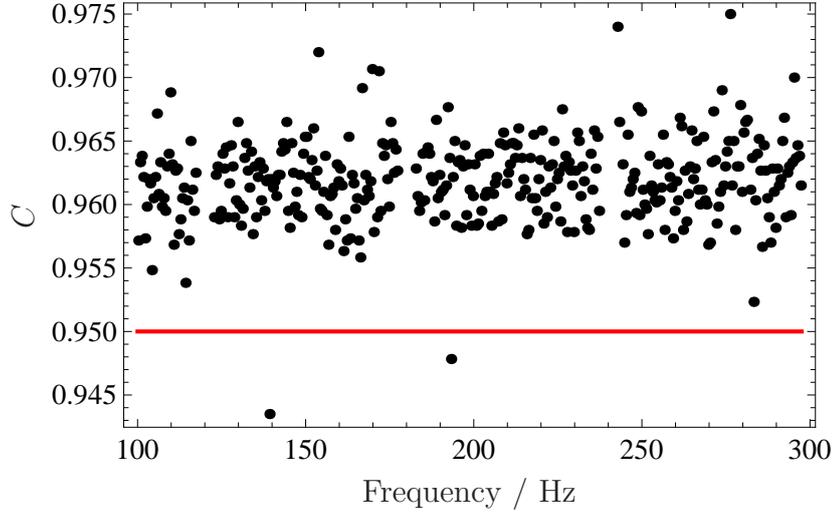


Figure 9.12: Confidence of the upper limits in each upper limit band (black points), as determined by software injections. The target confidence of 95% is plotted in red.

modelled by the following equation:

$$1 - C = \int_0^{2\mathcal{F}_{\text{nl}}} d(2\mathcal{F}) \int_{-1}^1 d(\cos \iota) \int_{-\frac{\pi}{4}}^{\frac{\pi}{4}} d\psi \int_0^{\mu_{\text{max}}} d\mu \quad \text{p}\left(\chi_4^2[(1 - \mu)\rho^2]; 2\mathcal{F}\right) \text{p}(\cos \iota) \text{p}(\psi) \text{p}(\mu), \quad (9.8)$$

where μ is the mismatch, and $\text{p}(\cos \iota)$, $\text{p}(\psi)$, and $\text{p}(\mu)$ are the probabilities of particular values of $\cos \iota$, ψ , and μ respectively. The optimal signal-to-noise ratio is most conveniently given by (e.g. Prix 2007a)

$$\rho^2(h_0^C) = A(A_1^2 + A_3^2) + B(A_2^2 + A_4^2) + C(2A_1A_2 + 2A_3A_4), \quad (9.9)$$

where the A , B , and C coefficients (equation 5.43) include the noise spectral density S_n , and the A_i parameters (equations 5.39) are functions of h_0^C .

Equation 9.8 is very similar to equation 8.15, for the false dismissal rate, except that we explicitly account for the loss in signal-to-noise ratio due to the mismatch between a software-injected signal and the nearest search template. The quantities $\cos \iota$ and ψ are uniformly distributed, i.e. $\text{p}(\cos \iota)$ and $\text{p}(\psi)$ are constants. The probability distribution $\text{p}(\mu)$ is given by the properties of the template bank; here, we use the template bank algorithm of Chapter 7, and $\text{p}(\mu)$ is identical to the distributions plotted in Figures 7.4a.

We solve equation 9.8 for h_0^C numerically. Its right hand side is evaluated using the technique of Monte Carlo integration, which asserts that

$$\int_{\mathcal{S}} d(x, y, \dots) f(x, y, \dots) \approx \langle f(\text{random } x, y, \dots) \rangle \times \text{volume}(\mathcal{S}), \quad (9.10)$$

where $\langle \rangle$ imply averaging. The strain h_0 is initially guessed to be an *ad hoc* combination of the A , B , and C coefficients, and is updated using the Newton-Raphson root-finding method:

$$h_0^{\text{new}} = h_0^{\text{old}} - f(h_0^{\text{old}}) \left[\frac{\partial f}{\partial h_0}(h_0^{\text{old}}) \right]^{-1}, \quad (9.11)$$

where

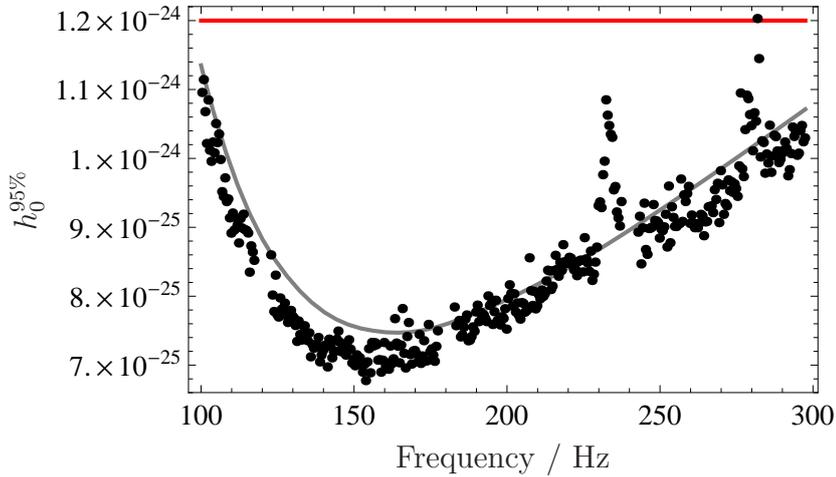
$$f(h_0) = \text{RHS of equation 9.8} - \text{LHS of equation 9.8}. \quad (9.12)$$

When $f(h_0)$ is sufficiently close to zero, the upper limit h_0^C has been found. This procedure has been implemented by the author in LALApps, in the file `ComputeFStatAnalyticMonteCarloUpperLimit.c` under `src/pulsar/FDS_isolated`.

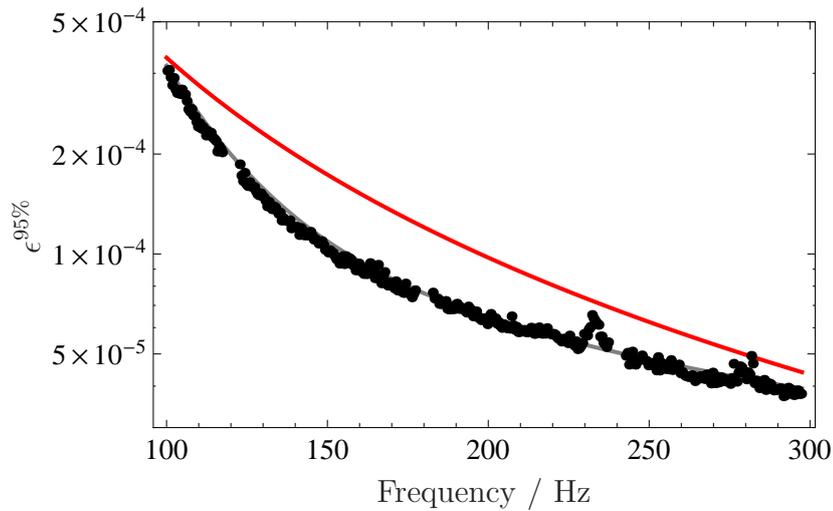
Once an upper limit $h_0^{95\%}$ is arrived at for each upper limit band, one round of software injections (described in section 9.5) are then used to confirm that the estimated $h_0^{95\%}$ is correct. Figure 9.12 plots the confidence $C(h_0)$ for each upper limit band, as determined by the software injections. With a few exceptions (corresponding to noisy bands) the confidences of the upper limits are higher than the targeted 95%. This implies that, according to the software injections, the upper limits arrived at by solving equation 9.8 understate their true confidence, and are therefore conservative estimates of the strain upper limit.

9.5.2 Upper limits on h_0 and ϵ

Figure 9.13 presents the 95% confidence upper limits on the strain h_0 of gravitational waves from the neutron star in Cas A, and on the star's ellipticity ϵ . As stated above, no upper limits are set within 2 Hz of the 60 Hz harmonics. The indirect upper limits, and the predicted observable upper limits (see Chapter 8) are plotted for comparison; the observational upper limits beat the indirect limits, as expected. The upper limits on h_0 derived from the search perform slightly better than was predicted; for example, at ~ 150 Hz, the search places an upper limits of $\sim 7 \times 10^{-25}$, compared to the predicted $\sim 7.5 \times 10^{-25}$. This small improvement is likely due to improved



(a) Upper limit on strain



(b) Upper limit on ellipticity

Figure 9.13: Upper limits at 95% confidence derived from the search (black points), the indirect limits derived in section 8.2 (red lines), and the predicted upper limits of section 8.6 (grey lines), on (a) the strain of gravitational waves from Cas A $h_0^{95\%}$, and (b) the ellipticity of Cas A $\epsilon^{95\%}$.

performance of the LIGO detectors, relative to the design sensitivity curve used in the predictions. The deterioration in the upper limits between 230 and 240 Hz, and at 280 Hz, are due to excessive broadband noise in H1 (see Figure 9.2). The noise in these bands has been investigated and their sources identified.³

The upper limits on ellipticity range from approximately 4×10^{-5} to 4×10^{-4} over the frequencies searched. Various models of quark or hybrid stars, combined with the increased breaking index of Horowitz & Kadau (2009), predict maximum ϵ within this range, as reviewed in section 2.4. This search would therefore have been able to detect a maximally deformed neutron star in Cas A whose equation of state was at the speculative end of theoretical possibilities. Since, however, this search has not detected any convincing gravitational wave signal, we can make no statement regard the internal physics of the Cas A CCO, since its ellipticity may simply be much less than maximum.

Abbott *et al.* (2010) has recently set upper limits on known pulsars using the entire LIGO S5 data set. New strain upper limits were set on the Crab pulsar of at best $\sim 2 \times 10^{-25}$ (depending on how the data is analysed to account for a timing glitch in the pulsar), which is 13–15% of the Crab indirect upper limit. Abbott *et al.* (2010) also set an upper limit on the pulsar PSR J0537-6910 which is very close to its indirect limit. The PowerFlux search of the early S5 data (Abbott *et al.* 2009b) set all-sky upper limits below 10^{-24} over the Cas A search band, and a first search of the S5 data using Einstein@Home (Abbott *et al.* 2009d) set an all-sky upper limit of 3×10^{-24} between 125 and 225 Hz, for 90% confidence.

The observed strain upper limit on the Crab, $\sim 2 \times 10^{-25}$ (Abbott *et al.* 2010), is ~ 3.5 times lower than our best upper limit on Cas A, at ~ 150 Hz. It was obtained using ~ 50 times more data from H1 and L1 than was used in the Cas A search; the Crab search also used a similar amount of data from the Hanford 2-km interferometer H2. The gravitational wave frequency of the Crab (assumed to be twice the rotation frequency) is very near to 60 Hz; at that frequency, the strain noise in the LIGO detectors is, very roughly, an order of magnitude worse than at 150 Hz (Abbott *et al.* 2009e); thus, the ratio S_n/T_{obs} is about a factor of 5 larger for Cas A than for the Crab. To compare the upper limits, we take the ratio S_n/T_{obs} to be unity for Cas A, and assume that the sensitivity factor for Cas A (~ 35) is roughly 3 times the sensitivity factor for the Crab search; the \mathcal{F} statistic sensitivity factor for a single template search is ~ 11 (Abbott *et al.* 2007b), although Abbott *et al.* (2010) uses a different search method. We then average the ratio S_n/T_{obs}

³ N. Christensen and R. Schofield, private communications.

over the detectors used in each search (H1 and L1 for Cas A; H1, L1, and H2 for the Crab); the upper limit is then proportional to the sensitivity factor times the square root of this average (see equation 8.13). The ratio of the upper limits on Cas A (at 150 Hz) and the Crab is then

$$\frac{\text{Cas A}}{\text{Crab}} \approx \frac{35\sqrt{(1+1)/2}}{11\sqrt{1/5 \times (1+1+2)/3}} \approx 6; \quad (9.13)$$

we assume the noise strain in H2 is twice that of H1 and L1. This ratio is in very rough agreement with the ratio of the actual upper limits (~ 3.5); thus, our upper limits on Cas A are at least consistent, to within an order of magnitude, with those of other periodic gravitational wave searches, given the length and sensitivity of the data analysed.

Because Cas A is a young neutron star, it is worth considering interpreting the upper limits obtained here in terms of r -mode emission (see section 2.4). The r -mode emission mechanism is different in a number of ways to the nonaxisymmetric deformation mechanism targeted by this search, and these differences are not explicitly reflected in the implementation of the \mathcal{F} statistic. Nevertheless, it may be possible to account for these differences after the fact, and in so doing derive the first observational upper limits on gravitational waves from r modes. Further details are presented in Abadie *et al.* (2010).

9.6 Summary

This chapter presented the results of the search for gravitational waves from Cas A proposed in Chapter 8. We selected the search data set from the LIGO S5 run, performed the search, and detailed the post-processing of the results, including the vetoing of some results due to contamination by instrumental noise. We presented the largest value of the F statistic, $2\mathcal{F}^* \approx 65.2$, found by the search, and concluded from it that we had not detected a gravitational wave signal from Cas A. Finally, we presented observational upper limits for Cas A on the gravitational wave strain and ellipticity, which beat their respective indirect upper limits.

9.A Additional material

This appendix includes additional material pertaining to this chapter.

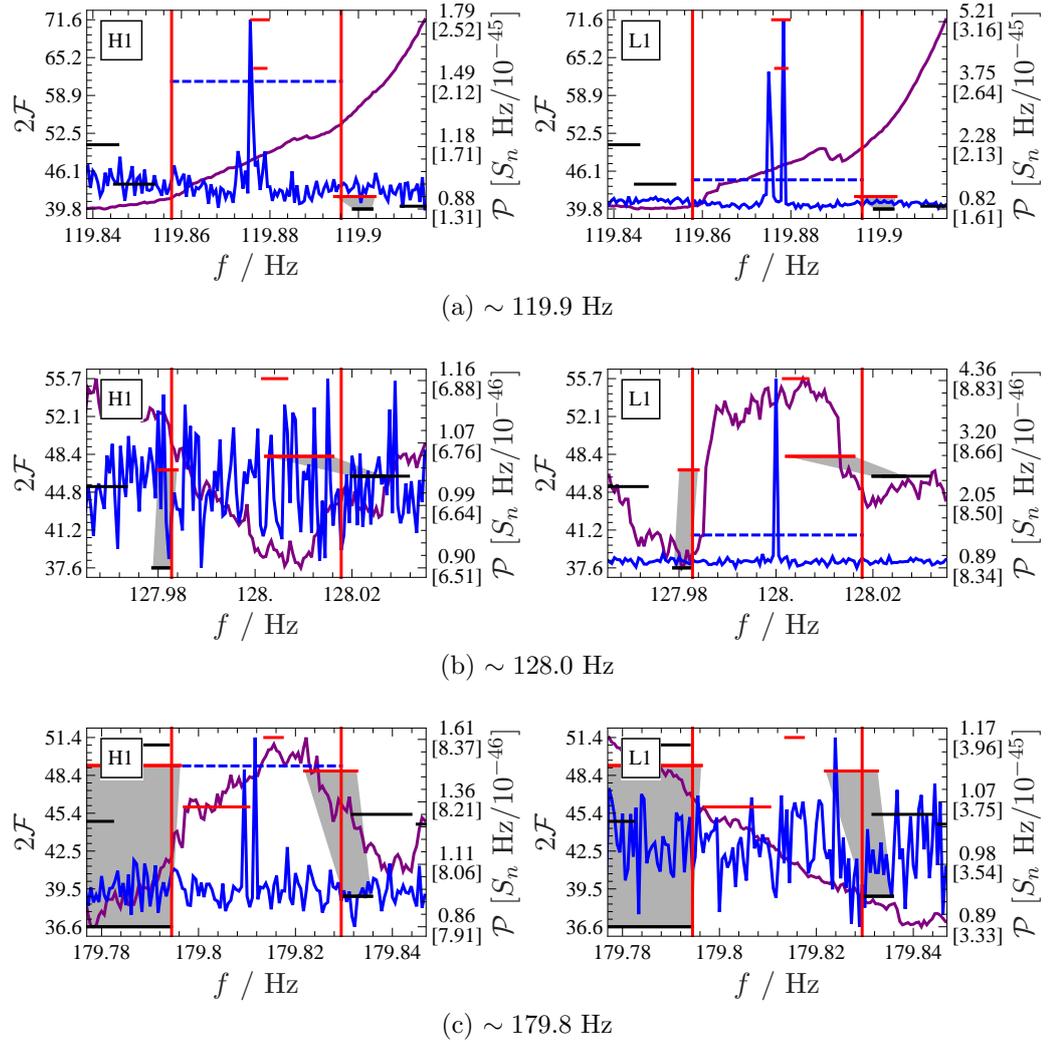


Figure 9.14: Veto bands triggered by instrumental lines at (a) ~ 119.9 Hz, (b) ~ 128.0 Hz, and (c) ~ 179.8 Hz. Figures (a) and (c) are associated with harmonics of 60 Hz, while Figure (b) is a harmonic of 16 Hz.

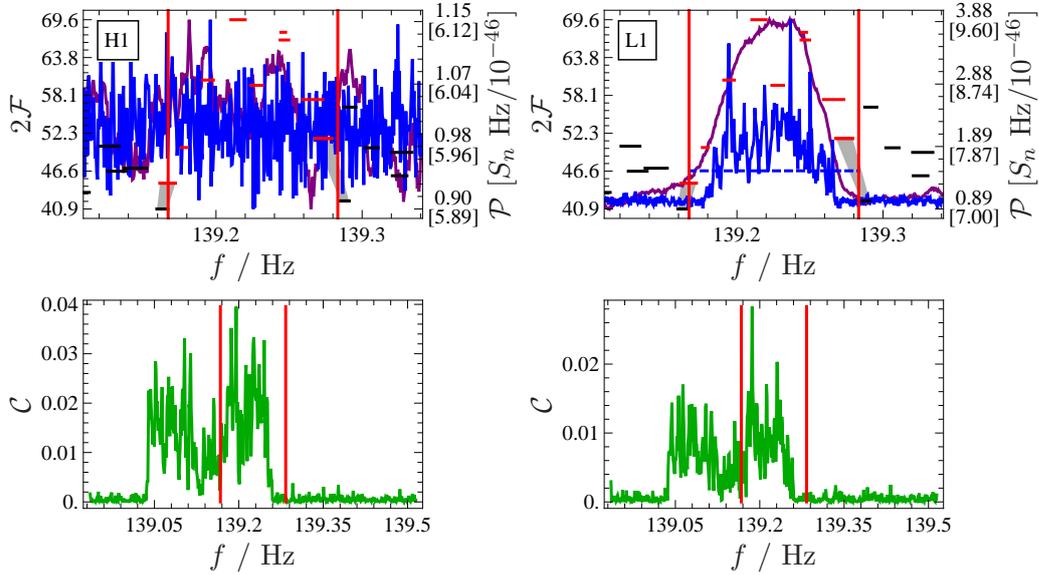


Figure 9.15: Veto band triggered by an instrumental L1 line at ~ 139.2 Hz. (top) Veto band plots, see section 9.3.3. (bottom) Coherence during March 2007 between L1:LSC-DARM_ERR and: (left) L0:PEM-ISCT4_ACCZ, and (right) L0:PEM-LVEA_MAGZ.

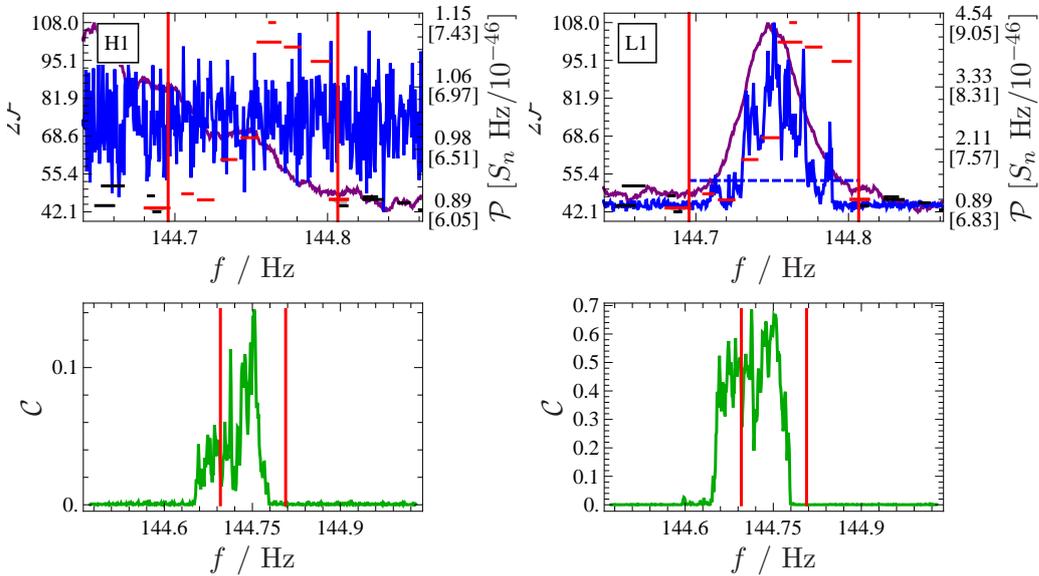


Figure 9.16: Veto band triggered by an instrumental L1 line at ~ 144.8 Hz. (top) Veto band plots, see section 9.3.3. (bottom) Coherence during March 2007 between L1:LSC-DARM_ERR and: (left) L0:PEM-BSC4_MIC, and (right) L0:PEM-EX_MAGZ.

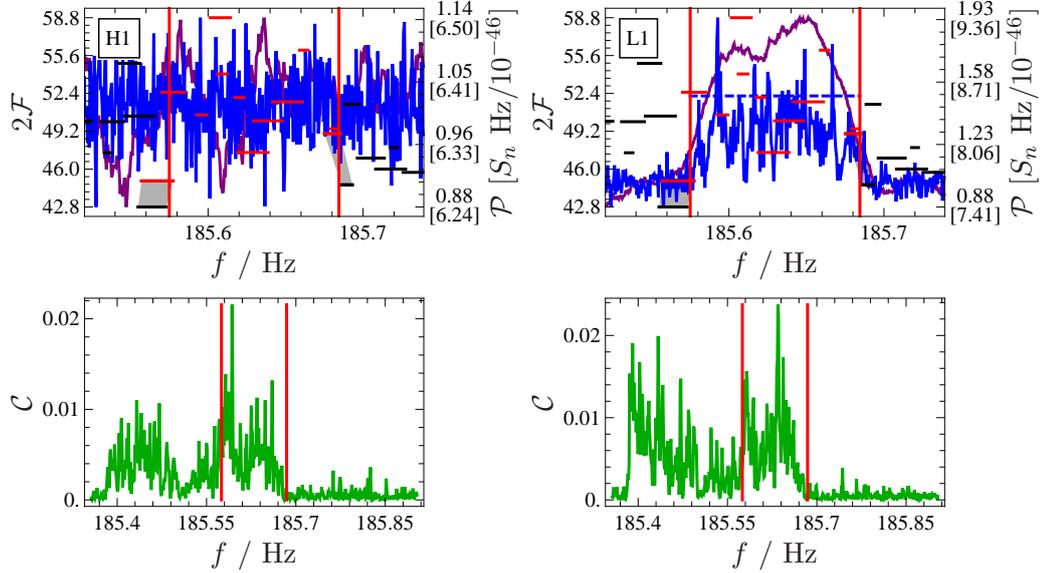


Figure 9.17: Veto band triggered by an instrumental L1 line at ~ 185.6 Hz. (top) Veto band plots, see section 9.3.3. (bottom) Coherence during March 2007 between L1:LSC-DARM_ERR and: (left) L0:PEM-LVEA_MAGZ, and (right) L0:PEM-RADIO_LVEA.

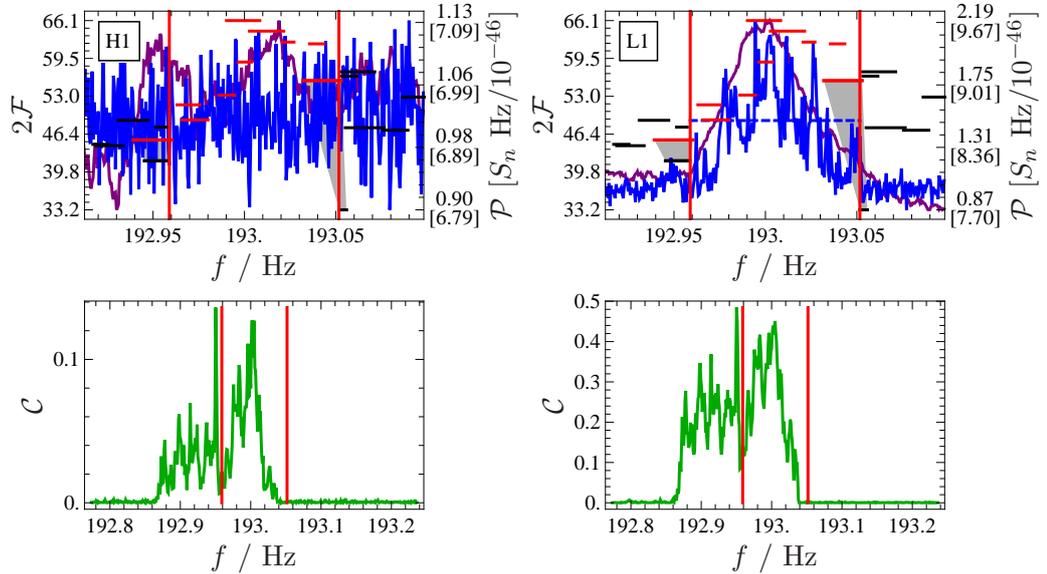


Figure 9.18: Veto band triggered by an instrumental L1 line at ~ 193.0 Hz. (top) Veto band plots, see section 9.3.3. (bottom) Coherence during March 2007 between L1:LSC-DARM_ERR and: (left) L0:PEM-BSC4_MIC, and (right) L0:PEM-EX_MAGZ.

9.A.1 Known instrumental lines

Figures 9.14–9.18 present veto band triggered by seven instrumental lines. The lines in Figure 9.14 are well-known harmonics of the 60 Hz AC mains power and of the 16 Hz data acquisition frequency (Abbott *et al.* 2004a). The lines in Figures 9.15–9.18 were found by Nelson Christensen and colleagues; see section 9.3.3 for further information. In each veto band, coherence (equation 9.7) is clearly seen between L1:LSC-DARM_ERR and a PEM channel. The physical origins of the coherent noise were not investigated further.

Chapter 10

Conclusion

This thesis considered a number of problems relevant to the physics of neutron stars and gravitational waves. We performed simulations of magnetic mountains on accreting neutron stars, explored methods for reducing the computational cost of periodic gravitational wave searches, and presented a search for periodic gravitational waves targeting the supernova remnant Cassiopeia A. Below, the achievements of this thesis are reviewed, and recommendations for further work are proposed.

Chapter 2 introduced the physics of neutron stars and gravitational waves, and an important link between them. In Chapters 3–4, we studied the problem of the burial of the magnetic field of accreting neutron stars, and the formation of magnetically confined mountains. In Chapter 3, we reviewed previous work on magnetic burial; we then presented a numerical procedure capable of building magnetic mountains with realistic masses. We justified the injection of the accreted matter from below, which is (with some subtleties) equivalent to injection from above in ideal magnetohydrodynamics. In Chapter 4, we presented the results of simulations of magnetically confined mountains, with masses up to $\sim 0.1M_{\odot}$, grown on a hard surface, as well as sinking into a soft fluid base. We presented an illustrative example of the growth of a mountain, and compared in detail the final configurations of hard- and soft-surface mountains. We found that the ellipticity of a mountain grown on a hard surface approaches $\sim 2 \times 10^{-4}$ for accreted masses greater than $\sim 10^{-3}M_{\odot}$, and that sinking reduces the ellipticity by between 25% and 60%. We compared our simulations to the work of Choudhuri & Konar (2002), and discussed the consequences for gravitational waves from low-mass x-ray binaries.

This work should be seen as one step in a progression of refinements to our numerical model of magnetic burial, which have included the self-consistent calculation of the mass-flux distribution (Payne & Melatos 2004), the hydro-

magnetic stability of the equilibrium state in axisymmetric (Payne & Melatos 2007) and three-dimensional (Vigelius & Melatos 2008) simulations, and the addition of resistive relaxation (Vigelius & Melatos 2009b). The next key refinement would be to incorporate a more realistic equation of state; the assumption of an isothermal equation of state breaks down during the later stages of accretion, where $M_a \gtrsim 10^{-3} M_\odot$ (Vigelius & Melatos 2009a). The equation of state also determines the equilibrium density profile of the star, which is important in determining to what depth the mountain sinks before stabilising. The simulations performed here suggest that the sinking of the mountain, while substantial, does not obliterate it; this is encouraging for the prospect of detecting gravitational waves generated by magnetic mountains.

Chapter 5 summarised results from the searches for gravitational waves conducted to date, and reviewed the analysis of the periodic gravitational waves expected from spinning neutron stars. In Chapter 6, we generalised the PowerFlux semi-coherent data analysis method to estimate the amplitude and polarisation parameters of periodic gravitational wave signals. We used simulated signals injected into Gaussian noise to compare the parameter estimation and detection efficiencies of the generalised PowerFlux methods we obtained against the standard PowerFlux methods. We found that, in general, the standard circular PowerFlux method is the most efficient; the relative performance of the remaining methods is dependent on the declination of the injected signals.

While we were unsuccessful in inventing a more efficient semi-coherent search method, it is good to know that the most efficient variety of PowerFlux (at least from within the methods examined here) is that already being used in all-sky searches of LIGO data. It would be interesting to understand more fully why circular PowerFlux is the most efficient, and under what conditions. The PowerFlux maximum likelihood statistic, mentioned in Mendell & Wette (2008), is another intriguing possibility for further research.

In Chapter 7, we presented a template bank generation algorithm for coherent periodic gravitational wave searches. We demonstrated how a sphere covering on an optimally thin lattice, together with a metric on the search parameter space, is used to construct a template bank with a minimum number of points. We also show how to place templates along the edges of the parameter space to ensure complete coverage, and how to estimate the total number of required templates. The performance of an implementation of the algorithm was successfully tested. The algorithm was utilised in the search for gravitational waves from Cassiopeia A presented in Chapters 8–9.

There are a number of improvements that could be made to the algorithm. For instance, we refer to the slice through a three-dimensional tiling over f , \dot{f} , and \ddot{f} shown in Figure 8.3. We see that the extent of the param-

eter space in the \ddot{f} direction is small compared to the bounding box of the mismatch ellipse of a single template; it seems likely, therefore, that in this instance the tiling in \ddot{f} could be compressed to a single template. Templates laid over a two-dimensional f - \dot{f} space would still extend into the \ddot{f} dimension, and the mutual overlap of their mismatch ellipses could be sufficient to give complete coverage of the \ddot{f} dimension. The maximum size of the \ddot{f} dimension for which this would be possible would depend on the geometry of the sphere covering. It would also be interesting to extend the algorithm to cover non-flat parameter spaces, such as the parameter space of sky positions for periodic gravitational wave signals. A coarse covering could be used to divide the parameter space into patches where the metric can be considered locally flat; each patch would then be covered by a lattice-based template bank. It remains to be seen whether such an algorithm would provide a more efficient coverage than existing methods.

In Chapters 8–9, we presented a search for periodic gravitational waves targeting the central compact object in the supernova remnant Cassiopeia A. In Chapter 8 we reviewed the wealth of astronomical observations of the remnant, and motivated a gravitational wave search targeting the young neutron star at its centre. We presented the proposed search of 12 days of LIGO S5 data using the `ComputeFStatistic_v2` implementation of the \mathcal{F} statistic, estimated the sensitivity of the search, and demonstrated that the search would beat an indirect upper limit on gravitational waves based on energy conservation. We also predicted the distribution of the largest value of the \mathcal{F} statistic found by the search assuming a signal is not detected. In Chapter 9 we presented the implementation of the search: from the selection of data from the S5 run, and the division of the search into multiple jobs for distribution over a computer cluster, to the post-processing of the results, including the removal of some values of the \mathcal{F} statistic due to contamination by instrumental noise. The largest value of the \mathcal{F} statistic found by the search, after post-processing, was consistent with the expected distribution assuming no signal; we conclude, therefore, that we have not detected gravitational waves from Cassiopeia A. Finally, we present the upper limits on gravitational waves derived from the search, which beat the indirect limit as expected.

Cassiopeia A now belongs to a select collection of astronomical objects (including, notably, the Crab pulsar) for which this has been achieved. Even with its initial design configuration, LIGO is already producing scientifically interesting results; this bodes well for the gravitational wave astronomy expected to be achievable by advanced interferometers. In addition, there are a number of possible improvements to the Cassiopeia A search. First, the coherent \mathcal{F} statistic method in this search is not the most sensitive possi-

ble; a hierarchical search combining coherent and semi-coherent stages would provide a significant improvement in sensitivity. Second, more sensitive data is now available: in July 2009, LIGO began its sixth science run (S6), featuring instrumental improvements (many of them brought forward from the advanced LIGO design), which are expected to yield a twofold improvement in strain sensitivity (LIGO Laboratory 2009c). Third, more computational power may be utilised; in particular, Einstein@Home would be an interesting platform on which to consider conducting a future search targeting Cassiopeia A. The methodology developed for the Cassiopeia A search may also be applied to other promising targets, such as the supernova remnant Vela Junior.

Advanced LIGO plans to begin observations in 2015 (LIGO Laboratory 2009a). This date is closely followed by two interesting anniversaries: the one-hundredth anniversary of general relativity (Einstein 1916), and the fiftieth anniversary of the discovery of pulsars (Hewish *et al.* 1968). By then, we may very well possess new, and potentially revolutionary, insights into the unknown physics of neutron stars, and a complete validation of the general theory of relativity. The next few years promise to be interesting times for both fields.

Bibliography

- J. Abadie *et al.* (LIGO Scientific Collaboration), 2010. First search for gravitational waves from the youngest known neutron star. *The Astrophysical Journal* in press.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2004a. Analysis of first LIGO science data for stochastic gravitational waves. *Physical Review D* 69(12), 122004. <http://dx.doi.org/10.1103/PhysRevD.69.122004>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2004b. 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. <http://dx.doi.org/10.1103/PhysRevD.69.082004>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2005a. First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform. *Physical Review D* 72(10), 102004. <http://dx.doi.org/10.1103/PhysRevD.72.102004>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2005b. Limits on Gravitational-Wave Emission from Selected Pulsars Using LIGO Data. *Physical Review Letters* 94(18), 181103. <http://dx.doi.org/10.1103/PhysRevLett.94.181103>.
- B. Abbott *et al.* (LIGO Scientific Collaboration and ALLEGRO Collaboration), 2007a. First cross-correlation analysis of interferometric and resonant-bar gravitational-wave data for stochastic backgrounds. *Physical Review D* 76(2), 022001. <http://dx.doi.org/10.1103/PhysRevD.76.022001>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2007b. Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Physical Review D* 76(8), 082001. <http://dx.doi.org/10.1103/PhysRevD.76.082001>.

- B. Abbott *et al.* (LIGO Scientific Collaboration), 2007c. Searching for a Stochastic Background of Gravitational Waves with the Laser Interferometer Gravitational-Wave Observatory. *The Astrophysical Journal* 659(2), 918. <http://dx.doi.org/10.1086/511329>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2007d. Upper limit map of a background of gravitational waves. *Physical Review D* 76(8), 082003. <http://dx.doi.org/10.1103/PhysRevD.76.082003>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2007e. Upper limits on gravitational wave emission from 78 radio pulsars. *Physical Review D* 76(4), 042001. <http://dx.doi.org/10.1103/PhysRevD.76.042001>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008a. All-sky search for periodic gravitational waves in LIGO S4 data. *Physical Review D* 77(2), 022001. <http://dx.doi.org/10.1103/PhysRevD.77.022001>.
- B. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), 2008b. Astrophysically triggered searches for gravitational waves: status and prospects. *Classical and Quantum Gravity* 25(11), 114051. <http://dx.doi.org/10.1088/0264-9381/25/11/114051>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008c. Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar. *The Astrophysical Journal* 683(1), L45. <http://dx.doi.org/10.1086/591526>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008d. Implications for the Origin of GRB 070201 from LIGO Observations. *The Astrophysical Journal* 681(2), 1419. <http://dx.doi.org/10.1086/587954>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008e. Search for Gravitational-Wave Bursts from Soft Gamma Repeaters. *Physical Review Letters* 101(21), 211102. <http://dx.doi.org/10.1103/PhysRevLett.101.211102>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008f. Search for gravitational waves associated with 39 gamma-ray bursts using data from the second, third, and fourth LIGO runs. *Physical Review D* 77(6), 062004. <http://dx.doi.org/10.1103/PhysRevD.77.062004>.
- B. Abbott *et al.* (LIGO Scientific Collaboration), 2008g. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D* 77(6), 062002. <http://dx.doi.org/10.1103/PhysRevD.77.062002>.

- B. Abbott *et al.* (LIGO Scientific Collaboration), 2009a. Einstein@Home search for periodic gravitational waves in LIGO S4 data. *Physical Review D* 79(2), 022001. <http://dx.doi.org/10.1103/PhysRevD.79.022001>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009b. All-Sky LIGO Search for Periodic Gravitational Waves in the Early Fifth-Science-Run Data. *Physical Review Letters* 102(11), 111102. <http://dx.doi.org/10.1103/PhysRevLett.102.111102>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), 2009c. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature* 460, 990. <http://dx.doi.org/10.1038/nature08278>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009d. Einstein@Home search for periodic gravitational waves in early S5 LIGO data. *Physical Review D* 80(4), 042003. <http://dx.doi.org/10.1103/PhysRevD.80.042003>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009e. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics* 72(7), 076901. <http://dx.doi.org/10.1088/0034-4885/72/7/076901>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009f. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Physical Review D* 80(10), 102001. <http://dx.doi.org/10.1103/PhysRevD.80.102001>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009g. Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data. *Physical Review D* 80(6), 062001. <http://dx.doi.org/10.1103/PhysRevD.80.062001>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009h. Search for gravitational waves from low mass binary coalescences in the first year of LIGO's S5 data. *Physical Review D* 79(12), 122001. <http://dx.doi.org/10.1103/PhysRevD.79.122001>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009i. Search for gravitational waves from low mass compact binary coalescence in 186 days of LIGO's fifth science run. *Physical Review D* 80(4), 047101. <http://dx.doi.org/10.1103/PhysRevD.80.047101>.

- B. P. Abbott *et al.* (LIGO Scientific Collaboration), 2009j. Search for high frequency gravitational-wave bursts in the first calendar year of LIGO's fifth science run. *Physical Review D* 80(10), 102002. <http://dx.doi.org/10.1103/PhysRevD.80.102002>.
- B. P. Abbott *et al.*, 2009k. Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm. *The Astrophysical Journal* 701(2), L68. <http://dx.doi.org/10.1088/0004-637X/701/2/L68>.
- B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), 2010. Searches for Gravitational Waves from Known Pulsars with Science Run 5 LIGO Data. *The Astrophysical Journal* 713(1), 671. <http://dx.doi.org/10.1088/0004-637X/713/1/671>.
- F. Acernese *et al.*, 2008. Virgo status. *Classical and Quantum Gravity* 25(18), 184001. <http://dx.doi.org/10.1088/0264-9381/25/18/184001>.
- T. Akgün and I. Wasserman, 2008. Toroidal magnetic fields in type II superconducting neutron stars. *Monthly Notices of the Royal Astronomical Society* 383, 1551. <http://dx.doi.org/10.1111/j.1365-2966.2007.12660.x>.
- M. Alford, C. Kouvaris, and K. Rajagopal, 2004. Gapless Color-Flavor-Locked Quark Matter. *Physical Review Letters* 92(22), 222001. <http://dx.doi.org/10.1103/PhysRevLett.92.222001>.
- N. Andersson, 1998. A New Class of Unstable Modes of Rotating Relativistic Stars. *The Astrophysical Journal* 502(2), 708. <http://dx.doi.org/10.1086/305919>.
- N. Andersson and G. L. Comer, 2007. Relativistic Fluid Dynamics: Physics for Many Different Scales. *Living Reviews in Relativity* 10(1). <http://www.livingreviews.org/lrr-2007-1>.
- N. Andersson, K. Glampedakis, and B. Haskell, 2009. Oscillations of dissipative superfluid neutron stars. *Physical Review D* 79(10), 103009. <http://dx.doi.org/10.1103/PhysRevD.79.103009>.
- Z. Arzoumanian, D. F. Chernoff, and J. M. Cordes, 2002. The Velocity Distribution of Isolated Radio Pulsars. *The Astrophysical Journal* 568(1), 289. <http://dx.doi.org/10.1086/338805>.
- P. Astone *et al.*, 2003. Methods and results of the IGEC search for burst gravitational waves in the years 1997–2000. *Physical Review D* 68(2), 022001. <http://dx.doi.org/10.1103/PhysRevD.68.022001>.

- P. Astone *et al.*, 2007. Results of the IGEC-2 search for gravitational wave bursts during 2005. *Physical Review D* 76(10), 102001. <http://dx.doi.org/10.1103/PhysRevD.76.102001>.
- ATNF, 2009 (cited Oct. 14). The Australia Telescope National Facility Pulsar Catalogue. <http://www.atnf.csiro.au/research/pulsar/psrcat>.
- W. Baade and R. Minkowski, 1954. Identification of the Radio Sources in Cassiopeia, Cygnus A, and Puppis A. *The Astrophysical Journal* 119, 206. <http://dx.doi.org/10.1086/145812>.
- W. Baade and F. Zwicky, 1934. Cosmic Rays from Super-novae. *Proceedings of the National Academy of Sciences of the United States of America* 20, 259. <http://dx.doi.org/10.1073/pnas.20.5.259>.
- J. G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower, 2007. Comparisons of binary black hole merger waveforms. *Classical and Quantum Gravity* 24(12), S25. <http://dx.doi.org/10.1088/0264-9381/24/12/S03>.
- R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, 1996. Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters. *Physical Review D* 53, 3033. <http://dx.doi.org/10.1103/PhysRevD.53.3033>.
- D. Baskaran, L. P. Grishchuk, and W. Zhao, 2010. Primordial Gravitational Waves and Cosmic Microwave Background Radiation. <http://arxiv.org/abs/1004.0804v1>.
- C. Bassa, Z. Wang, A. Cumming, and V. M. Kaspi (eds.), 2008. 40 Years of Pulsars: Millisecond Pulsars, Magnetars and More. *American Institute of Physics Conference Series*, vol. 983. <http://adsabs.harvard.edu/abs/2008AIPC..983.....B>.
- B. Bertotti, D. Brill, and R. Krotkov, 1962. “Experiments on Gravitation”. L. Witten (ed.), *Gravitation: an introduction to current research*, chap. 1. New York: John Wiley & Sons.
- J. Betzwieser, G. Mendell, and K. Riles, 2008 (Dec. 20). SFT V3 $h(t)$ Segment Generation. <https://www.lsc-group.phys.uwm.edu/twiki/bin/view/CW/SFTV3hoftSegmentGeneration>. LSC authorisation required.
- L. Bildsten, 1998. Gravitational Radiation and Rotation of Accreting Neutron Stars. *The Astrophysical Journal* 501, L89. <http://dx.doi.org/10.1086/311440>.

- L. Bildsten and G. Ushomirsky, 2000. Viscous Boundary-Layer Damping of r-Modes in Neutron Stars. *The Astrophysical Journal* 529(1), L33. <http://dx.doi.org/10.1086/312454>.
- G. S. Bisnovatyi-Kogan and B. V. Komberg, 1974. Pulsars and close binary systems. *Soviet Astronomy* 18, 217. <http://adsabs.harvard.edu/abs/1974SvA....18..217B>.
- L. Blanchet, 2006. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Reviews in Relativity* 9(4). <http://www.livingreviews.org/lrr-2006-4>.
- J. M. Blondin and K. Freese, 1986. Is the 1.5-ms pulsar a young neutron star? *Nature* 323, 786. <http://dx.doi.org/10.1038/323786a0>.
- S. Bonazzola and E. Gourgoulhon, 1996. Gravitational waves from pulsars: emission by the magnetic-field-induced distortion. *Astronomy & Astrophysics* 312, 675. <http://adsabs.harvard.edu/abs/1996A&A...312..675B>.
- R. Bondarescu, S. A. Teukolsky, and I. Wasserman, 2009. Spinning down newborn neutron stars: Nonlinear development of the r-mode instability. *Physical Review D* 79(10), 104003. <http://dx.doi.org/10.1103/PhysRevD.79.104003>.
- P. R. Brady and T. Creighton, 2000. Searching for periodic sources with LIGO. II. Hierarchical searches. *Physical Review D* 61(8), 082001. <http://dx.doi.org/10.1103/PhysRevD.61.082001>.
- J. Braithwaite and Å. Nordlund, 2006. Stable magnetic fields in stellar interiors. *Astronomy & Astrophysics* 450, 1077. <http://dx.doi.org/10.1051/0004-6361:20041980>.
- J. Braithwaite and H. C. Spruit, 2006. Evolution of the magnetic field in magnetars. *Astronomy & Astrophysics* 450, 1097. <http://dx.doi.org/10.1051/0004-6361:20041981>.
- E. F. Brown, 2000. Nuclear Heating and Melted Layers in the Inner Crust of an Accreting Neutron Star. *The Astrophysical Journal* 531, 988. <http://dx.doi.org/10.1086/308487>.
- E. F. Brown and L. Bildsten, 1998. The Ocean and Crust of a Rapidly Accreting Neutron Star: Implications for Magnetic Field Evolution and Thermonuclear Flashes. *The Astrophysical Journal* 496, 915. <http://dx.doi.org/10.1086/305419>.

- R. Buras, M. Rampp, H.-T. Janka, and K. Kifonidis, 2003. Improved Models of Stellar Core Collapse and Still No Explosions: What Is Missing? *Physical Review Letters* 90(24), 241101. <http://dx.doi.org/10.1103/PhysRevLett.90.241101>.
- A. Burrows, E. Livne, L. Dessart, C. D. Ott, , and J. Murphy, 2007. Features of the Acoustic Mechanism of Core-Collapse Supernova Explosions. *The Astrophysical Journal* 655(1), 416. <http://dx.doi.org/10.1086/509773>.
- L. Cadonati *et al.*, 2009. Status of NINJA: the Numerical INJection Analysis project. *Classical and Quantum Gravity* 26(11), 114008. <http://dx.doi.org/10.1088/0264-9381/26/11/114008>.
- D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, 2003. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature* 424, 42. <http://dx.doi.org/10.1038/nature01732>.
- D. Chakrabarty, M. J. Pivovarov, L. E. Hernquist, J. S. Heyl, and R. Narayan, 2001. The Central X-Ray Point Source in Cassiopeia A. *The Astrophysical Journal* 548(2), 800. <http://dx.doi.org/10.1086/318994>.
- N. Chamel and P. Haensel, 2008. Physics of Neutron Star Crusts. *Living Reviews in Relativity* 11(10). <http://www.livingreviews.org/lrr-2008-10>.
- A. R. Choudhuri and S. Konar, 2002. Diamagnetic screening of the magnetic field in accreting neutron stars. *Monthly Notices of the Royal Astronomical Society* 332, 933. <http://dx.doi.org/10.1046/j.1365-8711.2002.05362.x>.
- N. Christensen, S. Caride, G. Ely, and T. Isogai, 2008. S5 Coherence Studies. LIGO technical report, no. G080210-00-Z. <http://www.ligo.caltech.edu/docs/G/G080210-00/G080210-00.pdf>.
- A. Colaiuda, V. Ferrari, L. Gualtieri, and J. A. Pons, 2008. Relativistic models of magnetars: structure and deformations. *Monthly Notices of the Royal Astronomical Society* 385, 2080. <http://dx.doi.org/10.1111/j.1365-2966.2008.12966.x>.
- Condor. <http://www.cs.wisc.edu/condor>.
- J. H. Conway and N. J. A. Sloane, 1988a. Sphere Packings, Lattices and Groups. *Grundlehren der mathematischen Wissenschaften*, no. 290. New York: Springer-Verlag. ISBN: 038796617X.

- J. H. Conway and N. J. A. Sloane, 1988b. “Certain Important Lattices and Their Properties”. Conway & Sloane (1988a), chap. 4.
- J. H. Conway and N. J. A. Sloane, 1988c. “Coverings, Lattices and Quantizers”. Conway & Sloane (1988a), chap. 2.
- J. H. Conway and N. J. A. Sloane, 1988d. “Sphere Packings and Kissing Numbers”. Conway & Sloane (1988a), chap. 1.
- A. Cumming, 2005. “Magnetic Field Evolution During Neutron Star Recycling”. F. A. Rasio and I. H. Stairs (eds.), Binary Radio Pulsars, p. 311. *Astronomical Society of the Pacific Conference Series*, vol. 328. <http://adsabs.harvard.edu/abs/2005ASPC..328..311C>.
- A. Cumming, P. Arras, and E. Zweibel, 2004. Magnetic Field Evolution in Neutron Star Crusts Due to the Hall Effect and Ohmic Decay. *The Astrophysical Journal* 609, 999. <http://dx.doi.org/10.1086/421324>.
- A. Cumming, E. Zweibel, and L. Bildsten, 2001. Magnetic Screening in Accreting Neutron Stars. *The Astrophysical Journal* 557, 958. <http://dx.doi.org/10.1086/321658>.
- C. Cutler, 2002. Gravitational waves from neutron stars with large toroidal B fields. *Physical Review D* 66(8), 084025. <http://dx.doi.org/10.1103/PhysRevD.66.084025>.
- C. Cutler, I. Gholami, and B. Krishnan, 2005. Improved stack-slide searches for gravitational-wave pulsars. *Physical Review D* 72(4), 042004. <http://dx.doi.org/10.1103/PhysRevD.72.042004>.
- C. Cutler and B. F. Schutz, 2005. Generalized \mathcal{F} -statistic: Multiple detectors and multiple gravitational wave pulsars. *Physical Review D* 72(6), 063006. <http://dx.doi.org/10.1103/PhysRevD.72.063006>.
- A. De Luca, 2008. “Central Compact Objects in Supernova Remnants”. Bassa *et al.* (2008), p. 311. <http://adsabs.harvard.edu/abs/2008AIPC..983.....B>.
- V. Dergachev and K. Riles, 2005. Description of PowerFlux algorithms and implementation. LIGO technical report, no. T050186-00-Z. <http://www.ligo.caltech.edu/docs/T/T050186-00.pdf>.
- S. V. Dhurandhar, D. G. Blair, and M. E. Costa, 1996. Possibility of detecting gravitational waves from millisecond pulsars by resonant bar antennas. *Astronomy & Astrophysics* 311, 1043. <http://adsabs.harvard.edu/abs/1996A&A...311.1043D>.

- H. Dimmelmeier, C. D. Ott, A. Marek, and H.-T. Janka, 2008. Gravitational wave burst signal from core collapse of rotating stars. *Physical Review D* 78(6), 064056. <http://dx.doi.org/10.1103/PhysRevD.78.064056>.
- E. Dwek and R. G. Arendt, 2008. Infrared Echoes Reveal the Shock Breakout of the Cas A Supernova. *The Astrophysical Journal* 685, 976. <http://dx.doi.org/10.1086/589988>.
- A. Einstein, 1916. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik* 354(7), 769. <http://dx.doi.org/10.1002/andp.19163540702>. Original manuscript and translation available from the Albert Einstein Archives [<http://www.albert-einstein.org>], archival call no. 120-788.
- C.-A. Faucher-Giguère and V. M. Kaspi, 2006. Birth and Evolution of Isolated Radio Pulsars. *The Astrophysical Journal* 643(1), 332. <http://dx.doi.org/10.1086/501516>.
- R. A. Fesen, M. C. Hammell, J. Morse, R. A. Chevalier, K. J. Borkowski, M. A. Dopita, C. L. Gerardy, S. S. Lawrence, J. C. Raymond, and S. van den Bergh, 2006a. The Expansion Asymmetry and Age of the Cassiopeia A Supernova Remnant. *The Astrophysical Journal* 645, 283. <http://dx.doi.org/10.1086/504254>.
- R. A. Fesen, G. G. Pavlov, and D. Sanwal, 2006b. Near-Infrared and Optical Limits for the Central X-Ray Point Source in the Cassiopeia A Supernova Remnant. *The Astrophysical Journal* 636, 848. <http://dx.doi.org/10.1086/498087>.
- L. S. Finn, 2009. Response of interferometric gravitational wave detectors. *Physical Review D* 79(2), 022002. <http://dx.doi.org/10.1103/PhysRevD.79.022002>.
- L. S. Finn and D. F. Chernoff, 1993. Observing binary inspiral in gravitational radiation: One interferometer. *Physical Review D* 47(6), 2198. <http://dx.doi.org/10.1103/PhysRevD.47.2198>.
- C. L. Fryer, 1999. Mass Limits For Black Hole Formation. *The Astrophysical Journal* 522, 413. <http://dx.doi.org/10.1086/307647>.
- C. L. Fryer and V. Kalogera, 2001. Theoretical Black Hole Mass Distributions. *The Astrophysical Journal* 554, 548. <http://dx.doi.org/10.1086/321359>.

- N. K. Glendenning, 1992. First-order phase transitions with more than one conserved charge: Consequences for neutron stars. *Physical Review D* 46(4), 1274. <http://dx.doi.org/10.1103/PhysRevD.46.1274>.
- J. P. H. Goedbloed and S. Poedts, 2004. Principles of Magnetohydrodynamics. Cambridge: Cambridge University Press. ISBN: 0521623472.
- G. H. Golub and C. F. V. Loan, 1983. Matrix Computations. *Johns Hopkins Series in the Mathematical Sciences*, no. 3. Baltimore: Johns Hopkins University Press. ISBN: 0801830109.
- E. V. Gotthelf and J. P. Halpern, 2008. “CCO Pulsars as Anti-Magnetars: Evidence of Neutron Stars Weakly Magnetized at Birth”. Bassa *et al.* (2008), p. 320. <http://adsabs.harvard.edu/abs/2008AIPC...983.....B>.
- E. V. Gotthelf and J. P. Halpern, 2009. Discovery of a 112 ms X-Ray Pulsar in Puppis A: Further Evidence of Neutron Stars Weakly Magnetized at Birth. *The Astrophysical Journal* 695, L35. <http://dx.doi.org/10.1088/0004-637X/695/1/L35>.
- H. Grote (LIGO Scientific Collaboration), 2008. The status of GEO 600. *Classical and Quantum Gravity* 25(11), 114043. <http://dx.doi.org/10.1088/0264-9381/25/11/114043>.
- J. M. Hameury, S. Bonazzola, J. Heyvaerts, and J. P. Lasota, 1983. Magnetohydrostatics in the polar caps of the gamma-ray burst sources. *Astronomy & Astrophysics* 128, 369. <http://adsabs.harvard.edu/abs/1983A&A...128..369H>.
- I. W. Harry, B. Allen, and B. S. Sathyaprakash, 2009. Stochastic template placement algorithm for gravitational wave data analysis. *Physical Review D* 80(10), 104014. <http://dx.doi.org/10.1103/PhysRevD.80.104014>.
- J. W. Hartman, D. Bhattacharya, R. Wijers, and F. Verbunt, 1997. A study of the evolution of radio pulsars through improved population synthesis. *Astronomy & Astrophysics* 322, 477. <http://adsabs.harvard.edu/abs/1997A&A...322..477H>.
- B. Haskell, N. Andersson, D. I. Jones, and L. Samuelsson, 2007. Are Neutron Stars with Crystalline Color-Superconducting Cores Relevant for the LIGO Experiment? *Physical Review Letters* 99(23), 231101. <http://dx.doi.org/10.1103/PhysRevLett.99.231101>.

- B. Haskell, L. Samuelsson, K. Glampedakis, and N. Andersson, 2008. Modelling magnetically deformed neutron stars. *Monthly Notices of the Royal Astronomical Society* 385, 531. <http://dx.doi.org/10.1111/j.1365-2966.2008.12861.x>.
- J. C. Hayes, M. L. Norman, R. A. Fiedler, J. O. Bordner, P. S. Li, S. E. Clark, A. ud Doula, and M.-M. Mac Low, 2006. Simulating Radiating and Magnetized Flows in Multiple Dimensions with ZEUS-MP. *The Astrophysical Journal Supplement Series* 165, 188. <http://dx.doi.org/10.1086/504594>.
- A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, 2003. How Massive Single Stars End Their Life. *The Astrophysical Journal* 591, 288. <http://dx.doi.org/10.1086/375341>.
- A. Heger, S. E. Woosley, and H. C. Spruit, 2005. Presupernova Evolution of Differentially Rotating Massive Stars Including Magnetic Fields. *The Astrophysical Journal* 626, 350-363. <http://dx.doi.org/10.1086/429868>.
- A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, 1968. Observation of a Rapidly Pulsating Radio Source. *Nature* 217, 709. <http://dx.doi.org/10.1038/217709a0>.
- C. J. Horowitz and K. Kadau, 2009. Breaking Strain of Neutron Star Crust and Gravitational Waves. *Physical Review Letters* 102, 191102. <http://dx.doi.org/10.1103/PhysRevLett.102.191102>.
- D. W. Hughes, 1980. Did Flamsteed see the Cassiopeia A supernova? *Nature* 285, 132. <http://dx.doi.org/10.1038/285132a0>.
- S. A. Hughes, 2003. Listening to the universe with gravitational-wave astronomy. *Annals of Physics* 303(1), 142. [http://dx.doi.org/10.1016/S0003-4916\(02\)00025-8](http://dx.doi.org/10.1016/S0003-4916(02)00025-8).
- R. Hulse and J. Taylor, 1975. Discovery of a pulsar in a binary system. *The Astrophysical Journal* 195, L51. <http://dx.doi.org/10.1086/181708>.
- U. Hwang, J. M. Laming, C. Badenes, F. Berendse, J. Blondin, D. Cioffi, T. DeLaney, D. Dewey, R. Fesen, K. A. Flanagan, C. L. Fryer, P. Ghavamian, J. P. Hughes, J. A. Morse, P. P. Plucinsky, R. Petre, M. Pohl, L. Rudnick, R. Sankrit, P. O. Slane, R. K. Smith, J. Vink, and J. S. Warren, 2004. A Million Second Chandra View of Cassiopeia A. *The Astrophysical Journal* 615, L117. <http://dx.doi.org/10.1086/426186>.

- H.-T. Janka, 2004. “Neutron Star Formation and Birth Properties”. F. Camilo and B. M. Gaensler (eds.), *Young Neutron Stars and Their Environments*, p. 3. *IAU Symposium*, vol. 218. <http://adsabs.harvard.edu/abs/2004IAUS...218....3J>.
- P. Jaranowski and A. Królak, 1999. Data analysis of gravitational-wave signals from spinning neutron stars. II. Accuracy of estimation of parameters. *Physical Review D* 59(6), 063003. <http://dx.doi.org/10.1103/PhysRevD.59.063003>.
- P. Jaranowski and A. Królak, 2000. Data analysis of gravitational-wave signals from spinning neutron stars. III. Detection statistics and computational requirements. *Physical Review D* 61(6), 062001. <http://dx.doi.org/10.1103/PhysRevD.61.062001>.
- P. Jaranowski, A. Królak, and B. F. Schutz, 1998. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Physical Review D* 58(6), 063001. <http://dx.doi.org/10.1103/PhysRevD.58.063001>.
- N. L. Johnson and S. Kotz, 1970. *Continuous Univariate Distributions*, vol. 2, chap. 28. Boston: Houghton Mifflin. ISBN: 0471446270.
- D. I. Jones and N. Andersson, 2002. Gravitational waves from freely precessing neutron stars. *Monthly Notices of the Royal Astronomical Society* 331(1), 203. <http://dx.doi.org/10.1046/j.1365-8711.2002.05180.x>.
- J. Kanner, T. L. Huard, S. Marka, D. C. Murphy, J. Piscionere, M. Reed, and P. Shawhan, 2008. LOOC UP: locating and observing optical counterparts to gravitational wave bursts. *Classical and Quantum Gravity* 25(18), 184034. <http://dx.doi.org/10.1088/0264-9381/25/18/184034>.
- D. L. Kaplan, S. R. Kulkarni, and S. S. Murray, 2001. Search for a Near-Infrared Counterpart to the Cassiopeia A X-Ray Point Source. *The Astrophysical Journal* 558, 270. <http://dx.doi.org/10.1086/322459>.
- Y. Kim, G. H. Rieke, O. Krause, K. Misselt, R. Indebetouw, and K. E. Johnson, 2008. Structure of the Interstellar Medium around Cas A. *The Astrophysical Journal* 678, 287. <http://dx.doi.org/10.1086/533426>.
- B. Knippel and A. Sedrakian, 2009. Gravitational radiation from crystalline color-superconducting hybrid stars. *Physical Review D* 79(8), 083007. <http://dx.doi.org/10.1103/PhysRevD.79.083007>.

- B. Knispel and B. Allen, 2008. Blandford's argument: The strongest continuous gravitational wave signal. *Physical Review D* 78(4), 044031. <http://dx.doi.org/10.1103/PhysRevD.78.044031>.
- K. D. Kokkotas, 2008. Gravitational Wave Astronomy. <http://arxiv.org/abs/0809.1602>.
- K. D. Kokkotas and B. Schmidt, 1999. Quasi-Normal Modes of Stars and Black Holes. *Living Reviews in Relativity* 2(2). <http://www.livingreviews.org/lrr-1999-2>.
- S. Konar and D. Bhattacharya, 1997. Magnetic field evolution of accreting neutron stars. *Monthly Notices of the Royal Astronomical Society* 284, 311. <http://adsabs.harvard.edu/abs/1997MNRAS.284..311K>.
- S. Konar and D. Bhattacharya, 1999. Magnetic field evolution of accreting neutron stars - III. *Monthly Notices of the Royal Astronomical Society* 308, 795. <http://dx.doi.org/10.1046/j.1365-8711.1999.02781.x>.
- D. Konenkov and U. Geppert, 2001. The evolution of the core and surface magnetic fields in isolated neutron stars. *Monthly Notices of the Royal Astronomical Society* 325, 426. <http://dx.doi.org/10.1046/j.1365-8711.2001.04469.x>.
- R. Kosiński and M. Hanasz, 2006. On the influence of cooling and heating processes on Parker instability. *Monthly Notices of the Royal Astronomical Society* 368, 759. <http://dx.doi.org/10.1111/j.1365-2966.2006.10142.x>.
- M. Kramer, 2005. "Pulsars". L. Gurvits, S. Frey, and S. Rawlings (eds.), Radio Astronomy from Karl Jansky to Microjansky, p. 219. *EAS Publications Series*, vol. 15. <http://dx.doi.org/10.1051/eas:2005155>.
- M. Kramer, K. Xilouris, A. Jessner, D. Lorimer, R. Wielebinski, and A. Lyne, 1997. Origin of pulsar radio emission. I. High frequency data. *Astronomy & Astrophysics* 322, 846. <http://adsabs.harvard.edu/abs/1997A&A...322..846K>.
- O. Krause, S. M. Birkmann, T. Usuda, T. Hattori, M. Goto, G. H. Rieke, and K. A. Misselt, 2008. The Cassiopeia A Supernova Was of Type IIb. *Science* 320(5880), 1195. <http://dx.doi.org/10.1126/science.1155788>.
- O. Krause, G. H. Rieke, S. M. Birkmann, E. Le Floc'h, K. D. Gordon, E. Egami, J. Biegging, J. P. Hughes, E. T. Young, J. L. Hinz, S. P. Quanz,

- and D. C. Hines, 2005. Infrared Echoes near the Supernova Remnant Cassiopeia A. *Science* 308(5728), 1604. <http://dx.doi.org/10.1126/science.1112035>.
- B. Krishnan and A. M. Sintes, 2007. Hough search with improved sensitivity. LIGO technical report, no. T070124-00-Z. <http://www.ligo.caltech.edu/docs/T/T070124-00.pdf>.
- LALApps. LSC Algorithm Library Applications. <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalapps.html>.
- L. D. Landau, 1967. "On the theory of stars". D. ter Haar (ed.), *Collected Papers of L. D. Landau*, chap. 8. New York: Gordon and Breach. Originally published in *Physikalische Zeitschrift der Sowjetunion*, volume 1, page 285 (1932).
- M. Landry and G. Mendell, 2007 (Sept. 13). S5 calibration, $h(t)$, DQ flags and SFT production. <https://www.lsc-group.phys.uwm.edu/cgi-bin/enote.pl?nb=puls5general&action=view&page=26>. LSC authorisation required.
- A. Lazzarini, 2007. Update from LIGO Laboratory. LIGO technical report, no. G070649-00-M. <http://www.ligo.caltech.edu/docs/G/G070649-00/G070649-00.pdf>.
- LIGO Laboratory, 2009a (July 14). Advanced LIGO. <http://www.ligo.caltech.edu/advLIGO>.
- LIGO Laboratory, 2009b (cited Sept. 3). Download site for archived publication-quality interferometer sensitivity curves. http://www.ligo.caltech.edu/~jzweizig/distribution/LSC_Data.
- LIGO Laboratory, 2009c (June 15). Firm Date Set for Start of S6. <http://ligonews.blogspot.com/2009/06/firm-date-set-for-start-of-s6.html>.
- L. Lindblom and B. J. Owen, 2002. Effect of hyperon bulk viscosity on neutron-star r-modes. *Physical Review D* 65(6), 063006. <http://dx.doi.org/10.1103/PhysRevD.65.063006>.
- L. Lindblom, B. J. Owen, and S. M. Morsink, 1998. Gravitational Radiation Instability in Hot Young Neutron Stars. *Physical Review Letters* 80(22), 4843. <http://dx.doi.org/10.1103/PhysRevLett.80.4843>.

- L. Lindblom, B. J. Owen, and G. Ushomirsky, 2000. Effect of a neutron-star crust on the r-mode instability. *Physical Review D* 62(8), 084030. <http://dx.doi.org/10.1103/PhysRevD.62.084030>.
- L.-M. Lin, 2007. Constraining crystalline color superconducting quark matter with gravitational-wave data. *Physical Review D* 76(8), 081502(R). <http://dx.doi.org/10.1103/PhysRevD.76.081502>.
- D. R. Lorimer, 2008. Binary and Millisecond Pulsars. *Living Reviews in Relativity* 11(8). <http://www.livingreviews.org/lrr-2008-8>.
- R. V. E. Lovelace, M. M. Romanova, and G. S. Bisnovatyi-Kogan, 2005. Screening of the Magnetic Field of Disk Accreting Stars. *The Astrophysical Journal* 625, 957. <http://dx.doi.org/10.1086/429532>.
- LSC (LIGO Scientific Collaboration), 2009 (cited Oct. 26). LSC Channel Wiki. https://ldas-jobs.ligo.caltech.edu/cgi-bin/chanwiki?LSC_Channel_Wiki. LSC authorisation required.
- J. Madsen, 2000. Probing Strange Stars and Color Superconductivity by *r*-Mode Instabilities in Millisecond Pulsars. *Physical Review Letters* 85(1), 10. <http://dx.doi.org/10.1103/PhysRevLett.85.10>.
- G. M. Manca and M. Vallisneri, 2010. Cover art: Issues in the metric-guided and metric-less placement of random and stochastic template banks. *Physical Review D* 81(2), 024004. <http://dx.doi.org/10.1103/PhysRevD.81.024004>.
- R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, 2005. The Australia Telescope National Facility Pulsar Catalogue. *The Astronomical Journal* 129, 1993. <http://dx.doi.org/10.1086/428488>.
- M. Mannarelli, K. Rajagopal, and R. Sharma, 2007. Rigidity of crystalline color superconducting quark matter. *Physical Review D* 76(7), 074026. <http://dx.doi.org/10.1103/PhysRevD.76.074026>.
- A. Marek and H.-T. Janka, 2009. Delayed Neutrino-Driven Supernova Explosions Aided by the Standing Accretion-Shock Instability. *The Astrophysical Journal* 694(1), 664. <http://dx.doi.org/10.1088/0004-637X/694/1/664>.
- M. A. McLaughlin, J. M. Cordes, A. A. Deshpande, B. M. Gaensler, T. H. Hankins, V. M. Kaspi, and J. S. Kern, 2001. Upper Limits on Periodic, Pulsed Radio Emission from the X-Ray Point Source in Cassiopeia A. *The Astrophysical Journal* 547, L41. <http://dx.doi.org/10.1086/318891>.

- A. Melatos, 2000. Radiative precession of an isolated neutron star. *Monthly Notices of the Royal Astronomical Society* 313, 217. <http://dx.doi.org/10.1046/j.1365-8711.2000.03031.x>.
- A. Melatos and D. J. B. Payne, 2005. Gravitational Radiation from an Accreting Millisecond Pulsar with a Magnetically Confined Mountain. *The Astrophysical Journal* 623, 1044. <http://dx.doi.org/10.1086/428600>.
- A. Melatos and E. S. Phinney, 2001. Hydromagnetic Structure of a Neutron Star Accreting at Its Polar Caps. *Publications of the Astronomical Society of Australia* 18, 421. <http://dx.doi.org/10.1071/AS01056>.
- G. Mendell, 2008 (Jan. 25). http://www.ldas-cit.ligo.caltech.edu/ldas_outgoing/GapChecker/S5hoft/archive/MISSING_L1_RDS_C03_L1.Fri_Jan_25_16_34_23_PST_2008.txt. LSC authorisation required.
- G. Mendell and K. Wette, 2006. Parameter Estimation Using Short Fourier Transforms. LIGO technical report, no. T060286-00-Z. <http://www.ligo.caltech.edu/docs/T/T060286-00.pdf>.
- G. Mendell and K. Wette, 2008. Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves. *Classical and Quantum Gravity* 25(11), 114044. <http://dx.doi.org/10.1088/0264-9381/25/11/114044>.
- S. Mereghetti, A. Tiengo, and G. L. Israel, 2002. The X-Ray Source at the Center of the Cassiopeia A Supernova Remnant. *The Astrophysical Journal* 569, 275. <http://dx.doi.org/10.1086/339277>.
- C. Messenger, R. Prix, and M. A. Papa, 2009. Random template banks and relaxed lattice coverings. *Physical Review D* 79(10), 104017. <http://dx.doi.org/10.1103/PhysRevD.79.104017>.
- S. D. Mohanty, 2003. Efficient Algorithm for computing a Running Median. T030168-00-D. <http://www.ligo.caltech.edu/docs/T/T030168-00.pdf>.
- S. S. Murray, S. M. Ransom, M. Juda, U. Hwang, and S. S. Holt, 2002. Is the Compact Source at the Center of Cassiopeia A Pulsed? *The Astrophysical Journal* 566, 1039. <http://dx.doi.org/10.1086/338224>.
- NCI National Facility. National Computational Infrastructure National Facility. <http://nf.nci.org.au>. Formerly the APAC (Australian Partnership for Advanced Computing) National Facility.

- J. R. Oppenheimer and G. M. Volkoff, 1939. On Massive Neutron Cores. *Physical Review* 55(4), 374. <http://dx.doi.org/10.1103/PhysRev.55.374>.
- J. P. Ostriker and J. E. Gunn, 1969. On the Nature of Pulsars. I. Theory. *The Astrophysical Journal* 157, 1395. <http://dx.doi.org/10.1086/150160>.
- C. D. Ott, A. Burrows, T. A. Thompson, E. Livne, and R. Walder, 2006. The Spin Periods and Rotational Profiles of Neutron Stars at Birth. *The Astrophysical Journal Supplement Series* 164, 130. <http://dx.doi.org/10.1086/500832>.
- B. J. Owen, 1996. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Physical Review D* 53(12), 6749. <http://dx.doi.org/10.1103/PhysRevD.53.6749>.
- B. J. Owen, 2005. Maximum Elastic Deformations of Compact Stars with Exotic Equations of State. *Physical Review Letters* 95(21), 211101. <http://dx.doi.org/10.1103/PhysRevLett.95.211101>.
- B. J. Owen, 2006. Detectability of periodic gravitational waves by initial interferometers. *Classical and Quantum Gravity* 23(8), S1. <http://dx.doi.org/10.1088/0264-9381/23/8/S01>.
- B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, 1998. Gravitational waves from hot young rapidly rotating neutron stars. *Physical Review D* 58(8), 084020. <http://dx.doi.org/10.1103/PhysRevD.58.084020>.
- C. Palomba, 2000. Pulsars ellipticity revised. *Astronomy & Astrophysics* 354, 163. <http://adsabs.harvard.edu/abs/2000A&A...354..163P>.
- G. G. Pavlov and G. J. M. Luna, 2009. A Dedicated Chandra ACIS Observation of the Central Compact Object in the Cassiopeia A Supernova Remnant. *The Astrophysical Journal* 703(1), 910. <http://dx.doi.org/10.1088/0004-637X/703/1/910>.
- G. G. Pavlov, V. E. Zavlin, B. Aschenbach, J. Trümper, and D. Sanwal, 2000. The Compact Central Object in Cassiopeia A: A Neutron Star with Hot Polar Caps or a Black Hole? *The Astrophysical Journal* 531, L53. <http://dx.doi.org/10.1086/312521>.
- D. J. B. Payne and A. Melatos, 2004. Burial of the polar magnetic field of an accreting neutron star - I. Self-consistent analytic and numerical

- equilibria. *Monthly Notices of the Royal Astronomical Society* 351, 569. <http://dx.doi.org/10.1111/j.1365-2966.2004.07798.x>.
- D. J. B. Payne and A. Melatos, 2007. Burial of the polar magnetic field of an accreting neutron star - II. Hydromagnetic stability of axisymmetric equilibria. *Monthly Notices of the Royal Astronomical Society* 376, 609. <http://dx.doi.org/10.1111/j.1365-2966.2007.11451.x>.
- H. J. Pletsch, 2008. Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection. *Physical Review D* 78(10), 102005. <http://dx.doi.org/10.1103/PhysRevD.78.102005>.
- R. Prix, 2006. \mathcal{F} -statistic bias in noise-estimator. <http://www.aei.mpg.de/~repr/EnoteEntries/FstatNoiseBias.pdf>.
- R. Prix, 2007a. Search for continuous gravitational waves: Metric of the multidetector \mathcal{F} -statistic. *Physical Review D* 75(2), 023004. <http://dx.doi.org/10.1103/PhysRevD.75.023004>.
- R. Prix, 2007b. Template-based searches for gravitational waves: efficient lattice covering of flat parameter spaces. *Classical and Quantum Gravity* 24(19), S481. <http://dx.doi.org/10.1088/0264-9381/24/19/S11>.
- R. Prix, 2008 (Oct. 3). The F-statistic and its implementation in ComputeFStatistic_v2. <https://dcc.ligo.org/DocDB/0001/T0900149/001/CFSv2-LIGO-T0900149-v1.pdf>.
- R. Prix, 2009. "Gravitational Waves from Spinning Neutron Stars". W. Becker (ed.), *Neutron Stars and Pulsars*, chap. 24. *Astrophysics and Space Science Library*, vol. 357. Berlin: Springer-Verlag. http://dx.doi.org/10.1007/978-3-540-76965-1_24. Preprint available at https://dcc.ligo.org/public/0000/P060039/002/LIGO-P060039-v2_HeraeusReview.pdf.
- R. Prix and Y. Itoh, 2005. Global parameter-space correlations of coherent searches for continuous gravitational waves. *Classical and Quantum Gravity* 22, S1003. <http://dx.doi.org/10.1088/0264-9381/22/18/S14>.
- R. Prix and B. Krishnan, 2009. Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics. *Classical and Quantum Gravity* 26(20), 204013. <http://dx.doi.org/10.1088/0264-9381/26/20/204013>.

- S. M. Ransom, 2002. “Fast Search Techniques for High Energy Pulsars”. P. O. Slane and B. M. Gaensler (eds.), *Neutron Stars in Supernova Remnants*, p. 361. *Astronomical Society of the Pacific Conference Series*, vol. 271. <http://adsabs.harvard.edu/abs/2002ASPC..271..361R>.
- J. E. Reed, J. J. Hester, A. C. Fabian, and P. F. Winkler, 1995. The Three-dimensional Structure of the Cassiopeia A Supernova Remnant. I. The Spherical Shell. *The Astrophysical Journal* 440, 706. <http://dx.doi.org/10.1086/175308>.
- S. P. Reynolds, K. J. Borkowski, D. A. Green, U. Hwang, I. Harrus, and R. Petre, 2008. The Youngest Galactic Supernova Remnant: G1.9+0.3. *The Astrophysical Journal* 680, L41. <http://dx.doi.org/10.1086/589570>.
- M. Rheinhardt and U. Geppert, 2002. Hall-Drift Induced Magnetic Field Instability in Neutron Stars. *Physical Review Letters* 88, 101103. <http://dx.doi.org/10.1103/PhysRevLett.88.101103>.
- K. Riles *et al.*, 2008 (July). S5 Pulsar Hardware Injections. https://www.lsc-group.phys.uwm.edu/ligovirgo/cw/protected/S5_injections/pulsar_injections.html. LSC authorisation required.
- K. Riles *et al.*, 2009 (cited Oct. 17). S5 Segment Data Quality Repository. <http://gallatin.physics.lsa.umich.edu/~keithr/S5DQ>.
- R. W. Romani, 1990. A unified model of neutron-star magnetic fields. *Nature* 347, 741. <http://dx.doi.org/10.1038/347741a0>.
- M. M. Romanova, A. K. Kulkarni, and R. V. E. Lovelace, 2008. Unstable Disk Accretion onto Magnetized Stars: First Global Three-dimensional Magnetohydrodynamic Simulations. *The Astrophysical Journal* 673, L171. <http://dx.doi.org/10.1086/527298>.
- M. Ruderman, 1969. Neutron Starquakes and Pulsar Periods. *Nature* 223(597). <http://dx.doi.org/10.1038/223597b0>.
- M. Ruderman, T. Zhu, , and K. Chen, 1998. Neutron Star Magnetic Field Evolution, Crust Movement, and Glitches. *The Astrophysical Journal* 492, 267. <http://dx.doi.org/10.1086/305026>.
- E. Ryan, R. M. Wagner, and S. G. Starrfield, 2001. New Optical Constraints on the Presence of a Compact Central Object in Cassiopeia A. *The Astrophysical Journal* 548, 811. <http://dx.doi.org/10.1086/319034>.

- P. R. Saulson, 1994. Fundamentals of Interferometric Gravitational Wave Detectors. Singapore: World Scientific. ISBN: 9810218206.
- A. Schuermann and F. Vallentin, 2008. Methods in the Local Theory of Packing and Covering Lattices. <http://arxiv.org/abs/math/0412320v2>.
- A. Schuermann and F. Vallentin, 2009 (cited July 10). Geometry of Lattices and Algorithms. http://www.math.uni-magdeburg.de/lattice_geometry.
- B. F. Schutz and B. Sathyaprakash, 2009. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Reviews in Relativity* 12(2). <http://www.livingreviews.org/lrr-2009-2>.
- A. C. Searle, P. J. Sutton, M. Tinto, and G. Woan, 2008. Robust Bayesian detection of unmodelled bursts. *Classical and Quantum Gravity* 25(11), 114038. <http://dx.doi.org/10.1088/0264-9381/25/11/114038>.
- D. A. Shaddock, 2008. Space-based gravitational wave detection with LISA. *Classical and Quantum Gravity* 25(11), 114012. <http://dx.doi.org/10.1088/0264-9381/25/11/114012>.
- S. L. Shapiro and S. A. Teukolsky, 1983. Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects. New York: John Wiley & Sons. ISBN: 0471873179.
- P. Shawhan, 2009. (Some Additional) Multi-Messenger Astronomy Opportunities for the LSC and Virgo. G080438-00-Z. <http://www.ligo.caltech.edu/docs/G/G080438-00/G080438-00.pdf>.
- A. M. Sintes and B. Krishnan, 2006. Improved Hough search for gravitational wave pulsars. *Journal of Physics: Conference Series* 32, 206. <http://dx.doi.org/10.1088/1742-6596/32/1/031>.
- G. Srinivasan, D. Bhattacharya, A. G. Muslimov, and A. J. Tsygan, 1990. A novel mechanism for the decay of neutron star magnetic fields. *Current Science* 59, 31. <http://adsabs.harvard.edu/abs/1990CSci...59...31S>.
- D. C. Srivastava and S. K. Sahay, 2002. Data analysis of continuous gravitational wave: Fourier transform – I. *Monthly Notices of the Royal Astronomical Society* 337(1), 305. <http://dx.doi.org/10.1046/j.1365-8711.2002.06032.x>.
- N. Stergioulas, 2003. Rotating Stars in Relativity. *Living Reviews in Relativity* 6(3). <http://www.livingreviews.org/lrr-2003-3>.

- J. M. Stone and M. L. Norman, 1992a. ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests. *The Astrophysical Journal Supplement Series* 80, 753. <http://dx.doi.org/10.1086/191680>.
- J. M. Stone and M. L. Norman, 1992b. ZEUS-2D: A Radiation Magnetohydrodynamics Code for Astrophysical Flows in Two Space Dimensions. II. The Magnetohydrodynamic Algorithms and Tests. *The Astrophysical Journal Supplement Series* 80, 791. <http://dx.doi.org/10.1086/191681>.
- R. E. Taam and E. P. J. van den Heuvel, 1986. Magnetic field decay and the origin of neutron star binaries. *The Astrophysical Journal* 305, 235. <http://dx.doi.org/10.1086/164243>.
- R. Takahashi, K. Arai, D. Tatsumi, M. Fukushima, T. Yamazaki, M.-K. Fujimoto, K. Agatsuma, Y. Arase, N. Nakagawa, A. Takamori, K. Tsubono, R. DeSalvo, A. Bertolini, S. Marka, and V. Sannibale (TAMA Collaboration), 2008. Operational status of TAMA300 with the seismic attenuation system (SAS). *Classical and Quantum Gravity* 25(11), 114036. <http://dx.doi.org/10.1088/0264-9381/25/11/114036>.
- H. Tananbaum, 1999. IAU Circular, no. 7246. <http://www.cfa.harvard.edu/iauc/07200/07246.html>.
- T. M. Tauris, E. P. J. van den Heuvel, , and G. J. Savonije, 2000. Formation of Millisecond Pulsars with Heavy White Dwarf Companions: Extreme Mass Transfer on Subthermal Timescales. *The Astrophysical Journal* 530(2), L93. <http://dx.doi.org/10.1086/312496>.
- J. H. Taylor and J. M. Weisberg, 1982. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *The Astrophysical Journal* 253, 908. <http://dx.doi.org/10.1086/159690>.
- J. H. Taylor and J. M. Weisberg, 1989. Further experimental tests of relativistic gravity using the binary pulsar PSR 1913 + 16. *The Astrophysical Journal* 345, 434. <http://dx.doi.org/10.1086/167917>.
- K. S. Thorne, 1980. Multipole expansions of gravitational radiation. *Reviews of Modern Physics* 52(2), 299. <http://dx.doi.org/10.1103/RevModPhys.52.299>.
- K. S. Thorne, 1987. "Gravitational Radiation". S. W. Hawking and W. Israel (eds.), *Three Hundred Years of Gravitation*, pp. 330-458. Cambridge: Cambridge University Press.

- V. Urpin, U. Geppert, and D. Konenkov, 1998. Magnetic and spin evolution of neutron stars in close binaries. *Monthly Notices of the Royal Astronomical Society* 295, 907. <http://dx.doi.org/10.1046/j.1365-8711.1998.01375.x>.
- G. Ushomirsky, C. Cutler, and L. Bildsten, 2000. Deformations of accreting neutron star crusts and gravitational wave emission. *Monthly Notices of the Royal Astronomical Society* 319(3), 902. <http://dx.doi.org/10.1046/j.1365-8711.2000.03938.x>.
- C. Van Den Broeck, 2005. The gravitational wave spectrum of non-axisymmetric, freely precessing neutron stars. *Classical and Quantum Gravity* 22(9), 1825. <http://dx.doi.org/10.1088/0264-9381/22/9/022>.
- E. P. J. van den Heuvel and O. Bitzaraki, 1995. The magnetic field strength versus orbital period relation for binary radio pulsars with low-mass companions: evidence for neutron-star formation by accretion-induced collapse? *Astronomy & Astrophysics* 297, L41. <http://adsabs.harvard.edu/abs/1995A&A...297L..41V>.
- C. A. van Eysden and A. Melatos, 2008. Gravitational radiation from pulsar glitches. *Classical and Quantum Gravity* 25(22), 225020. <http://dx.doi.org/10.1088/0264-9381/25/22/225020>.
- M. Vigelius, 2008. Gravitational Radiation from Accreting X-ray Binaries. Doctoral thesis. The University of Melbourne.
- M. Vigelius and A. Melatos, 2008. Three-dimensional stability of magnetically confined mountains on accreting neutron stars. *Monthly Notices of the Royal Astronomical Society* 386, 1294. <http://dx.doi.org/10.1111/j.1365-2966.2008.13139.x>.
- M. Vigelius and A. Melatos, 2009a. Improved estimate of the detectability of gravitational radiation from a magnetically confined mountain on an accreting neutron star. *Monthly Notices of the Royal Astronomical Society* 395, 1972. <http://dx.doi.org/10.1111/j.1365-2966.2009.14690.x>.
- M. Vigelius and A. Melatos, 2009b. Resistive relaxation of a magnetically confined mountain on an accreting neutron star. *Monthly Notices of the Royal Astronomical Society* 395, 1985. <http://dx.doi.org/10.1111/j.1365-2966.2009.14698.x>.

- R. V. Wagoner, 1984. Gravitational radiation from accreting neutron stars. *The Astrophysical Journal* 278, 345. <http://dx.doi.org/10.1086/161798>.
- Z. Wang, D. L. Kaplan, and D. Chakrabarty, 2007. A Search for fallback Disks in Four Young Supernova Remnants. *The Astrophysical Journal* 655, 261. <http://dx.doi.org/10.1086/509869>.
- A. L. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, 2008. Detecting gravitational wave emission from the known accreting neutron stars. *Monthly Notices of the Royal Astronomical Society* 389, 839. <http://dx.doi.org/10.1111/j.1365-2966.2008.13594.x>.
- K. Wette, B. J. Owen, B. Allen, M. Ashley, J. Betzwieser, N. Christensen, T. D. Creighton, V. Dergachev, I. Gholami, E. Goetz, R. Gustafson, D. Hammer, D. I. Jones, B. Krishnan, M. Landry, B. Machenschalk, D. E. McClelland, G. Mendell, C. J. Messenger, M. A. Papa, P. Patel, M. Pitkin, H. J. Pletsch, R. Prix, K. Riles, L. S. de la Jordana, S. M. Scott, A. M. Sintes, M. Trias, J. T. Whelan, and G. Woan, 2008. Searching for gravitational waves from Cassiopeia A with LIGO. *Classical and Quantum Gravity* 25(23), 235011. <http://dx.doi.org/10.1088/0264-9381/25/23/235011>.
- K. Wette, M. Vigelius, and A. Melatos, 2010. Sinking of a magnetically confined mountain on an accreting neutron star. *Monthly Notices of the Royal Astronomical Society* 402, 1099. <http://dx.doi.org/10.1111/j.1365-2966.2009.15937.x>.
- D. M. Whitbeck, 2006. Observational Consequences of Gravitational Wave Emission from Spinning Compact Sources. Doctoral thesis. The Pennsylvania State University. <http://etda.libraries.psu.edu/theses/approved/WorldWideIndex/ETD-1414/index.html>.
- R. A. M. J. Wijers, 1997. Evidence against field decay proportional to accreted mass in neutron stars. *Monthly Notices of the Royal Astronomical Society* 287, 607. <http://adsabs.harvard.edu/abs/1997MNRAS.287..607W>.
- C. M. Will, 1993. *Theory and Experiment in Gravitational Physics*, revised edn. Cambridge: Cambridge University Press. ISBN: 0521439736.
- C. M. Will, 2006 (cited July 12). The Confrontation between General Relativity and Experiment. *Living Reviews in Relativity* 9(3). <http://www.livingreviews.org/lrr-2006-3>.

- Y. Wu, C. D. Matzner, and P. Arras, 2001. R-Modes in Neutron Stars with Crusts: Turbulent Saturation, Spin-down, and Crust Melting. *The Astrophysical Journal* 549, 1011. <http://dx.doi.org/10.1086/319446>.
- R. X. Xu, 2003. Solid Quark Stars? *The Astrophysical Journal* 596, L59. <http://dx.doi.org/10.1086/379209>.
- K. Yamamoto, T. Uchiyama, S. Miyoki, M. Ohashi, K. Kuroda, H. Ishitsuka, T. Akutsu, S. Telada, T. Tomaru, T. Suzuki, N. Sato, Y. Saito, Y. Higashi, T. Haruyama, A. Yamamoto, T. Shintomi, D. Tatsumi, M. Ando, H. Tagoshi, N. Kanda, N. Awaya, S. Yamagishi, H. Takahashi, A. Araya, A. Takamori, S. Takemoto, T. Higashi, H. Hayakawa, W. Morii, and J. Akamatsu, 2008. Current status of the CLIO project. *Journal of Physics: Conference Series* 122, 012002. <http://dx.doi.org/10.1088/1742-6596/122/1/012002>.
- M. D. Young, R. N. Manchester, and S. Johnston, 1999. A radio pulsar with an 8.5-second period that challenges emission models. *Nature* 400, 848. <http://dx.doi.org/10.1038/23650>.
- ZEUS-MP. <http://lca.ucsd.edu/portal/codes/zeusmp2>. Version 2.1.2.
- C. M. Zhang and Y. Kojima, 2006. The bottom magnetic field and magnetosphere evolution of neutron star in low-mass X-ray binary. *Monthly Notices of the Royal Astronomical Society* 366, 137. <http://dx.doi.org/10.1111/j.1365-2966.2005.09802.x>.
- W. Zhang, S. E. Woosley, and A. Heger, 2008. Fallback and Black Hole Production in Massive Stars. *The Astrophysical Journal* 679, 639. <http://dx.doi.org/10.1086/526404>.
- M. Zimmermann, 1980. Gravitational waves from rotating and precessing rigid bodies. II. General solutions and computationally useful formulas. *Physical Review D* 21(4), 891. <http://dx.doi.org/10.1103/PhysRevD.21.891>.
- M. Zimmermann and E. Szedenits, 1979. Gravitational waves from rotating and precessing rigid bodies: Simple models and applications to pulsars. *Physical Review D* 20(2), 351. <http://dx.doi.org/10.1103/PhysRevD.20.351>.
- J. Zweizig *et al.*, 2009 (cited Oct. 17). S5 Data Quality. http://www.ligo.caltech.edu/~jzweizig/S5_Data_Quality.