

**Analysis of spatial mode sensitivity of a gravitational wave
interferometer and a targeted search for gravitational
radiation from the Crab pulsar**

by

Joseph Betzwieser

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 2007


© Joseph Betzwieser, MMVII. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Author
Department of Physics
August 31, 2007

Certified by
Nergis Mavalvala
Associate Professor of Physics
Thesis Supervisor

Accepted by
Thomas J. Greytak
Associate Department Head for Education



ARCHIVES

Analysis of spatial mode sensitivity of a gravitational wave interferometer and a targeted search for gravitational radiation from the Crab pulsar

by

Joseph Betzwieser

Submitted to the Department of Physics
on August 31, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Over the last several years the Laser Interferometer Gravitational Wave Observatory (LIGO) has been making steady progress in improving the sensitivities of its three interferometers, two in Hanford, Washington, and one in Livingston, Louisiana. These interferometers have reached their target design sensitivities and have since been collecting data in their fifth science run for well over a year.

On the way to increasing the sensitivities of the interferometers, difficulties with increasing the input laser power, due to unexpectedly high optical absorption, required the installation of a thermal compensation system. We describe a frequency resolving wave-front sensor, called the phase camera, which was used on the interferometer to examine the heating effects and corrections of the thermal compensation system. The phase camera was also used to help understand an output mode cleaner which was temporarily installed on the Hanford 4 km interferometer.

Data from the operational detectors was used to carry out two continuous gravitational wave searches directed at isolated neutron stars. The first, targeted RX J1856.5-3754, now known to be outside the LIGO detection band, was used as a test of a new multi-interferometer search code, and compared it to a well tested single interferometer search code and data analysis pipeline. The second search is a targeted search directed at the Crab pulsar, over a physically motivated parameter space, to complement existing narrow time domain searches. The parameter space was chosen based on computational constraints, expected final sensitivity, and possible frequency differences due to free precession and a simple two component model. An upper limit on the strain of gravitational radiation from the Crab pulsar of 1.6×10^{-24} was found with 95% confidence over a frequency band of 6×10^{-3} Hz centered on twice the Crab pulsar's electromagnetic pulse frequency of 29.78 Hz. At the edges of the parameter space, this search is approximately 10^5 times more sensitive than the time domain searches. This is a preliminary result, presently under review by the LIGO Scientific Collaboration.

Thesis Supervisor: Nergis Mavalvala
Title: Associate Professor of Physics

Acknowledgments

I have been lucky to be able to work with an incredible group of people doing what some people said was impossible ten or twenty years ago. During my time at MIT and at the Hanford site I have met, worked with, and learned from too many people to thank properly. However, I'd still like to try to thank a few by name.

To Luca, thanks for showing me the ropes at Hanford.

To Keita, thanks for your all insightful suggestions.

To Fred, thanks for all the words of wisdom.

To Dave, thanks for teaching me the ins and outs of optics.

To Gregg, thanks for teaching me how to make a ringdown measurement.

To Myron, thanks for showing me how to use a machine shop the right way.

To Paul, Josh, and Richard, thanks for sharing your electronics expertise.

To the whole Hanford crew, thanks for showing me how to put together, break, fix, and finally run a LIGO interferometer.

To Mike, Greg, MAP, and the rest of the continuous wave search group, thanks for teaching me how to hear the universe.

To the group at MIT, thanks for sharing your experience and your friendship.

To Nergis, thanks for your unwavering support and guidance.

I'd also like to thank Yoon and the rest of my family for the support and love they have shown me these many years.

- Joe October 3, 2007

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

I	Introduction	15
1	Generating gravitational waves	17
1.1	General relativity and gravitational waves	17
1.1.1	Gravitational wave effects on test masses	18
1.1.2	Generation of gravitational waves	19
1.2	Continuous gravitational wave sources	21
1.2.1	Potential source: RX J1856.5-3754	23
1.2.2	Potential source: the Crab pulsar	25
2	Laser interferometer gravitational wave detectors	27
2.1	Laser interferometers	27
2.1.1	Introduction	27
2.1.2	Laser and mode cleaners	28
2.1.3	Fabry-Perot arms	28
2.1.4	Power recycling cavity	30
2.1.5	AS port	30
2.2	Problems	32
2.2.1	Thermal heating and thermal compensation	32
2.2.2	Mode and power mismatches and the AS_I signal	35
II	Spatial mode analysis	37
3	The phase camera: a frequency sensitive wavefront sensor	39
3.1	Mechanical details	39

3.2	Testing the phase camera and the output mode cleaner	41
3.2.1	Motivation for the output mode cleaner	41
3.2.2	Theory and modeling	42
3.2.3	Phase camera measurements and comparisons	54
4	Improving and applying the phase camera	65
4.1	Phase camera with reference beam	65
4.1.1	Experimental setup	65
4.1.2	Demodulation board redesign	67
4.1.3	Fitting and understanding the data	67
4.2	H1 in different heating states	73
4.2.1	Discussion of phase camera results	74
III	Targeted continuous gravitational wave searches	83
5	Neutron stars, spindown models, and estimating upper limits of sources	85
5.1	Neutron stars and pulsars	85
5.2	Estimated deformations of neutron stars	87
5.3	Energy considerations	89
5.4	Differences between GW and EM pulses	92
5.4.1	Free precession	92
5.4.2	Two component model	97
5.5	Application to RX J1856.5-3754	99
5.6	Application to the Crab pulsar	102
6	Search methods	105
6.1	Search codes and constraints	105
6.2	Detector response function	106
6.3	Signal form in terms of Doppler and amplitude parameters	110
6.4	The F statistic: a coherent search method	112
6.5	Coherent vs incoherent templates and unknown parameters	116
6.6	Exact grid spacing	120
6.7	Interferometer sensitivities	122

6.8	Estimating limits of detection	122
7	Searches and results	125
7.1	RX J1856.5-3754 directed search	125
7.1.1	Parameter space chosen	125
7.1.2	Search results	126
7.1.3	Version 1 coincidence step	128
7.1.4	Upper limit injections	129
7.1.5	Final comparison	130
7.2	Crab wide parameter search	133
7.2.1	Chosen search parameter space and results	133
7.2.2	Upper limits	137
7.2.3	Checks on the search result and upper limits	138
7.2.4	Constraining energy loss in the Crab pulsar	156
IV	Conclusion	160
	Conclusion	161
	Appendices	163
A	Table of acronyms	163
B	Phasecamera demodulation and control board schematic	165

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

1-1	A passing gravitational wave's effects	20
1-2	Neutron star with a non-axisymmetric bump	23
2-1	LIGO interferometer	29
2-2	TCS system	34
3-1	Phase camera and galvo control loop	40
3-2	OMC in relation to the interferometer	48
3-3	Modal decomposition of OMC transmission	51
3-4	Phase map generated from OMC model	53
3-5	Michelson bright fringe on transmission of the OMC	55
3-6	Direct reflection off ITMY on transmission of OMC	56
3-7	Michelson bright fringe directly to phase camera	58
3-8	Direct reflection off ITMX directly to phase camera	59
3-9	Qphase fit of Michelson bright fringe data	61
4-1	Carrier, sidebands and reference beam	66
4-2	Reference beam setup for the phase camera	70
4-3	Carrier signal: horizontal and vertical profiles	71
4-4	Fit of carrier on sideband signal	72
4-5	"Reversed TCS" state images	78
4-6	"No TCS" state images	79
4-7	"Tuned TCS" state images	80
4-8	"Power up" state images	81
5-1	Diagram of angles for precession	93

5-2	Precession for small λ	95
5-3	Precession for small θ	96
5-4	Ellipticity as a function of current frequency	101
5-5	1st Spin down as function of current frequency	102
6-1	Reference angles for definition of interferometer location	107
7-1	RX J1856.5-3754: version 1 largest $2F$ values for H1	127
7-2	RX J1856.5-3754: version 1 largest $2F$ values for L1	127
7-3	RX J1856.5-3754: version 2 largest $2F$ values	128
7-4	RX J1856.5-3754: coincident largest $2F$ values	129
7-5	RX J1856.5-3754: upper limits	131
7-6	RX J1856.5-3754: ratio of $V1/V2$ upper limits	132
7-7	Sensitivity scaling as a function of templates	135
7-8	Confidence vs injection runs	138
7-9	Strain vs SFT for all 3 interferometers	140
7-10	Estimated strain vs frequency for all 3 interferometers	142
7-11	Gaussian noise strain vs SFT	143
7-12	Estimated Gaussian noise strain vs frequency	143
7-13	Expected versus actual number of templates for all 3 interferometers	145
7-14	Expected versus actual number of templates for a Gaussian noise data set	146
7-15	L1: weighted noise from 59 to 60 Hz	148
7-16	H1: weighted noise from 59 to 60 Hz	148
7-17	Distribution of templates with $2F$ greater than 25 for all 3 interferometers	150
7-18	Distribution of templates with $2F$ greater than 25 for Gaussian noise data	151
7-19	Response of the Hanford detector to varying ψ and i	154
7-20	Response of the Livingston detector to varying ψ and i	154
7-21	Confidence versus injected signal strength	155
7-22	Ellipticity versus moment of inertia	157
7-23	Ratio of GW energy loss to total spin down energy loss	158

List of Tables

3.1	OMC parameters	44
3.2	Astigmatic beam decomposition	49
3.3	Fit and model of OMC scan	50
3.4	Amplitude modulation comparisons	60
3.5	Model and fit comparisons for both the Michelson bright fringe and direct reflection	63
4.1	NSPOB and sideband power ratios	75
4.2	Heating states and phase camera power ratios	76
7.1	Crab parameters at the start of S5	136
7.2	Positions of Hanford and Livingston detectors	153
A.1	Acronym definitions	163

Part I

Introduction

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Generating gravitational waves

This chapter describes gravitational waves and considers their generation from isolated neutron stars.

1.1 General relativity and gravitational waves

From the simple experimental observation of the equivalence of inertial and gravitational masses and the desire to incorporate the principle of relativity into a theory of gravity was developed one of the most elegant and successful theories of physics, that of Einstein's General Relativity [20]. It tells us that masses move along geodesics and that the curvature of space-time along which mass moves generates the perceived gravitational force. In turn space-time curvature is generated by the presence of mass and energy. This fact leads General Relativity to predict the propagation of gravitational waves on space-time [40].

Consider the curvature encoded in the space-time metric $g_{\mu\nu}$ where

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu. \quad (1.1)$$

Here ds is the space-time interval or proper time of an observer who is at rest with the coordinate system, dx^μ and dx^ν are the indexed coordinates of space-time and the Einstein summing convention is assumed. From this most general case, let us consider the weak-field limit where space can be considered mostly flat with a small perturbation on top of it, which is a reasonable approximation to human scale laboratories on Earth. In this case the metric $g_{\mu\nu}$ can be written as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.2)$$

where $h_{\mu\nu}$ is the perturbation (and thus holds the gravitational wave information) and $\eta_{\mu\nu}$ is the Minkowski metric of flat space-time equal to

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (1.3)$$

In order for a gravitational wave to propagate, it must satisfy the wave equation

$$\square h_{\mu\nu} = \left(\nabla^2 - \frac{\partial}{\partial t^2} \right) h_{\mu\nu} = 0. \quad (1.4)$$

We neglect the static solutions to the wave equation, as we are interested in radiation which changes with time as opposed to static Newtonian potentials. We are free to choose the coordinate system in which to view the perturbation (gauge choice) and the one in which the math is most clear is the transverse traceless gauge, in which freely falling masses move along coordinate lines. In the transverse traceless gauge with the weak-field assumption, the solution of a wave propagating in the x_3 direction can be written as [40]

$$h_{\mu\nu}(x, t) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos\left(2\pi f_{GW} \left[\frac{x_3}{c} - t\right]\right) \quad (1.5)$$

where h_+ and h_\times are the two possible polarizations of the wave and f_{GW} is the frequency of the gravitational wave.

1.1.1 Gravitational wave effects on test masses

Consider a h_+ polarized gravitational wave acting on a pair of test masses at rest relative to each other. Place both a distance L from the arbitrary origin of their common Lorentz frame, one along the x_1 axis and the other along the x_2 axis. Apply the $g_{\mu\nu}$ metric to find the space-time interval ds between the origin and each of the masses as a function of time.

As a h_+ polarized gravitational wave propagates in the x_3 direction, the ds between the origin and the first mass will vary sinusoidally and be 180° out of phase with the sinusoidal variation of ds between the origin and the second mass.

One method to sample this change in ds is to propagate a laser beam from the origin to each test mass and back. One can integrate along the laser's path from the origin to each test mass to find a measurable quantity, the difference in the light's phase between the two round trip paths:

$$\Delta\Phi = 2\frac{2\pi}{\lambda} \int_0^L \sqrt{|g_{x_2x_2}|} dx_2 - 2\frac{2\pi}{\lambda} \int_0^L \sqrt{|g_{x_1x_1}|} dx_1 \quad (1.6)$$

where λ is the wavelength of the laser light and $\Delta\Phi$ is the difference in total accumulated round trip phase by the laser light along the two paths. If the laser light round trip time is small compared to the gravitational wave period, then the maximum total difference caused by a h_+ wave can be approximated by

$$\Delta\Phi \simeq 2h_+ \frac{2\pi L}{\lambda}. \quad (1.7)$$

One can do the same calculations with the h_\times polarization or any linear combination of the two.

1.1.2 Generation of gravitational waves

To gain insight into the generation of gravitational waves, we look to electromagnetic theory as a guide. Electromagnetic waves are governed by the same wave equation except that the electromagnetic field \mathbf{E} replaces the perturbation $h_{\mu\nu}$. If we restrict ourselves to waves whose wavelength is much larger than the source we can use a multipole expansion. In the same way acceleration of charge is necessary for electromagnetic waves, acceleration of mass (and thus second derivatives of the moments) are necessary for gravitational waves. The monopole of the expansion for both electromagnetic waves and gravitational waves does not radiate due to the conservation of charge on the one hand and the conservation of mass and energy on the other. We can move on to consider the gravitational equivalent to the electric dipole [40]

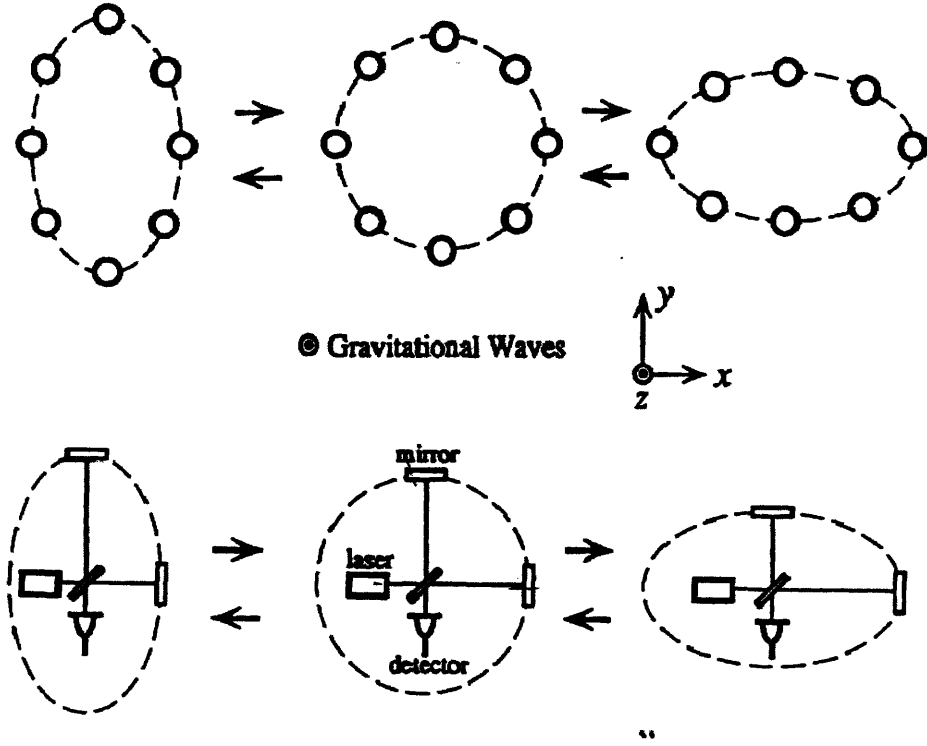


Figure 1-1: Effects of a passing gravitational wave on free masses and a laser interferometer.

$$\mathbf{d}_g = \int \rho(\mathbf{r}) \mathbf{r} dx_1 dx_2 dx_3 . \quad (1.8)$$

However, it is easy to see that the second derivative of \mathbf{d}_g will be zero because of the conservation of momentum ($m\dot{\mathbf{x}} = \text{constant}$). The next higher moments are equivalent to the magnetic dipole and the electric quadrupole. The gravitational equivalent to the magnetic dipole is

$$\mu_g = \int \rho(\mathbf{r}) \mathbf{r} \times \mathbf{v}(\mathbf{r}) dx_1 dx_2 dx_3 . \quad (1.9)$$

Again in this case there is no radiation, this time due to the conservation of angular momentum ($m\dot{\mathbf{x}} \times \mathbf{r} = \text{constant}$). Finally, with the quadrupole moment we lack any more conservation laws to apply and thus have the first term in the expansion that can generate radiation. The typically quoted reduced gravitational quadrupole moment is

$$I_{\mu\nu} = \int \rho(\mathbf{r}) dx_\mu dx_\nu - \frac{1}{3} \delta_{\mu\nu} r^2 dx_1 dx_2 dx_3 \quad (1.10)$$

Thus in direct analogy to electromagnetism's potential [35]

$$A_\mu = \frac{\dot{d}_\mu}{R} \quad (1.11)$$

where d_μ is a charge dipole moment and R is the distance from the dipole to the point where the potential is measured at, we can write

$$h_{\mu\nu} = \frac{2}{R} \ddot{I}_{\mu\nu} . \quad (1.12)$$

Note that we are working in units where the gravitational constant, G , and the speed of light, c , are set to 1. The factor of 2 comes from the correct tensor calculations which are presented in detail in [40].

1.2 Continuous gravitational wave sources

For the majority of this thesis we will focus on one type of source of changing quadrupole moment, a non-axisymmetric isolated neutron star. One can imagine such a star with a bump on its surface or an equatorial bulge producing an eccentricity as shown in Figure 1-2. Consider a non-axisymmetric neutron star with the principle moments I_{11} , I_{22} , I_{33} . Let us assume it is rotating about the x_3 axis with frequency f_{ROT} . In this case the second derivative of the quadrupole moment tensor $I_{\mu\nu}$ is [56]

$$\ddot{I}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -16\pi^2 f_{ROT}^2 (I_{11} - I_{22}) \cos(4\pi f_{ROT} t) & -32\pi^2 f_{ROT}^2 (I_{11} - I_{22}) \sin(4\pi f_{ROT} t) & 0 \\ 0 & -32\pi^2 f_{ROT}^2 (I_{11} - I_{22}) \sin(4\pi f_{ROT} t) & 16\pi^2 f_{ROT}^2 (I_{11} - I_{22}) \cos(4\pi f_{ROT} t) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} . \quad (1.13)$$

$I_{11} - I_{22}$ can be replaced in the above with ϵI_{33} by using the definition of equatorial ellipticity

$$\epsilon = \frac{I_{11} - I_{22}}{I_{33}} . \quad (1.14)$$

Using this $\ddot{I}_{\mu\nu}$ in Eq. (1.12) and choosing to observe the generated wave from a new coordinate system with the x_3 axis oriented at an angle i away from the axis of rotation we get [56]

$$h_{\mu\nu} = \frac{32\pi^2 f_{rot}^2 \epsilon I_{33}}{R} \times \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\cos(4\pi f_{ROT}t)(1 + \cos^2(i)) & -2\sin(4\pi f_{ROT}t)\cos(i) & 0 \\ 0 & -2\sin(4\pi f_{ROT}t)\cos(i) & \cos(4\pi f_{ROT}t)(1 + \cos^2(i)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1.15)$$

This can be compared to Eq. (1.5) which allows one to read off h_+ and h_\times and also note that $f_{GW} = 2f_{ROT}$.

From the above Eq. (1.15) it is easy to determine what properties a detectable continuous wave source would have. First of all, it must have some eccentricity. However there are theoretical estimates of the maximum possible eccentricities supportable by neutron stars. These estimates are based off of possible equations of state for neutron stars, which are equations which define the relationship between the temperature, pressure and density within the star. By surveying the more plausible equations of state one finds a maximum ellipticity of $\sim 10^{-6}$ [51] can be supported. Exotic equations of states for things such as hybrid or quark stars, where the core is partially or completely made up of quark matter, can have maximum ellipticities a factor of 100 larger. However, it should be noted that such equations of state are considered by many astrophysicists to be speculative, and so the ellipticities we infer from them may not be quite as plausible [42]. A more in-depth discussion of ellipticities of neutron stars is covered in Section 5.2.

Secondly, the source needs to be rotating rapidly, due to the f^2 dependence of the strain. It also needs to be within the LIGO detection band, which has a low end around 50 Hz determined by seismic noise. Known pulsars have a maximum spin periods of a few milliseconds, yielding maximum gravitational wave frequencies of about 1 kHz. In the case of an isolated neutron star this means it needs to be young and not have radiated away

too much of its angular momentum. Thirdly, an ideal source should be as close as possible since the strain at a detector is inversely proportional to the distance to the source.

By considering the known properties of neutron stars we can now make a quick estimate of the maximum strength of gravitational waves reaching detectors on Earth. Typical radii are of order 10 km and typical masses are around $1.4 M_{\odot}$ [40]. This leads to a typical moment of inertia 10^{45} g cm^2 assuming a rigid rotator. Using 1 kpc as a distance standard and simply note any change in distance just produces an inverse change in the detect strain. The above equations can be converted from natural units to experimentalist units by including correct factors of G and c . Using the typical maximum values for frequency and ellipticity yields a GW amplitude of the order $h \sim 10^{-24} (1 \text{ kpc}/R)(\epsilon/10^{-6})(f/500\text{Hz})^2$ detected at Earth from a source 1 kpc away.

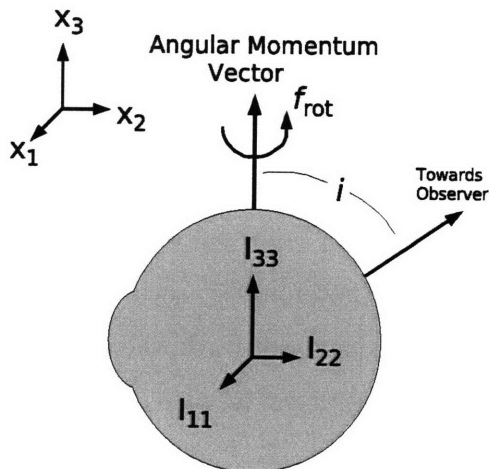


Figure 1-2: Diagram of a non-axisymmetric neutron star with a bump. The angle i is the inclination angle between the axis of rotation and the direction to the detector.

Next we consider two isolated neutron stars that could make good candidates for detection by LIGO.

1.2.1 Potential source: RX J1856.5-3754

A potential candidate that meets several of the previously listed criteria is the isolated neutron star RX J1856.5-3756. RX J1856.5-3754 is the closest known neutron star, at a distance of ~ 120 pcs [31] [53], as determined by parallax measurements. While it is seen in the optical by the Hubble space telescope, it was actually discovered in the X-ray by the Rosat All Sky Survey in 1996 [54]. It does not emit detectably at radio wavelengths,

and thus was not in the original set of pulsar candidates being examined by LIGO. It has a remarkably featureless x-ray spectrum which is well fit to a black body spectrum of temperature 63.5 eV, although there is an excess in the UV-optical region of the spectrum by about a factor of 7 from just the x-ray data fit [13]. A naive black body fit to the x-ray data suggests that the star is of the order 4km in radius while a similarly naive black body fit to the optical data leads to a star of radius of 14km [13]. The extreme mismatch between these two fits indicates that one must model the effects of the neutron star's atmosphere to understand the electromagnetic emission. At the time of considering RX J1856.5-3754 as a source Chandra and XMM observations had not been able to detect any periodic variation in the x-ray emission, having placed an upper limit of 1.5% on the pulsing fraction of x-rays relative to the total x-ray flux [12].

A lack of spin period to focus on requires a blind search, although there are good estimates of its age allowing us to limit the parameter space of the search, when coupled with spin down energy calculations. Since the star is losing rotational energy as it radiates away gravitational radiation, an equation relating the current rotation frequency and the rate of change of that rotation frequency (the spin down rate) can be written. If sufficient time has passed since the birth of the neutron star, such that the current rotation frequency is small compared to the rotation frequency at birth, the ratio between the current rotation frequency and the current spin down rate will become a function only of the age of the star. If the age is known then at any given frequency, spin downs only up to a certain maximum need be considered, limiting the spin down parameter space necessary to be searched. This is covered in more depth in Section 5.3.

Age estimates for RX J1856.5-3754 come from considering its temperature and how rapidly it cools, the lack of a visible super nova remnant, and possible associations with other stars from a common starting point. Its high temperature and x-ray brightness implies that RX J1856.5-3754 must be around 10^6 years old or younger, since the star has not lived long enough to cool too much by thermal emission. On the other hand, there is no visible supernova remnant, which implies an age greater than about 10^5 years, the length of time most supernova remnants remain visible. RX J1856.5-3754's proper motion of ~ 0.33 mas yr $^{-1}$ suggests that $\sim 5 \times 10^5$ years is correct. Tracing RX J1856's trajectory back in time finds it crossing the trajectory of run-away star ζ Oph 5×10^5 years in the past [52]. They both may have originated in the Upper Sco OB association, where their paths meet,

before the supernova of the progenitor of RX J1856.5-3754 kicked them out.

While carrying out a gravitational wave search directed at RX J1856.5-3754 it was found to have a spin period below the LIGO detection band by X-ray astronomers. It has a period of 7 seconds with only a variation in the X-rays of 1.5%, the smallest ever seen in an isolated X-ray pulsar [50]. However the methodology and resulting test of the code and search pipeline performed using it as a target are still of use and thus it is presented here and in Chapter 7.

1.2.2 Potential source: the Crab pulsar

The Crab pulsar and its surrounding nebula are remnants of a supernova that occurred in 1054 AD approximately 2 kpc distant [17]. Its been well studied over a broad band of electromagnetic frequencies from radio to gamma rays [38]. It produces very regular pulses across the electromagnetic spectrum whose frequency and time evolution are known to high accuracy. These electromagnetic wave pulses provide insight into the rotation rate and spin down rate of the Crab pulsar and also act as a starting point from which to narrow down directed searches for gravitational waves from the Crab pulsar.

The Crab has a known radio pulse frequency of ~ 30 Hz which implies the rotation frequency is close to that. It also known to be spinning down at the rate of $\sim 3 \times 10^{-10}$ Hz/s [38] [28]. This information lets us calculate a naive classical upper limit on the strain at an earth based detector of 1.4×10^{-24} and upper limit of 7.5×10^{-4} on the ellipticity of the Crab pulsar [43]. Thus the Crab is a good potential source because of its extreme youth and rate of spin down, despite its distance. These radio observations also let us narrow the parameter space down, but there is still some uncertainty left in how closely the gravitational radiation period matches the radio pulse period [56].

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Laser interferometer gravitational wave detectors

This chapter will summarize the basics of a laser interferometer detector. It will then discuss heating problems with certain optics which arose while trying to reach the designed input laser power and how these problems were addressed with a thermal compensation system. An in depth examination of the laser light inside the interferometer from before and after this compensation system was used can be found in Chapter 4.

2.1 Laser interferometers

2.1.1 Introduction

One of the simplest and most sensitive methods for measuring the difference in length along two paths is a Michelson interferometer. Figure 2-1 shows the full Michelson interferometer combined with the two Fabry-Perot arms, discussed in Section 2.1.3 and a power recycling cavity, discussed in Section 2.1.4. This setup is the heart of all three LIGO detectors. These interferometers have been undergoing construction and commissioning for the past decade and recently have gone through several periods of designated data collection. These periods, during which the interferometer is left to simply run without constant upgrades and commissioning, are referred to as science runs. The most recent science run, the 5th so far, started in November 2005 and was the first science run in which the interferometers were operating at their designed sensitivity level.

At its most basic, a LIGO interferometer consists of a laser light source which is then directed through a 50/50 beamsplitter. The transmitted and reflected beams travel along perpendicular paths until reaching two end test masses (ETMs), which are highly reflective mirrors. The reflected light from each mirror is then recombined at the beamsplitter. The recombined electric field at the anti-symmetric (AS) port depends on the difference in the optical paths of the two beams. The electric field at the AS port for this simple configuration can be written as

$$E_{AS} = E_{in} \left(r_{ex} t_{bs} r_{bs} e^{i2\phi_x} - r_{ey} t_{bs} r_{bs} e^{i2\phi_y} \right) \quad (2.1)$$

where E_{in} is the electric field incident on the beamsplitter from the laser source, ϕ_x and ϕ_y are the phases accumulated by the light in a one way trip down each arm, t_{bs} and r_{bs} are the field transmission and reflection of the beamsplitter, and r_{ex} and r_{ey} are the field reflection of the end mirrors [3]. In reference to the earlier calculations in Section 1.1.1, the $\Delta\Phi$ of Eq. (1.6) would be equal to $2(\phi_y - \phi_x)$.

2.1.2 Laser and mode cleaners

The main laser is a 10 Watt Nd:YAG laser from Lightwave, which operates at a wavelength of 1064 nm. It is both frequency and intensity stabilized. The laser output is passed through a pre-mode cleaner (PMC). This is a 21 cm long triangular cavity designed to filter the spatial mode and intensity noise above about 1 MHz. It then passes through the mode cleaner, another triangular cavity, with a 24.492 meter round trip. This cavity spatially filters the light, only allowing a single mode to pass, while higher order modes pick up additional phase and fall out of resonance. It also acts as a stable angular reference, as misalignments become translated to higher order modes within the cavity and also fall out of resonance [3].

2.1.3 Fabry-Perot arms

To increase the sensitivity of the basic Michelson interferometer, one would like to increase L , the length of the arms, up until the length is comparable to the wavelength of the gravitational waves, roughly 300 to 6000 km. At that point the approximations leading up to Eq. (1.7), which shows the signal grows linearly with L , break down. Effectively,




Figure 2-1: A diagram of a LIGO interferometer (taken from [3]). The purple light represents the input laser light. The blue line represents the resonant sidebands in the power recycling cavity while the red line represents the carrier signal resonant in the recycling cavity and the Fabry-Perot arms, as discussed in Sections 2.1.3, 2.1.4, 2.1.5.

the light is picks up phase shifts from multiple cycles of the wave which cancel each other, preventing any further gain in sensitivity. There are technical problems with building the arms even 300 km long. One way to effectively make the arms longer is to have the laser light make multiple trips down the same arm. As long as the time spent by the light in the cavity is small compared to the period of the gravitational waves, this will significantly increase the sensitivity of the instrument. However for changes occurring faster than this storage time, the resulting signal will be attenuated relative to the lower frequencies in the same configuration.

By turning the Michelson interferometer arms into resonating Fabry-Perot cavities, one can increase the number of round trips the laser light makes. This adds technical complexity in the form of an additional set of partially transmitting input test masses (ITMs) and the

necessary feedback control systems to keep the Fabry-Perot cavities on resonance.

2.1.4 Power recycling cavity

In the case of the Michelson interferometer, the light interferes destructively in the direction of the AS port and constructively in the direction of the incoming laser light. Instead of just dumping the outgoing power, one can construct a cavity to take the outgoing light and recycle it back into the interferometer. By adding a partially transmitting recycling mirror between the laser and the beamsplitter a cavity can be made which recycles the power back into the interferometer. This cavity also requires additional feedback control systems to remain on resonance and generally provides a factor of 50 gain to the carrier, and a factor of 26 gain to the sidebands when the input test masses are thermally tuned [5]. The recycling mirror transmission was chosen to match the total round trip losses by the carrier in the rest of the interferometer so as to couple all the carrier light into the interferometer [3].

The reason for using all available power is because of shot noise. The interferometer operates on a dark fringe and thus very few photons carry the gravitational wave signal to the AS port. Shot noise is due to the quantum fluctuations in the number of photons reaching the detector. Shot noise scales as \sqrt{N} , where N is the number of photons. However, the gravitational wave signal scales with the power, and therefore the number of photons, and so the signal to shot noise ratio scales as \sqrt{N} [3]. By increasing the total power we increase the interferometer sensitivity at the frequencies which are limited by shot noise.

2.1.5 AS port

The electric field at the AS port is detected with a photodiode, which produces a photoelectric current proportional to the average photon flux, or power, on the detector. For a simplified calculation of the power at the AS port we assume the end mirrors have identical reflectivity and are equal to 1 (i.e. $r_{ex} = r_{ey} = r_e = 1$) and that the beamsplitter is a perfect 50/50 (i.e. $r_{bs} \times t_{bs} = 1/2$). Using Eq. (2.1) we can then write the AS port power as

$$P_{AS} = E^* E = 4|E_{in}|^2 (r_e t_{bs} r_{bs})^2 \sin(\Delta\phi)^2 = |E_{in}|^2 \sin(\Delta\phi)^2 \quad (2.2)$$

where we define $\Delta\phi = \phi_y - \phi_x$ [3]. To produce signals that are at radio frequencies on the photodiodes, the light entering the interferometer is first phase modulated to produce several sets of sidebands. These sidebands are used in a Pound-Drever-Hall locking scheme [18] (see also [23]), to keep the different cavities on resonance (or locked). One set of sidebands is reflected by the mode cleaner, one set is reflected by the power recycling mirror, and the last set is reflected by the input test masses, and resonate in the recycling cavity. This last set of sidebands which resonates in the power recycling cavity is used to beat against the carrier and produce the final gravitational wave signal channel at the AS port. An asymmetry of 0.356 m in the lengths between the BS and the two ITMs, called the Schnupp asymmetry, causes the sidebands to come out the AS port. Nominally, the recycling cavity should be almost critically coupled for the sidebands such that almost all the sideband light should end up at the AS port.

These sidebands are produced by using a Pockels cell to phase modulate the input light. The laser light modulated once for one set of sidebands can be written as

$$\begin{aligned}
 E_{in} &= E_0 e^{i\Gamma \cos(\Omega_m t)} \\
 &\cong E_0 (J_0(\Gamma) + iJ_1(\Gamma)e^{+i\Omega_m t} + iJ_1(\Gamma)e^{-i\Omega_m t} \\
 &\quad - J_2(\Gamma)e^{+i2\Omega_m t} - J_2(\Gamma)e^{-i2\Omega_m t})
 \end{aligned} \tag{2.3}$$

$$\tag{2.4}$$

where Γ is the modulation depth in radians, Ω_m is the modulation angular frequency, J_n is the n^{th} order Bessel function of the first kind, and E_0 is the unmodulated field of the laser. In the case of the sidebands resonant in the power recycling cavity, $\Omega_m \approx 2\pi \times 24.5$ MHz and $\Gamma \approx 0.4$. The output of the photodiode is demodulated at the modulation frequency used to run the Pockels cell. The In-phase (cosine) and Quadrature-phase (sine) of the AS photodiode are separately recorded.

2.2 Problems

2.2.1 Thermal heating and thermal compensation

The ITMs were designed with the concept that a certain amount of absorption would be taking place leading to a curvature change due to thermal lensing, effectively the changing index of refraction due to increasing temperature, when operating in a high input power (6 Watt) state. The power gain of 50 in the recycling cavity and gain of 130 in the arms means roughly 300 W and 10 kW are circulating in the recycling cavity and arms, respectively, in this high input power state. The expected absorption was about 1 ppm in the HR coating and about 4-5 ppm/cm in the substrate. Through the initial science runs, during which the LIGO interferometers were left alone to take data as opposed to being worked on directly to improve sensitivity, the LIGO interferometers were operating below this expected high power state and consequently had a slightly wrong radius of curvature for the ITMs and ETMs.

However, during attempts to reach high power on the H1 interferometer, it became clear that optics were absorbing far more than expected. At their nominal radius of curvature at 6 Watts of input power, the gain in power in the sidebands is theoretically expected to be $G_{sb} = 30$ [5]. However, a maximum gain of $G_{sb} = 26.5$ was measured at 1.8 Watts of laser power into the mode cleaner, a much lower power than expected. At higher power, the sideband power gain reduced along with the sidebands decreasing in spatial size. This suggested that one or more optics of the power recycling cavity was absorbing too much power.

Using spot size measurements on H1 at different IFO ports and at different heating states [41], we were able to determine that ITMX was absorbing the most power, 34 ± 4 mWatt per Watt of power into the mode cleaner, although ITMY was only a factor 2.6 less than that. These in vacuum measurements were unable to distinguish between bulk substrate absorption or coating absorption. However, since the power hitting the first coating layer on the high reflectivity side of the ITM, which faces the Fabry-Perot arm cavity, is roughly 140 times the incident power, and the higher probability of surface contaminants, the absorption is most likely due to the surface rather than the bulk. If one assumed all the absorption were in the surface, then the ITMX coating would be absorbing 15 ± 1.8 ppm and the ITMY coating would be absorbing 5.6 ± 0.7 ppm. This is compared to the specification of 1 ppm.

One hypothesis is that contaminants were introduced to the surface during assembly and installation and resulted in this large difference from specification to the actual absorption. However, we don't have a definitive cause of the excess absorption. Based on these findings, ITMX was replaced with a spare in June of 2005. Since the ITMs still had absorption values above the design specifications, a system was necessary to compensate for the thermal heating effects of the main laser.

A thermal compensation system (TCS) was originally studied by Ryan Lawrence [36] and was designed to apply heating to the test masses as needed to correct the curvature. This eventually led to the full system installed on the LIGO interferometers [48] [6]. The final systems installed for each ITM, shown in Figure 2-2, consist of a 10 Watt CO_2 laser operating at $10.6 \mu\text{m}$ wavelength that illuminates one of two interchangeable mask patterns. The laser light passes through a telescope with magnification of 26.5 and a ZnSe viewport before reaching the high reflectivity side of the ITM (the side facing the arm cavity) of the ITM. The $10.6 \mu\text{m}$ laser light is almost completely absorbed by the fused silica substrate of the ITM.




Figure 2-2: A schematic of the Thermal Compensation System (TCS). This was installed on both ITMs on all 3 LIGO interferometers. It is made up of a 10.6 μm CO₂ laser, a set of interchangeable masks, and a series of projection optics that produce an image of the mask on the optic.

The two masks are an annulus pattern and a hole. The annulus is used to compensate for too much heat being deposited by the main LIGO laser light. The hole in the center is used to correct for insufficient heat to reach the desired curvature. These masks are installed in the Fourier plane of the projection system. A Bessel mask further downstream clips the higher order maxima of the Airy diffraction pattern, leaving only the central lobe of the Airy disk. A polarizer on a rotation stage is used to adjust the power.

This system was successfully able to correct for the high power heating effects and allowed the LIGO interferometers to reach operation at 6 Watt input laser power state. To control the long time dependence of the TCS on the radius of curvature, a servo system was devised and implemented. The full details of this system can be found in Stefan Ballmer's thesis [5].

2.2.2 Mode and power mismatches and the AS_I signal

The carrier light resonates mostly in the 4 km long arms of the interferometer, and thus their spatial mode is mostly dependant on the state of the arm cavities. The sidebands only resonate in the power recycling cavity and thus the sideband structure is determined by the state of the power recycling cavity. The deformation and thermal lensing of the ITMs, due to their excess absorption, affects primarily the power recycling cavity. This causes the two sidebands' power and spatial profiles to become mismatched with the carrier and each other. Since the gravitational wave signal is detected by the beat between the carrier and sidebands, poor matching will lead to excess noise.

To illustrate what poor matching can do, we write the incident light field on the AS photodiode as

$$\Psi = A_c e^{-i\omega_c t} + iA_+ e^{-i(\omega_c + \Omega_m)t} + iA_- e^{-i(\omega_c - \Omega_m)t} \quad (2.5)$$

where Ω_m is the sideband modulation frequency, ω_c is the carrier frequency, and A_c, A_+, A_- contains the modal content and the relative magnitudes of the carrier, upper sideband, and lower sideband respectively. When the photodiode signal is demodulated at frequency Ω_m , two signals are produced, the In-phase signal and the Quadrature-phase signal. Since these are recorded by the AS port photodiode, these are labeled AS_I and AS_Q. The interferometer is designed such that gravitational wave signals are only in AS_Q and the

AS_I signal is zero.

We can write the power in each demodulated signal as

$$P_{AS_I} = \int_0^{\frac{2\pi}{\Omega_m}} \Psi \Psi^* \cos(\Omega_m t) = (A_c A_{sb+}^* + A_c^* A_{sb-}) - (A_c^* A_{sb+} + A_c A_{sb-}^*) \quad (2.6)$$

and

$$P_{AS_Q} = \int_0^{\frac{2\pi}{\Omega_m}} \Psi \Psi^* \sin(\Omega_m t) = (A_c A_{sb+}^* + A_c^* A_{sb-}) + (A_c^* A_{sb+} + A_c A_{sb-}^*) \quad (2.7)$$

where P_{AS_I} and P_{AS_Q} are the powers in the AS_I and AS_Q signals.

Note that when the sidebands have identical amplitude and modal content, namely $A_{sb+} = A_{sb-}$, P_{AS_I} goes to zero, leaving only AS_Q with signal. However, if the overall magnitude is different, effectively amplitude modulation of the sidebands, P_{AS_I} no longer identically cancels. In addition if the modal content of the sidebands are not identical, and there are higher order modes present to beat with in the carrier, P_{AS_I} will again not cancel. The higher order modes in the carrier and sidebands will also add noise directly to the AS_Q signal as well.

If the mismatches between the carrier and sidebands are severe enough, the excess AS_I signal can cause saturation of the RF electronics and mixers, resulting in loss of any gravitational wave signal [3]. There is an electronic system in place to provide feed-back that cancels out RF photo current on the diode caused by the AS_I signal. This system can correct up to 10 mA_{pk} photo current per diode. The control signal for this system is recorded in the channel AS_I_Corr. However in high power configurations this correction system proved to be insufficient [5]. The excess AS_I problem was eventually solved with the thermal compensation system running a correction servo designed to null the AS_I_Corr signal. It is important to examine the power and modal nature of the individual components of the light in the recycling cavity with and without TCS correction, to ensure that it is correcting properly. This examination is accomplished through the use of the phase camera, a system whose basics and initial test case are described in Chapter 3. The examination of the thermal states and TCS correction of the interferometer is covered in Chapter 4.

Part II

Spatial mode analysis

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

The phase camera: a frequency sensitive wavefront sensor

The “phase camera” is a high-resolution wavefront sensor that measures the complete spatial profile of any frequency component of a laser field containing multiple frequencies. This technique is useful as a probe of the spatial overlap of the carrier field with each sideband component exiting the output port of a gravitational wave interferometer. It was first demonstrated to work at MIT [25] and then brought to the LIGO sites.

This chapter describes the functioning of the phase camera at a basic level, where it detects only the beat signal between the carrier and sidebands. At the Hanford site, the phase camera was first successfully used to test an output mode cleaner and confirm that it was working as designed. That test is presented here to show the phase camera working in a relatively simple to understand case on the actual interferometer. The addition of an independent reference beam, described in Section 4.1.1, was necessary for the more complex study of individual sideband and carrier components in different heating states described in Chapter 4.

3.1 Mechanical details

The phase camera is the combination of a position scanning galvanometer and a New Focus 1811A RF photodiode. The galvanometer moves two mirrors in conjunction such that the incoming light reflects off both in a spiral pattern and onto the photodiode at a rate of about one full spiral scan every half second. The two output channels of the photodiode,

DC power and RF power, are then sampled at a rate of up to 4000 points per full scan. The RF sampled points are then demodulated and passed to a computer, along with the position read back of the galvanometer as shown in Figure 3-1. The demodulation can be done at a frequency of 24.5 MHz, the difference between the carrier and sideband frequencies, and thus looks at the beat between the carrier and sidebands of the incoming light. A reference beam can be combined with the incoming light to produce beats with the carrier and sidebands individually, allowing their individual demodulation, and which is described in Section 4.1

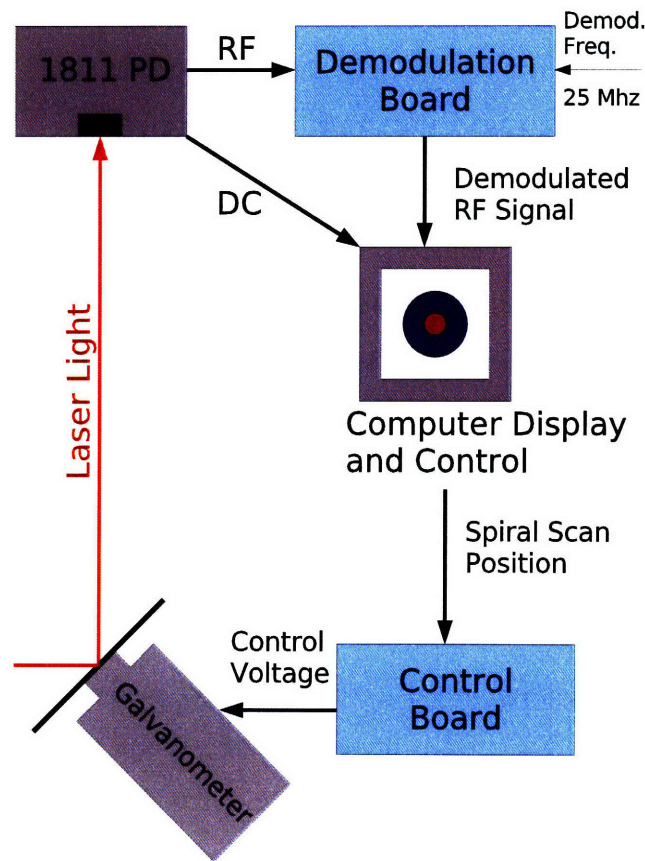


Figure 3-1: Phase camera control setup. The galvanometers move a pair of mirrors, directing the light in onto an 1811A RF photodiode. The RF signal from the photodiode, produced when the carrier and RF sidebands beat at 24.5 MHz, is demodulated, and then sent to a computer, along with the DC signal, where the signals are recorded. The computer sends the signal for the next position to a digital to analog converter board, which converts the signal into a control voltage which moves the mirrors to the next point in the spiral pattern.

3.2 Testing the phase camera and the output mode cleaner

The phase camera was used to help test and understand an output mode cleaner (OMC) on the Hanford 4km interferometer [10]. This also presented an opportunity to test the phase camera itself, to ensure that it was operating as expected.

3.2.1 Motivation for the output mode cleaner

Laser interferometers have problems associated with imperfect interference at the AS port, due to poor beam quality. Excess carrier light in higher order modes will produce excess shot noise while not contributing to the signal. Imperfect matching in the reflectivity of the two arms will produce a static, TEM_{00} carrier field at the AS port, which when it beats with the sideband light will produce noise in the signal channel. Any sideband field which does not spatially overlap with the carrier does not contribute to increasing the optical gain, but rather produces many detrimental effects like shot noise, acoustic sensitivity and so forth [3].

A potential solution to poor output beam quality is to place an output mode cleaner in the path of the AS port light before it reaches the AS photodiodes. By using a cavity with sufficiently large line width, both the RF sidebands and carrier TEM_{00} modes could be transmitted on the same resonance, while passively rejecting the higher order spatial modes. The more recently adopted solution is for future upgrades to the LIGO interferometers is to use only carrier light, without RF sidebands, a so called DC readout.

Experimental setup

The OMC was placed on the AS port of the Hanford 4k Interferometer as shown in Figure 3-2. This OMC was a solid spacer triangular cavity with a piezo-electric transducer (PZT), a device which converts electrical field into length changes, attached to one mirror for length control. It was designed to pass the TEM_{00} carrier and sidebands through the same resonance, while rejecting all the higher order modes. The cavity was held on resonance by applying a small dither signal to the OMC length at 40 kHz, and using a lock-in amplifier to generate an error signal from the transmitted light. A few useful parameters of the OMC are shown later in Table 3.1.

Incoming light

We have modeled and taken data for the OMC in the two simplifying cases of the laser reflected directly off an ITM onto the AS port, and of the interferometer being held on a Michelson bright fringe with the RM and ETMs misaligned. In these cases we assume that the modal content of the incoming carrier and sideband is essentially the same. However, a beam scan of the light on the table showed the light to be astigmatic, complicating the situation. The light had a first beam waist of $137\mu\text{m}$ along an axis rotated 25° counterclockwise relative to the OMC horizontal axis. The second beam waist of $107\mu\text{m}$ located 6.6 cm after the first waist was along an axis rotated 25° degrees relative to the OMC vertical axis. Initially, the OMC was mode matched to only the horizontal waist, treating that as the only waist. However, during the testing process it was moved such that the transmission through the OMC was maximized. All data present here is from this later situation.

Phase camera

In this experiment a 50/50 beamsplitter was placed at the output of the OMC, redirecting half of the transmitted light to the phase camera. On the path to the phase camera, shown in Figure 3-2, two lenses were placed so as to focus the light to the proper size to be scanned. This means the phase camera photodiode was approximately 0.45 m from a waist size of $150\mu\text{m}$, resulting in a Gouy phase of nearly $\pi/2$ radians. The Gouy phase is the phase difference between a Gaussian beam and a plane wave of the same optical frequency, acquired by the Gaussian beam as it propagates through a focus. This effective phase shift is different for higher order modes,

3.2.2 Theory and modeling

In the following sections a model for the astigmatic light from the interferometer reaching the OMC, and then passing through the OMC, will be developed. This requires defining a basis of modes affected by astigmatism, then transforming it into the OMC basis. By comparing the relative power in the modes actually measured at the output of the OMC with what we can determine if the OMC was built properly and performing as we expect.

Astigmatic beam

The Hermite Gaussian basis provides a complete set of solution for the fields which can propagate inside or outside of the cavity [33]. These solutions for the electromagnetic field, Ψ , which have been normalized for power, have the form

$$\begin{aligned} \Psi_{mn} = & \sqrt{\frac{2}{2^{(m+n)} m! n! \pi}} \frac{1}{w(z)} H_m \left(\frac{\sqrt{2}x}{w(z)} \right) H_n \left(\frac{\sqrt{2}y}{w(z)} \right) \\ & \times e^{(x^2+y^2)/w^2(z)} e^{-ik(x^2+y^2)/(2R(z))} e^{ikz} e^{i(m+n+1) \tan^{-1}[(z\lambda)/(\pi w_0^2)]} \end{aligned} \quad (3.1)$$

where m and n are TEM mode numbers, x and y are the distance from the center of the beam perpendicular to the direction of propagation, z is the distance from the waist in the direction of propagation, $w(z)$ is the radius of the beam at z , w_0 is the radius at the waist, $R(z)$ is the radius of curvature of the beam front at z , λ is the wavelength of light, k is the wave number ($2\pi/\lambda$), and H_m is the Hermite polynomial of order m .

This can be generalized for the incoming light by breaking up the symmetry between the x and y axes and allowing them to have different waists and waist positions. In the case of a Michelson locked on a bright fringe and also a direct reflection of the laser light off an ITM, the light should be a pure TEM₀₀, distorted only by the astigmatism. Thus, the incoming light can be written as

$$\begin{aligned} \Psi_{astigmatic}(x', y', z') = & \sqrt{\frac{2}{2^{(m+n)} m! n! \pi}} \sqrt{\frac{1}{w_x(z_x)}} \sqrt{\frac{1}{w_y(z_y)}} \\ & \times H_m \left(\frac{\sqrt{2}x}{w_x(z_x)} \right) H_n \left(\frac{\sqrt{2}y}{w_y(z_y)} \right) \\ & \times e^{-[x/w_x(z_x)]^2 - [y/w_y(z_y)]^2} e^{ik[(z_x+z_y)/2]} \\ & \times e^{i[m+(1/2)] \tan^{-1}[(z_x\lambda)/(\pi w_{0x}^2)]} e^{i[n+(1/2)] \tan^{-1}[(z_y\lambda)/(\pi w_{0y}^2)]} \end{aligned} \quad (3.2)$$

where z_x and z_y are the distances from the x aligned waist w_{0x} , and the y aligned waist w_{0y} , respectively.

The OMC basis

The OMC cavity imposes several boundary conditions on the general Hermite Gaussian solutions of Eq. (3.1). First, for resonance to occur, after each full round trip the phase and amplitude must match. Second, at the curved back mirror the wave front of the beam must match that of the cavity (neglecting the small incidence angle). Also note that, because of symmetry, the waist of the beam inside the cavity must occur midway between the two flat mirrors and equidistant from the curved back mirror. These conditions imply

$$w_0^4 = \left(\frac{\lambda}{\pi}\right)^2 L(R_0 - L) \quad (3.3)$$

and

$$\frac{c}{\lambda\nu_0} = (q + 1) + \frac{1}{\pi}(m + n + 1) \cos^{-1} \left(\sqrt{1 - \frac{L}{R_0}} \right) + \frac{1}{2} \text{Mod}_2(m) \quad (3.4)$$

where L is half the round trip length, R_0 is the radius of curvature of the back mirror, Mod_2 is modulo base 2, and ν_0 is the free spectral range (FSR). The FSR is equal to $c/2L$, where c is the speed of light, and is simply the distance in frequency space between successive transmission peaks of a given spatial mode. These transmission peaks are numbered by the q parameter, since $q+1$ is the number of half wavelengths inside the cavity. The modulo base 2 term comes from the fact that the cavity is triangular, and that for odd mode numbers in the direction parallel to the mode cleaner plane there is an additional π phase shift in the round trip (due to an odd number of reflections in the horizontal plane). Table 3.1 provides the values for these and several other related parameters of the OMC.

Table 3.1: OMC parameters

Cavity Finesse	30
Half of Round Trip Path (L)	50.75 mm
End Mirror Radius of Curvature (R_0)	75 mm
Free Spectral Range (ν_0)	2.95 GHz
Beam waist inside the OMC (w_0)	109 μm
Wavelength of Carrier light (λ)	1064 nm

The important part of Eq. (3.4) is that for a given frequency of light, each $(m + n)$ needs a different length of cavity to resonate, due to the Gouy phase shift. This means the OMC will decompose the incoming light into the Hermite Gaussian basis specified by its

geometrical parameters, and only let those modes whose $(m + n)$ satisfy Eq. (3.4) to pass, generally the carrier TEM₀₀ mode. All others will be reflected back from the OMC. However, because the mirrors of the OMC do not have perfect reflectivity, the range of frequencies which can pass is broadened, allowing most of the sideband TEM₀₀ modes to pass. This can be seen by looking at the transmitted light of the OMC, considering the reflection and transmission at each mirror. Assuming no losses in the cavity and no transmission on the end mirror, the complex amplitude of the transmitted light will be

$$\Psi_{transmitted} = \frac{t_1 t_2 \Psi_{incident}}{1 - r_1 r_2} e^{i\phi} \quad (3.5)$$

where

$$\phi = 2kL - 2(m + n + 1) \cos^{-1} \left(\sqrt{1 - \frac{L}{R_o}} \right) - \pi \text{Mod}_2(m), \quad (3.6)$$

$t_1 = t_2$ represent the transmission of the mirrors, $r_1 = r_2$ represent the reflectivity of the mirrors, and assuming $t_i^2 + r_i^2 = 1$, for $i = 1, 2$. The k represents the wavenumber of the light and varies due to the 24.5 MHz frequency difference between the carrier and sidebands, resulting in a different transmission ratio for the carrier and sidebands of the same spatial mode.

Decomposing into the OMC basis

In order to understand how the astigmatic beam of Eq. (3.2) passes through the OMC, we need to decompose it into the Hermite Gaussian basis of the OMC, Eq. (3.1). This was done in two ways, both of which resulted in the same decomposition. The first decomposition was done by taking the inner product between the astigmatic beam and a particular OMC mode, essentially the inner production between Eqs. (3.1) and (3.2). This inner product,

$$\int_{-\infty}^{\infty} \Psi_{mn} \Psi_{astigmatic} = a_{mn} e^{i\theta} \quad (3.7)$$

yields the amplitude and phase of the TEM_{mn} mode present in the beam.

The other method of decomposition started with creating a set of theoretical data with Eq. (3.2), with the amplitude of the astigmatic beam stored at a series of x and y coordinates. The simple model of

$$a_{00}e^{i\phi_{00}}\Psi_{00} + a_{01}e^{i\phi_{01}}\Psi_{01} + a_{10}e^{i\phi_{10}}\Psi_{10} + a_{02}e^{i\phi_{02}}\Psi_{02} + a_{11}e^{i\phi_{11}}\Psi_{11} + a_{20}e^{i\phi_{20}}\Psi_{20} , \quad (3.8)$$

where Ψ_{mn} is the mn^{th} OMC mode, was then fit to this artificial data set. The amplitude and phase of each of the modes (defined by Eq. (3.1)) is then left as a free parameter of the model. After fitting, we normalize the model by scaling so that the TEM₀₀ mode has an amplitude of one and a phase of zero. This second method is very similar to the one used to analyze the phase camera data. The results of the two decompositions are presented in the first two columns of Table 3.2.2. The matches between modes from the two methods are good, within ± 0.02 in amplitude and $\pm 11^\circ$ in phase.

There are two additional decompositions which can provide insight, that of an astigmatic beam without the 25° of rotation, where the two waists axes are aligned with the OMC axes, and that of an astigmatic beam misaligned in vertical displacement coming into the OMC. In the first case, the TEM₁₁ mode completely disappears, leaving only the expected TEM₀₂ and TEM₂₀ modes due to the mode mismatch. In the second case a vertical displacement of one third the waist size of the OMC was somewhat arbitrarily chosen because of measurements indicating alignment drift of this size on short time scales. This decomposition gives a feel for what deviations to expect between the OMC model and the phase camera data. The differences are mostly on the order of 15% for modes that are present in the perfect alignment and the inclusion of several misalignment modes, such as the TEM₀₁. Table 3.2.2 summarizes all of these decompositions.

Lastly, to calculate the effect of the astigmatic beam on transmission of the OMC, one simply needs to apply Eq. (3.5) to each mode for both the carrier and sidebands, then add all the modes back together for the carrier or sideband.

Decomposition versus Michelson bright fringe OMC scans

One straightforward experimental test of the astigmatic decomposition is to tune the OMC's length such that modes other than the TEM₀₀ become resonant and note the power of the transmitted light. In practice this is done by applying a slowly increasing voltage to the OMC's PZT to change the length of the OMC while continuously recording the power of

the light on transmission. By calculating the transverse mode spacing (TMS) from

$$\text{TMS} = \frac{1}{\pi} \arccos \left(\sqrt{1 - \frac{L}{R_0}} \right) \quad (3.9)$$

we can easily match peaks with sets of degenerate modes, since a TEM_{mn} mode will be shifted in frequency from the TEM_{00} mode by $(m+n) \times \text{TMS}$ in units of free spectral range. For the OMC, the TMS is 0.308 based off of the parameters from Table 3.1. To compare the astigmatic decomposition to the data taken in this way during a bright Michelson lock, we need to apply the OMC response function, Eq. (3.5), to the mode amplitude from Table 3.2.2 while letting L vary, which lets us go through a free spectral range of resonant frequencies. Then we take the power in each mode, sum them all, and plot as a function of free spectral range. The final formula used is thus

$$A \left(a_{00}^2 |T_{00}^{\text{carrier}}(dL)|^2 + a_{01}^2 |T_{01}^{\text{carrier}}(dL)|^2 + \dots \right) \quad (3.10)$$

where a_{mn} and θ_{mn} are amplitude and phase of mode mn from Table 3.2.2, $T_{mn}(dL)$ is the response of the OMC for mode mn for a given length change dL , and A is an overall scale parameter. Note that we ignore the effects of the sidebands in this calculation since for the bright Michelson lock the DC power is dominated by the carrier. Figure 3-3 shows a plot of the scan data taken during a bright Michelson lock that is overlaid with Eq. (3.10) as a function of dL . To produce amplitudes of the higher order modes that we can compare our models to, we fit the scan data to

$$\text{Power}(dL) = \sum_{m,n}^{m+n=4} A_{mn} |T_{mn}^{\text{carrier}}(dL)|^2 \quad (3.11)$$

where the amplitudes A_{mn} are the free parameters of the fit. A_{mn} are just the square of a_{mn} . Due to degeneracy between modes, we cannot completely distinguish A_{mn} terms, but we can compare the sums of the degenerate modes. Table 3.2.2 provides a comparison of these sums from the model and OMC scan data, with all the values normalized to a TEM_{00} transmission of 1. Both Figure 3-3 and Table 3.2.2 show a very strong correspondence between our model's prediction and the actual experimental result, especially when we consider that the model assumes perfect alignment and thus predicts no misalignment modes such as TEM_{10} and TEM_{01} .




Figure 3-2: The triangular OMC cavity is shown in relation to the full interferometer. It is placed at the AS port to clean up the light reaching the AS photodiode. The red light represents the carrier, while the blue light represents the sidebands. The phase camera was placed at pickoffs before and after the OMC to capture the state of the light. The phase camera images are from when the interferometer was held on a Michelson bright fringe and are expanded in Figures 3-7 and 3-5. During this state, the RM and ETMs were misaligned.

Mode	Integration Method		Fit Method		Integration Method Misaligned		Integration Method Non-rotated	
	Amp.	Phase	Amp.	Phase	Amp.	Phase	Amp.	Phase
00	1	0	1	0	1	0	1	0
01	-	-	-	-	0.28	-155	-	-
10	-	-	-	-	0.13	25	-	-
02	0.23	-67	0.24	-56	0.22	-43	0.28	-150
11	0.26	80	0.25	77	0.23	87	-	-
20	0.21	31	0.22	26	0.22	32	0.47	39
03	-	-	-	-	0.11	153	-	-
12	-	-	-	-	0.12	-63	-	-
21	-	-	-	-	0.05	-172	-	-
30	-	-	-	-	0.05	57	-	-
04	0.07	-113	0.07	-113	0.06	-86	0.19	-136
13	0.11	23	0.11	21	0.09	43	-	-
22	0.02	176	0.01	177	0.003	-161	0.02	38
31	0.10	111	0.10	103	0.09	117	-	-
40	0.06	62	0.06	51	0.061	64	0.15	45
05	-	-	-	-	0.041	101	-	-
14	-	-	-	-	0.06	-120	-	-
23	-	-	-	-	0.02	46	-	-
32	-	-	-	-	0.04	-46	-	-
41	-	-	-	-	0.02	-181	-	-
50	-	-	-	-	0.02	88	-	-

Table 3.2: The mode column lists the mode mn number. All the amplitude columns give the amplitude of the corresponding mode relative to the TEM_{00} which was fixed at 1. All the phase columns give the phase difference from the TEM_{00} mode, which was fixed at phase of 0° . The integration method decomposition was done by taking the inner product between the astigmatic beam and a particular OMC mode, as described by Eq. (3.7). The fit method amplitude and phase columns were done by creating a set of data then fit to a simple model described by Eq. (3.8). The misaligned integration method gives the mode amplitudes and phases entering the OMC misaligned in the vertical direction by one third of a waist size. The non-rotated integration method gives the results for an astigmatic beam without the 25° rotation, so that the waist axes are aligned with OMC axes.

Degenerate Mode	Modes	Data Fit (Power in modes)	Model (Power in modes)	Misaligned Model (Power in modes)
1 even	01	0.1396	-	0.0784
1 odd	10	-	-	0.0169
2 even	20,02	0.1040	0.1004	0.0968
2 odd	11	0.0633	0.0686	0.0529
3 even	03,21	0.0240	-	0.0146
3 odd	12,30	0.0232	-	0.0169
4 even	04,22,40	0.0066	0.0080	0.0073
4 odd	13,31	0.0214	0.0206	0.0162
5 even	05,23,41	0.0128	-	0.0025
5 odd	50,32,14	-	-	0.0056

Table 3.3: The degenerate mode column labels the resonances based on the order of the modes which pass through. The break in the degeneracy between modes of the same $m + n$ value is due to the triangular nature of the cavity. For odd mode numbers in the direction parallel to the mode cleaner plane there is an additional 0.5 FSR shift in resonance frequency. The modes columns lists all the modes which resonated together in the OMC. The power in modes is the sum of the power of all the resonant modes relative to the TEM_{00} , which is scaled to 1. The data fit column comes from fitting the scan data to Eq. (3.10). The model and misaligned models come from calculating the power from the amplitudes in Table 3.2.2 and then summing the appropriate modes together.




Figure 3-3: Plot of the OMC transmitted power versus detuning in units of FSR. The black line is the actual data, the blue line is the model prediction and the red line is a fit to the data. The fit was produced by fitting the data to equation Eq. (3.10).

What the phase camera should see

The phase camera sees the strength of the beat between the carrier and the sidebands, by demodulating at 24.5 MHz. The sidebands are generated from the carrier via phase modulation of the form

$$E_0 e^{i(\omega t + \Gamma \sin(\Omega_m t))} \cong E_0 \left[J_0(\Gamma) e^{i\omega t} + J_1(\Gamma) e^{i(\omega + \Omega_m)t} - J_1(\Gamma) e^{i(\omega - \Omega_m)t} \right] \quad (3.12)$$

where E_0 is the original carrier amplitude, ω is the angular frequency of the carrier, Ω_m is the frequency of the modulation, Γ is the modulation depth, and t is time. The approximation on the right hand side of Eq. (3.12) is simply an expansion, followed by collection of terms into Bessel functions of 0th and 1st order. The first, second and third terms on the right hand side are the carrier amplitude, the upper sideband amplitude, and the lower sideband amplitude, respectively. Since the carrier and sideband see the same surfaces for a bright Michelson lock or a direct reflection off an ITM, and the sideband generation should only produce a frequency and amplitude difference between the sideband and carrier, it is reasonable to assume that the modal content will be the same for both. In that case we can write the individual field amplitude of the carrier and sidebands as

$$A_{carrier} = E_0 J_0(\Gamma) \left(a_{00} e^{i\theta_{00}} T_{00}^{carrier} \Psi_{00} + a_{01} e^{i\theta_{01}} T_{01}^{carrier} \Psi_{01} + \dots \right)$$

$$A_{sb+} = E_0 J_1(\Gamma) \left(a_{00} e^{i\theta_{00}} T_{00}^{sb+} \Psi_{00} + a_{01} e^{i\theta_{01}} T_{01}^{sb+} \Psi_{01} + \dots \right) \quad (3.13)$$

$$A_{sb-} = -E_0 J_1(\Gamma) \left(a_{00} e^{i\theta_{00}} T_{00}^{sb-} \Psi_{00} + a_{01} e^{i\theta_{01}} T_{01}^{sb-} \Psi_{01} + \dots \right) \quad (3.14)$$

where $T_{00}^{carrier}$ is the OMC response, obtained from $\Psi_{transmitted}/\Psi_{incident}$ from Eq. (3.5) for the appropriate mode, and the Ψ_{mn} are from Eq. (3.1) with the Gouy phases fixed to match that of the experimental setup. The total field of the light can then be written as

$$A_{carrier} e^{i\omega t} + A_{sb+} e^{i(\omega + \Omega_m)t} + A_{sb-} e^{i(\omega - \Omega_m)t}. \quad (3.15)$$

Demodulating the power in phase at an angular frequency Ω_m yields the in-phase (Iphase), and 90° out of phase yields the quadrature-phase (Qphase) signals. These can be

explicitly written as

$$2\Re(A_{carrier}A_{sb+}^* + A_{carrier}^*A_{sb-}) = Q_{phase} \quad (3.16)$$

and

$$2\Im(A_{carrier}A_{sb+}^* + A_{carrier}^*A_{sb-}) = I_{phase} \quad (3.17)$$

By replacing the amplitudes a_{mn} and phases θ_{mn} with the appropriate model values we can create an I or Qphase map which is directly comparable to the measured phase camera phase maps. The phase map in Figure 3-4 is created using the parameters from the integration method of astigmatic beam decomposition listed in Table 3.2.2, with an overall demodulation phase rotation such that all the higher order mode “junk” is in one phase but not the other.




Figure 3-4: Phase map generated from the OMC model, with the integration method parameters from Table 3.2.2. This shows spatial amplitude of the Qphase demodulation. The demodulation phase rotation has been chosen such that all the higher order “junk” is in the Iphase, leaving the TEM_{11} mode as the largest contributor to the Qphase signal.

3.2.3 Phase camera measurements and comparisons

Phase camera data

The following images were taken by the phase camera in two different states. The first was when the OMC was locked with maximum light transmission while the interferometer was held on a bright Michelson fringe. The second was when all the mirrors of the interferometer were misaligned, except for ITMY and the BS, resulting in a direct reflection of light into the AS port. The images are reproducible for a given state of the interferometer.

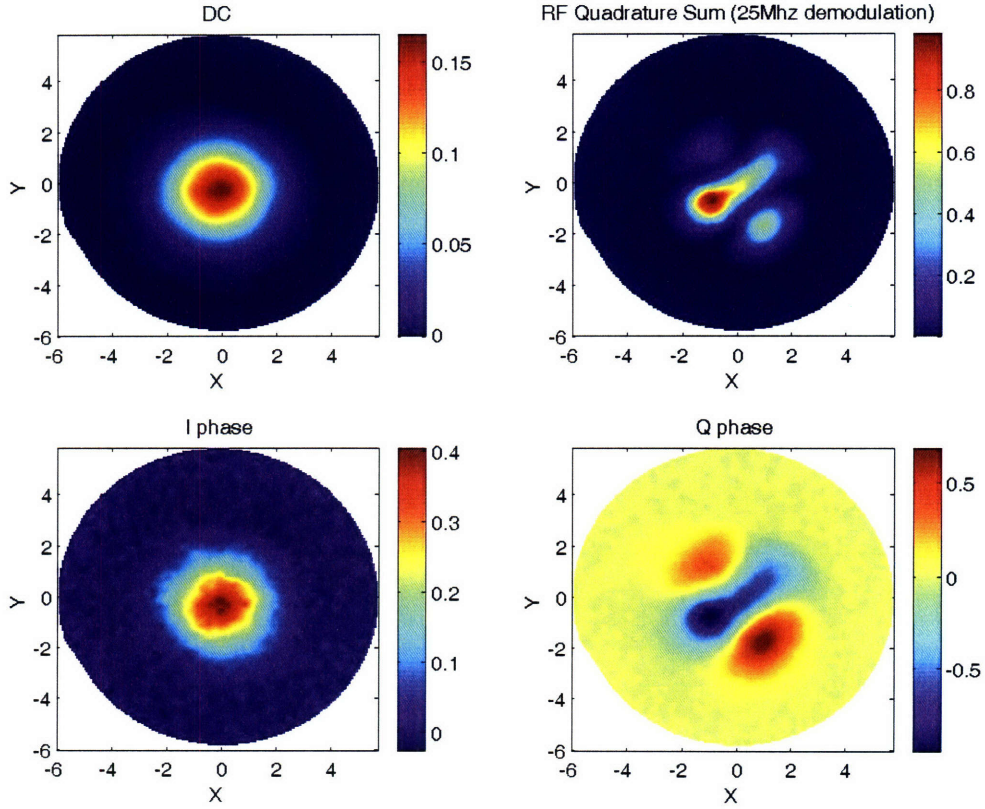


Figure 3-5: The above images from the phase camera were taken on transmission through the OMC when the interferometer was held on a Michelson bright fringe. The upper left image shows the spatial amplitude of the DC signal, which should be made up mostly of the clean TEM_{00} carrier signal. The lower left image is the spatial amplitude of the Iphase of the demodulated RF signal, demodulated at Ω_m , 24.5 MHz. The fact that we see a TEM_{00} mode in the Iphase implies there is some amplitude modulation or mismatch between the sidebands, since otherwise the beat the sidebands produce with the TEM_{00} carrier should cancel out. The lower right is the Qphase of the same signal. Our prediction for the Qphase image from modeling the astigmatic beam passing through the OMC is Figure 3-4, mostly the TEM_{00} mode which we see. The upper right is the overall amplitude of the signal calculated from $\sqrt{(Iphase)^2 + (Qphase)^2}$.

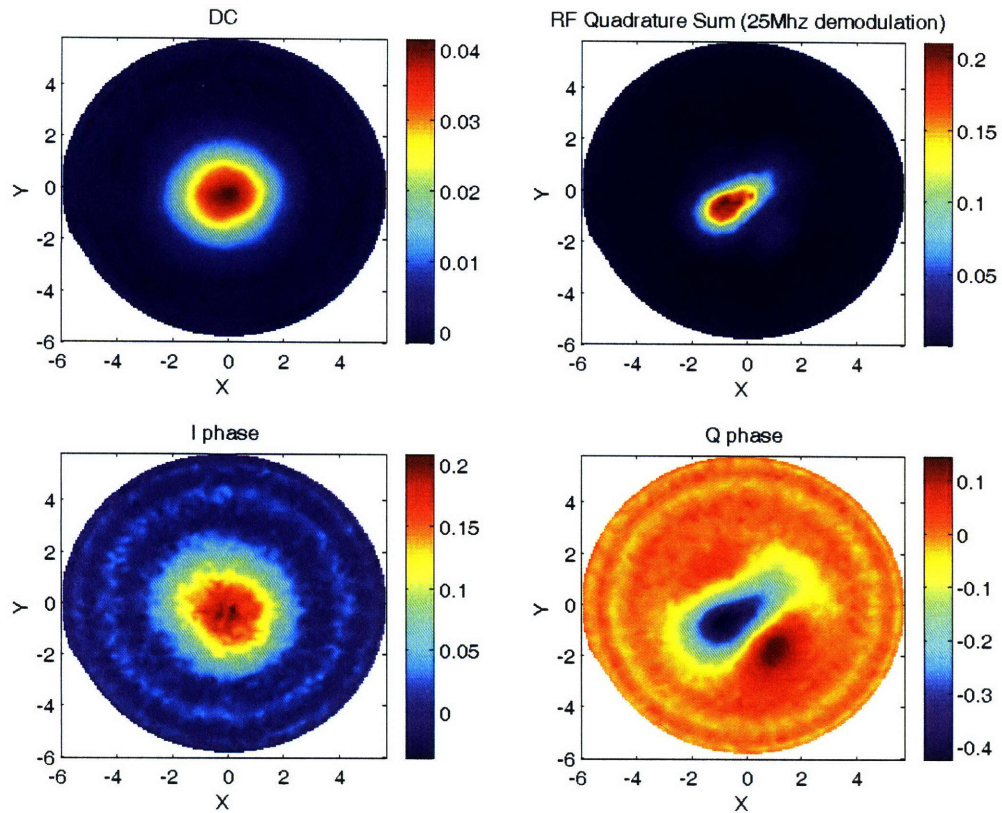


Figure 3-6: The above images from the phase camera were taken when only ITMY and the BS were aligned, producing a direction reflection of light off ITMY, through the BS, and to the OMC. The signals from which the images are generated are described in Figure 3-5. The factor of 4 reduction in overall power is due to the loss of 50% of the light at the beam splitter because of the ITMX misalignment. This also means the electronics background noise stands out more in these images when compared to Figure 3-5.

The DC power images show fairly clean TEM_{00} modes, indicating the OMC is in fact transmitting mostly the mode we want. The factor of 4 drop in DC power intensity from the Michelson bright fringe case to the direct reflection case is due to the loss of 50% of the light at the beam splitter because of the ITMX misalignment. This also means the direct reflection case is closer to the noise floor of the phase camera. The I and Qphases have been rotated to match the same overall phase rotation chosen for the model, so that one phase is expected to be flat and the other will contain the higher order modes transmitted. The fact that we see a clear TEM_{00} mode in one phase instead of a flat plane implies some amplitude modulation of the carrier and sidebands is being produced by or passing through the OMC. The TEM_{11} mode image in the Qphase is reassuring, as it looks similar to the model prediction from Figure 3-4.

To determine if the TEM_{00} mode present in the Iphase is due to the OMC or present on the light before reaching the OMC, we look at phase camera images when the OMC was removed. Figures 3-7 and 3-8 show the bright Michelson lock and an ITM direct reflection respectively, without the OMC in the path. In the bright Michelson case the amplitude modulation can be fully explained as being on the incoming light, since the ratio of the DC power to the Iphase power is roughly the same or less for the OMC present and not present cases as shown in Table 3.4. However, for the case of the direction reflection off an ITM, the amplitude modulation seems to be greater after the OMC. However, in this case the reflection was off ITMX rather than ITMY. The ratio of the Iphase signal is down by much more than the expected factor of 2 in amplitude when comparing between the Michelson bright fringe and direct reflection cases, while the DC amplitude is down by the expected factor of 4. This indicates the sidebands in the ITMX direct reflection case were affected by the reflection much differently from the ITMY case, and in such a way as to reduce the overall amplitude of the beat between the carrier and sidebands.

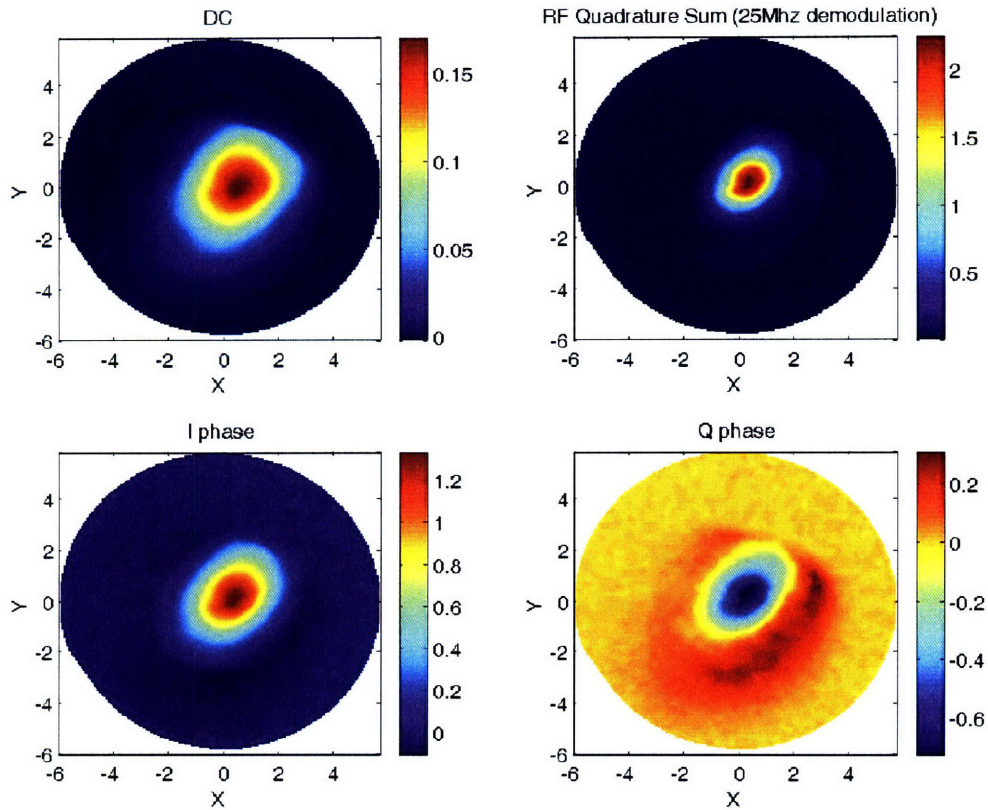


Figure 3-7: The above images were taken when the interferometer was held on a Michelson bright fringe and without the OMC in the light path. The signals from which the images are generated are described in Figure 3-5. Without an OMC to clean the signal, we see some combination of astigmatic modes. When compared to Figure 3-7, the overall power in the DC image is roughly the same, but the beat signal increases, which shows that the sidebands are more affected by the OMC than the carrier. This is expected since the OMC was designed to pass the carrier, and the sidebands are slightly offset in frequency, reducing their transmission.

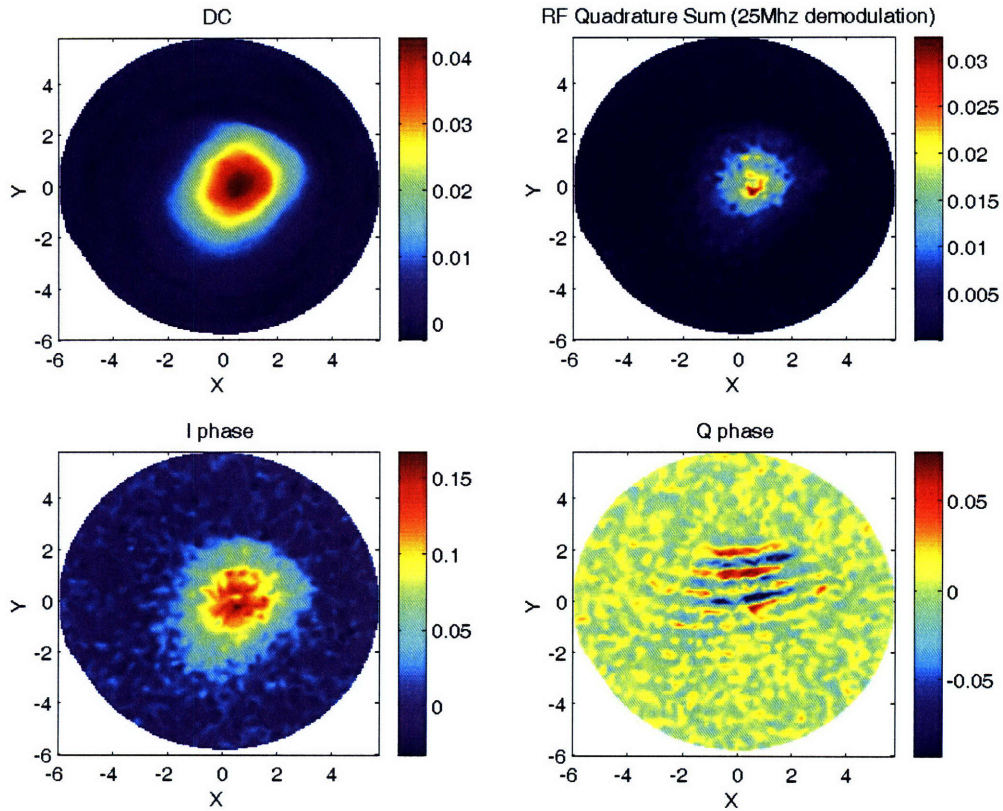


Figure 3-8: The above images from the phase camera were taken when only ITMX and the BS were aligned, producing a direction reflection of light off ITMX which was then directed to the phase camera. The signals from which the images are generated are described in Figure 3-5. In comparison to Figures 3-7, 3-6, and 3-5, we see that the carrier on sideband beat signal is lower than expected, while the DC light (representing mostly carrier) is relatively unchanged, indicating a difference in the sidebands between the ITMY reflection and the ITMX reflection.

Table 3.4: Amplitude modulation comparisons

	Data Set 1 Iphase/DC	Data Set 2 Iphase/DC	Data Set 3 Iphase/DC
Michelson Bright Fringe With OMC	2.28	2.48	2.45
Michelson Bright Fringe Without OMC	3.7	3.68	3.75
Direct Reflection With OMC	7.03	3.9	4.97
Direct Reflection Without OMC	3.12	3.18	3.2

Phase camera data and model comparison

To directly compare the models, we first need to set certain scale parameters. The easiest method of getting these scale parameters is to fit the DC power data to $A\Psi_{00}\Psi_{00}^*(x + x_{\text{offset}}, y + y_{\text{offset}}, w)$, letting the overall amplitude A , the center position offsets x_{offset} and y_{offset} , and w the waist size vary as free parameters. Fixing the offsets and waist size in our model first makes the final stage of fitting to the model much easier. We then make a nonlinear fit of Eqs. (3.16) and (3.17) to the data, treating $E_0^2 J_0(\Gamma)J_1(\Gamma)$ as a single amplitude free parameter A in addition to the a_{00} and θ_{00} free parameters. We use the values from Table 3.2.2 as our initial guesses, and restrict the parameter space to about $\pm 20\%$ of those values. To determine the quality of the fit, we first determine the noise in the data. By subtracting phase camera images from one another, and taking the standard deviation of the resulting difference values gives us our σ . We then use this to calculate our χ^2 (normalized) merit functions with the values between 1 and 4, depending on how noisy a particular image is, and whether it was the bright Michelson or direct reflection case. However, the residuals show some distinct structure excess, perhaps some TEM_{00} mixed with TEM_{20} in the case of the direct reflection fits. The two fits of the Qphase, one for the bright Michelson and one for the direct reflection in Figure 3.2.3 are typical of the sets.

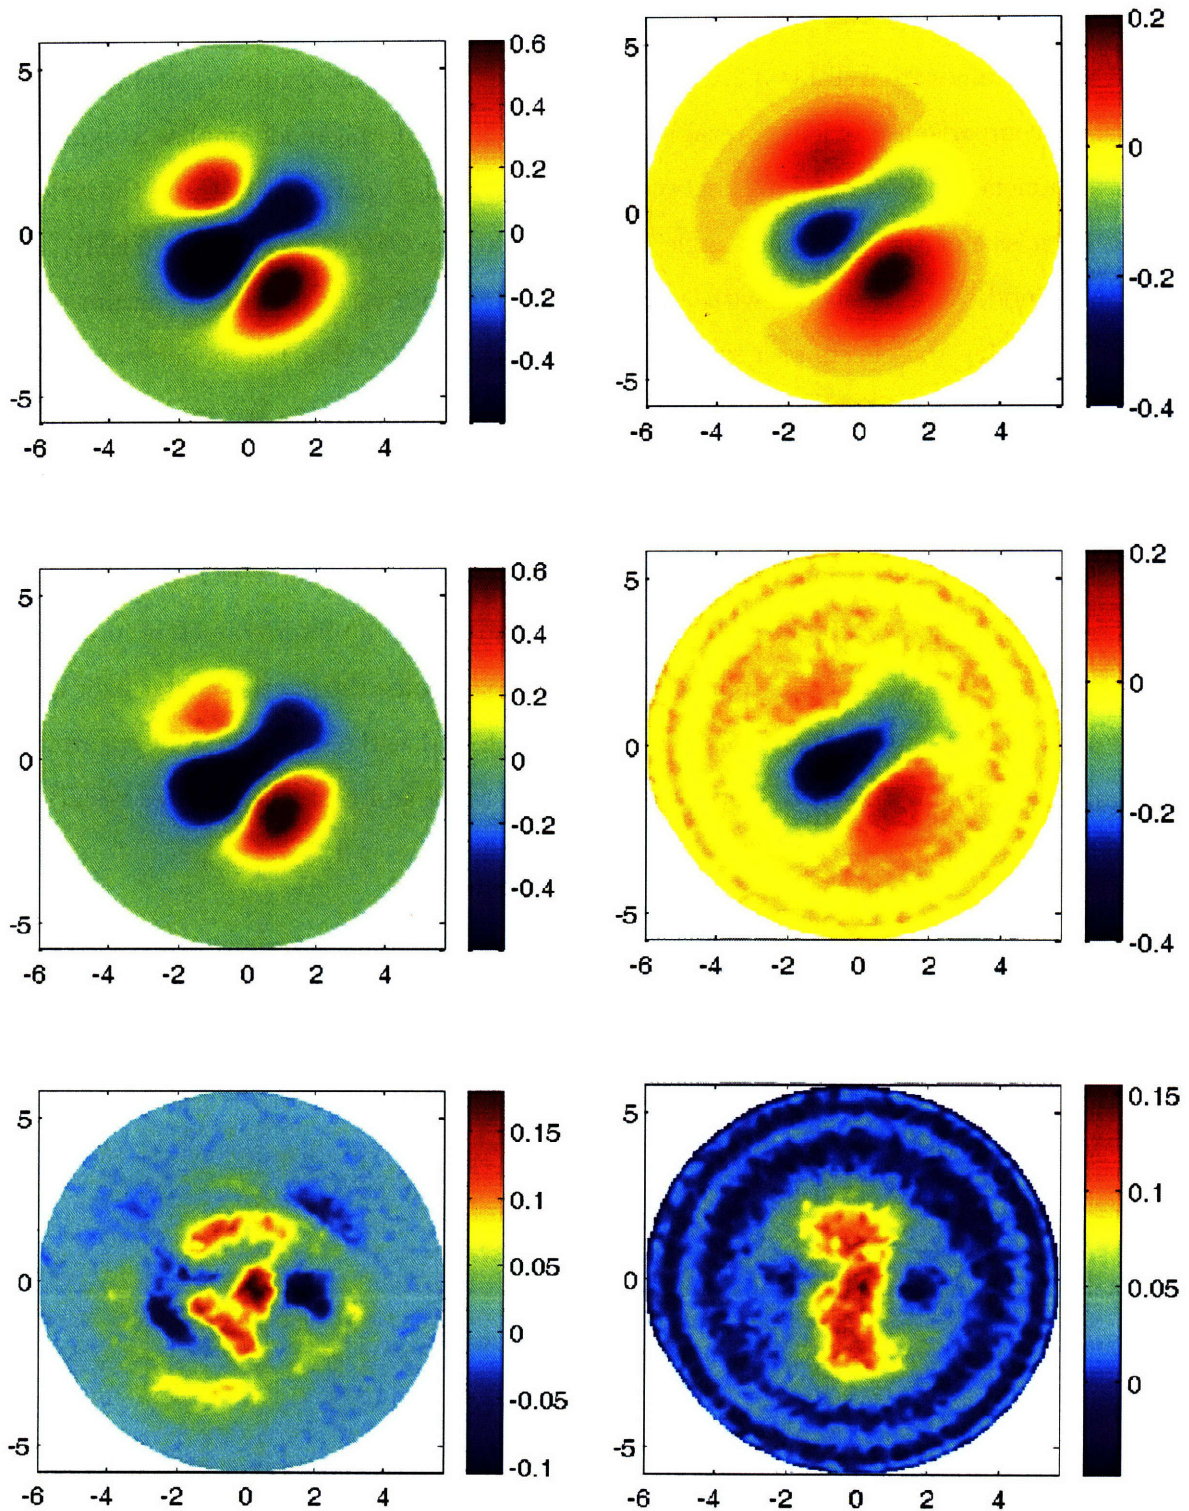


Figure 3-9: Left: Qphase fit of a Michelson bright fringe data, as seen by the phase camera on transmission of the OMC. Right: Qphase fit of a direct reflection off ITMY data, as seen by the phase camera on transmission through the OMC. The respective top image is a plot of the fit including modes of up to order 4. The respective middle image is of the phase camera data for the Qphase. The background electronics noise stands out more in the direct reflection images because of the lower overall power reaching the phasecamera.

Table 3.2.3 gives the parameter values for the two sets of measurements, along with the model values from Table 3.2.2, normalized to a TEM_{00} amplitude of 1 and a phase of 0. This comparison provides several insights. The large TEM_{01} mode implies a significant misalignment of the vertical axis between the incoming light and the OMC, on the order of the waist size or the divergence angle of the light at the OMC. The TEM_{02} , TEM_{11} , and TEM_{20} modes are consistent with being produced solely by the astigmatic nature of the light. Ignoring the effects of a simple misalignment, the fact that the model matches the data well implies the OMC is doing what it is expected to be doing. We get large amounts of higher order modes on transmission of the OMC mostly due to the known astigmatism, and that other additional problems or effects are not needed to explain the OMC output.

After testing finished on the OMC it was removed and returned to GEO. However, there are plans for a new OMC to eventually be added to the LIGO interferometers in a future upgrade. The phase camera system was next used to examine the interferometer light inside the recycling cavity in different heating states with and without the TCS in an attempt to understand how the ASI signal was being generated and corrected. Those results are presented in Chapter 3.

Mode	Model		Misaligned Model		Mich Bright		Direct Reflection	
	Amp.	Phase	Amp.	Phase	Amp.	Phase	Amp.	Phase
χ^2	-	-	-	-	0.802	-	3.889	-
A	-	-	-	-	258	-	99	-
00	1	0	1	0	1	0	1	0
01	-	-	0.28	-155	0.43	71	0.50	62
10	-	-	0.13	25	0.044	-15	0.066	2
02	0.23	-57	0.22	-43	0.30	-44	0.30	-72
11	0.26	80	0.23	87	0.256	100	0.20	100
20	0.21	31	0.22	32	0.27	11	0.27	11
03	-	-	0.11	153	0.023	12	0.03	12
12	-	-	0.12	-63	0.09	-70	0.04	62
21	-	-	0.05	-172	0.06	-17	0.06	-38
30	-	-	0.05	57	0.08	59	0.12	34
04	0.07	-113	0.06	-86	0.02	-90	0.02	-90
13	0.11	23	0.09	43	0.127	129	0.11	175
22	0.02	176	0.003	-161	0.03	191	0.03	191
31	0.10	111	0.09	117	0.10	131	0.08	91
40	0.06	62	0.061	64	0.04	82	0.04	82

Table 3.5: The χ^2 row indicates the χ^2 from the fit. The A factor indicates the factor needed to rescale the TEM_{00} amplitude to 1. The amplitude columns indicate the relative amplitude of the mode to the TEM_{00} and the phase columns indicate the difference in phase from the TEM_{00} , which was set to 0° .

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Improving and applying the phase camera

This chapter begins by describing the addition of an independent reference beam to the phase camera which is necessary for looking at the carrier and sideband components of the light individually. The chapter then focuses on how this improved phase camera system was used to investigate the heating problems of the LIGO 4k interferometer. The system was able to show a consistent progression through the different heating states of the interferometer and what the TCS was correcting that reduced the excess ASI noise and allowed the LIGO interferometers to operate at their designed operating power and sensitivity.

4.1 Phase camera with reference beam

4.1.1 Experimental setup

Interferometer light

Since the carrier resonates both in the arms and the power recycling cavity while the RF sidebands only resonate in the power recycling cavity, they effectively “see” different parts of the interferometer. The Fabry-Perot cavities can be described by their finesse, which is the FSR divided by the bandwidth of the resonances and is inversely proportional to the losses in the cavity, and their g factor, which is $(1 - L/R_{\text{ITM}})(1 - L/R_{\text{ETM}})$ where L is the length of the cavity, R_{ITM} is the radius of curvature of the ITM, and R_{ETM} is the radius of curvature of the ETM. In the same way that the OMC prevented higher order modes to

propagate because of the Gouy phase shift (see Section 3.2.2), the Fabry-Perot arm cavities suppress higher order modes. The cavities have high enough finesse, $F = 219$, combined with their $g = 0.33$, that the higher order modes are suppressed by a factor of up to 60 compared to a resonant TEM_{00} mode.

The almost flat input mirrors of the power recycling cavity support many modes, in which case the sidebands will tend to include higher order TEM modes. To understand better what is occurring inside the interferometer, it is helpful to be able to separate the carrier light from the sideband light. One way to do this is to introduce a reference beam and measure the its beat note with the carrier and each RF sideband independantly.

Reference beam

A reference beam shifted off by 75 MHz from the main carrier frequency and then mixed with the full interferometer light will produce a beat with the carrier at 75 MHz, a beat with the upper sideband at 99.5 MHz and a beat with the lower sideband at 50.5 MHz, as shown in Figure 4-1. For comparison, the carrier and sidebands beat at 24.5 MHz, and the sidebands beat with each other at 49 MHz.

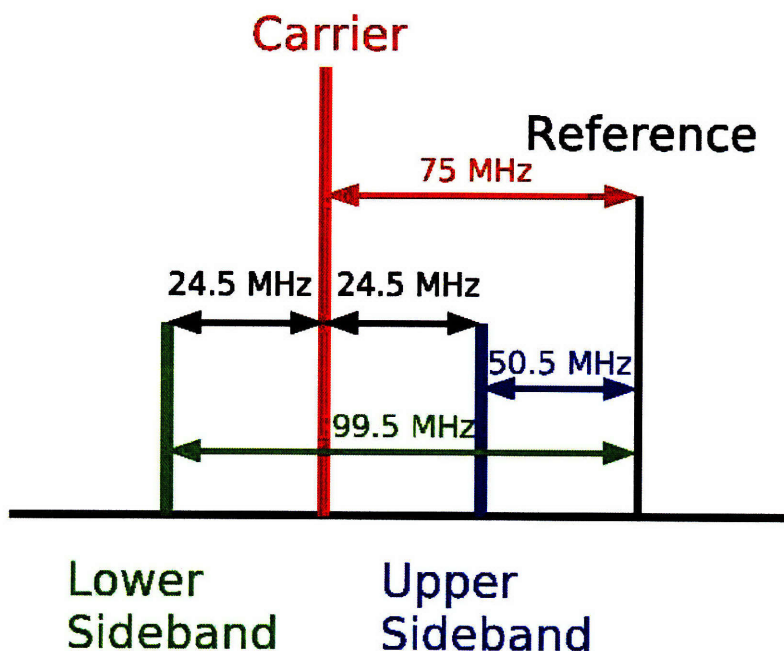


Figure 4-1: This diagram shows the relative positions of the carrier, sidebands, and shifted carrier reference beam in frequency. The carrier is 24.5 MHz from the sidebands and 75 MHz from the reference beam.

This reference light was generated by taking light directly from the output of the Light-wave laser and passing it through an acoustic optical modulator (AOM) driven at 75 MHz, as shown in Figure 4-2. This light was then coupled into a fiber optic cable which was run from the laser table to the phase camera setup. A fiber to free space coupler was attached to the fiber, and the resulting light was passed through a half-wave plate and a polarizing beam cube so as to match the polarization of the incoming interferometer light. A beam splitter was then used to combine the incoming interferometer light with the reference beam before going on to the rest of the phase camera system.

4.1.2 Demodulation board redesign

It was quickly discovered that the overall phase of the reference beam varied on the order of minutes, most likely due to thermal changes along the length of the fiber. In the time it took to switch the demodulation signal to the demodulation board the signal phase could vary, making it much harder to piece together the individual structure. The solution was to take all the data simultaneously, essentially recording the DC signal, the carrier and reference beat, the upper sideband and reference beat, and the lower sideband and reference beat at once. The images taken are simultaneously directly comparable, since they all have identical reference beam effects.

A redesign of the control and demodulation board was required, as the initial design only could demodulate and record one signal at time. The new board design can be found in Appendix B. Essentially, the demodulation channels and computer recording channels were increased by a factor of four, and the computer software was updated to handle the higher data rate.

The boards were assembled and tested at Hanford and performed satisfactorily before being installed in the phase camera system.

4.1.3 Fitting and understanding the data

The use of a reference beam creates several challenges in understanding the output of the data. The demodulated output of each channel is a combination of the reference beam and the signal of interest (either carrier or individual sideband). The reference beam also adds some noise. However, because of the new demodulation board where all the data is taken simultaneously, at each point recorded in the scan the reference beam must affect the

carrier and sidebands in exactly the same way. This provides us with enough information to separate the carrier and sidebands from the reference beam in software.

Since we record DC power of the light we know the total sum of the power of the four signals. This suggests the following method for calculating the amplitude effects of the reference beam. We know that

$$P_{DC} = A_{cr}^2 + A_{sb-}^2 + A_{sb+}^2 + A_{ref}^2 \quad (4.1)$$

where P_{DC} is the recorded DC power level. By taking the sum of the squares of the Iphase and Qphase for each demodulated signal we can calculate the total power in the IFO light times the reference beam power, writing

$$S_{total} \equiv A_{cr}^2 A_{ref}^2 + A_{sb-}^2 A_{ref}^2 + A_{sb+}^2 A_{ref}^2 = A_{ref}^2 (A_{cr}^2 + A_{sb-}^2 + A_{sb+}^2) \quad (4.2)$$

where the S_{total} is just the sum of the signals. Multiplying Eq. (4.1) by A_{ref}^2 and substituting in our known sum of signals, we get the following equation

$$A_{ref}^4 - A_{ref}^2 P_{DC} + S_{total} = 0. \quad (4.3)$$

This can be easily solved for A_{ref}^2 which can then be in turn be used to find just A_{cr} , A_{sb-} , and A_{sb+} . Figure 4-3 shows a horizontal and vertical cut of the power data before and after this calculation was performed, effectively $|A_{cr}|^2$ and $|A_{cr} A_{ref}|$, to make it easier to see the difference in the spatial power of the signal. The data was taken from a state with 6 Watts of power input to the interferometer, and when the TCS was running. It also shows the calculated reference beam power, $|A_{ref}|^2$, which was then scaled up by a factor of 10 so as to fit in the plot. One problem with this method is when the electronics noise becomes an appreciable part of the total signal coming from the phase camera, the above equations break down (and sometimes predict imaginary powers). In those regions of low signal to noise (roughly one sigma in the electronics noise) the data was ignored, resulting in blank sections in those plots that use Eq. (4.3) to calculate the effects of the reference beam.

One test of the effectiveness of this method is to take the individual sideband and carrier signals, and mix them in software in the same way they mix on the photodiode to produce an image which can be directly compared to the recorded one omega signal. For example,

the lower plot in Figure 4-4 shows a one omega signal in a thermally tuned state, such that the ASI noise was minimized with the use of the TCS. To begin, we define a set of Iphase and Qphase sideband signals as function of an overall rotation fitting parameter. For the lower side band we can write

$$\begin{aligned}
I_{sb-,rot}(\theta_{sb-}) + iQ_{sb-,rot}(\theta_{sb-}) = \\
(I_{sb-} \cos(\theta_{sb-}) + Q_{sb-} \sin(\theta_{sb-})Q_{sb-}) \\
+ i(Q_{sb-} \cos(\theta_{sb-}) - I_{sb-} \sin(\theta_{sb-}))
\end{aligned} \tag{4.4}$$

where the rotation angle θ_{sb-} rotates the overall phase of the entire lower sideband data set. The upper sideband is the same except with θ_{sb+} as its rotation angle fitting parameter. Lastly an overall amplitude factor a were used as a third fitting parameter in the least squares fit.

The full equation fit was

$$\begin{aligned}
a[(I_{sb-,rot}(\theta_{sb-}) + iQ_{sb-,rot}(\theta_{sb-}))(I_{cr} - iQ_{cr}) \\
+ (I_{sb+,rot}(\theta_{sb+}) - iQ_{sb+,rot}(\theta_{sb+}))(I_{cr} + iQ_{cr}) + \text{complex conjugate}]
\end{aligned} \tag{4.5}$$

That least squares fit produced the top plot in Figure 4-4 which appears to be in good agreement with the measured carrier on sideband signal in the lower plot.




Figure 4-2: This diagram shows the setup of the reference beam. A small portion of the main Lightwave laser was picked off and passed through an AOM driven at 75 MHz. This was then coupled to an optical fiber and taken to the phase camera on a different table where it was passed through a half-wave plate and a polarizing beam cube before being combined with the interferometer light on a beamsplitter.

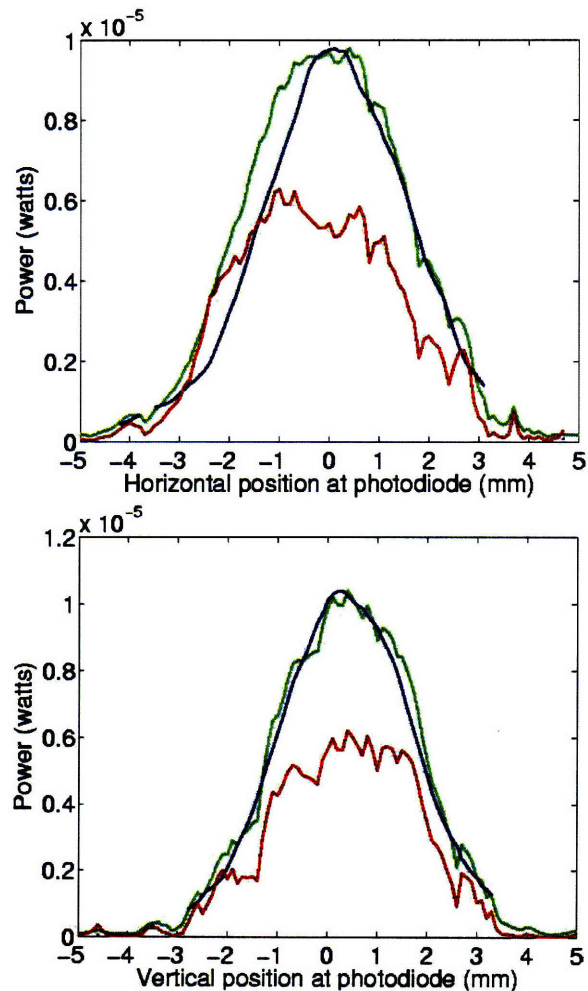


Figure 4-3: The top plot shows a horizontal profile of phasecamera data from when the interferometer had 6 Watts of power input and the TCS was on. The bottom plot shows a vertical profile for the same case. The blue line is the power in the carrier only. The red line is the calculated reference beam only, scaled up by a factor of 10. The green line is the carrier signal combined with the reference beam, effectively the measured 75 MHz signal from the phase camera, and has been scaled to match in power. The x -axis is in millimeters as measured at the position of the phase camera photodiode. By dividing out the reference beam and associated noise from it, the carrier shape gets closer to the Gaussian profile one expects of a TEM_{00} mode.

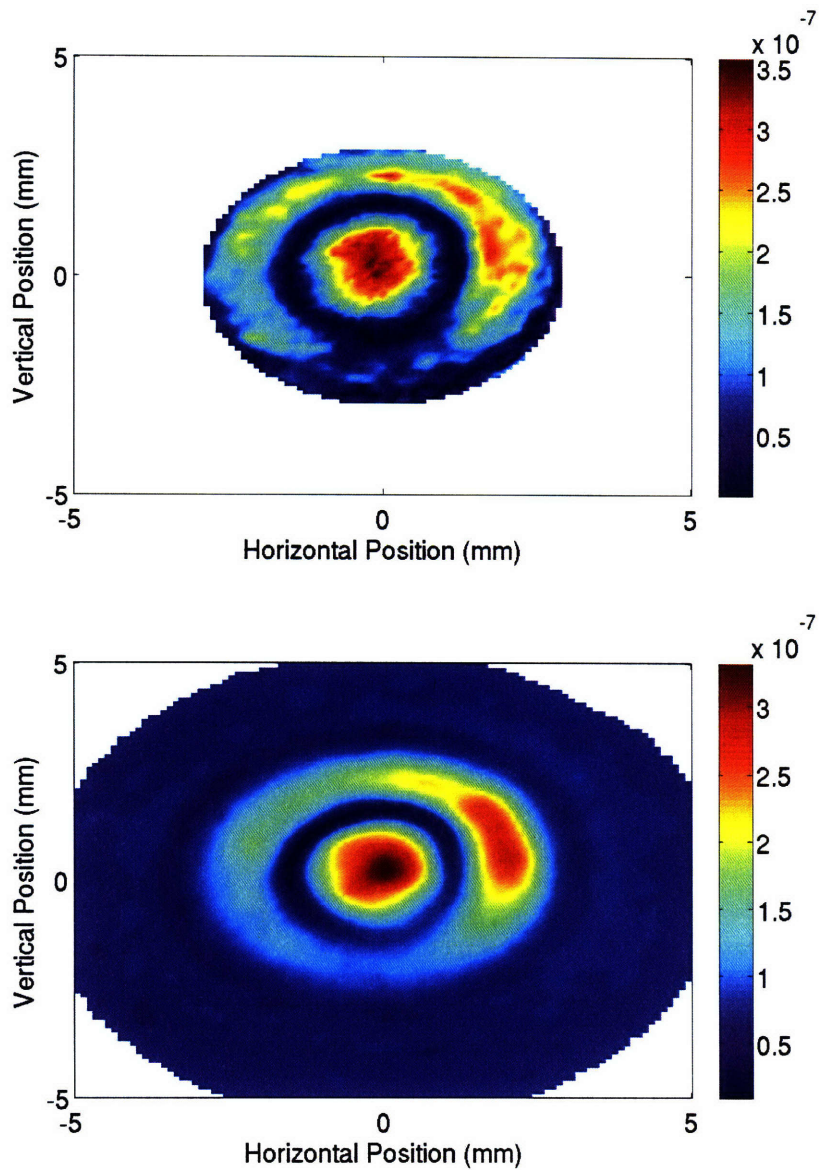


Figure 4-4: The top plot is a fit of the carrier mixed with sideband at 24.5 MHz created from the individual sideband and carrier signals measured with the reference beam. The fit was a least squares fit with three fitting parameters, specifically the phase rotation of the upper and lower sidebands, and the overall amplitude of the carrier on sideband signal. The lower plot is the directly measured 24.5 MHz carrier on sideband signal without reference beam. The color scale shows power in watts. The axes are in millimeters at the phase camera photodiode.

The data analysis on a single phase camera scan that uses the reference beam proceeds as follows. First the averaged dark noise is subtracted off from the data. The dark noise is just the signal recorded from the phase camera when no light is incident on the photodetector. Then the data is calibrated according to the photodetector response, including the effects of a 3 dB roll off at 125 MHz. The demodulation board calibration is then taken into account, which was measured before hand by using calibrated test signals in place of the photodiode. Lastly, the reference beam amplitude effects are calculated and removed as indicated previously.

4.2 H1 in different heating states

By using the phase camera to look at the inteferometer light as the optics go from receiving no heating, to optimal heating, and then to too much heating, we can understand what the effects of heating are on the interferometer light and determine how it affects the AS_I noise and the final sensitivity of the instrument.

Data was taken in four different heating states by the phase camera in the month before the S5 science run began. The data was taken from a pickoff inside the recycling cavity, giving a view of both resonant sidebands and the power recycled carrier, as opposed to the dark port carrier which should mostly cancel itself. In each case the interferometer was allowed to settle into the state for a half hour, roughly the thermal timescale of the ITMs.

The first, which we will refer to as the “No TCS” state, was with no thermal compensation at a low input power (only 1.8 watts incident on the interferometer). The second state, “Tuned TCS”, was at a low input power but with thermal compensation system set such that AS_I signal was nulled. The third, “Reversed TCS”, was at a low input power but with the thermal compensation system set to make the thermal problems worse, effectively switching the direction of the correction effects and mimicking what the interferometer does at high input power without correction. Lastly a set of data was taken while in the nominal high power configuration (6 watts incident on the interferometer) with an early iteration of the thermal compensation servo running, which we will refer to as the “Power up with TCS” state.

Table 4.2 lists some relevant recorded parameters of each state. The NSPOB channel refers to a pickoff on the beamsplitter which measures the sideband on sideband signal at

49 MHz. This signal is then normalized by the input power. The AS_DC is refers to one of the AS port's photodiode DC currents, which is a measure of the total light incident on the photodiode. AS_I_Corr is the Iphase AS port photodiode correction signal which is used to prevent saturation of the AS port photodiode.

4.2.1 Discussion of phase camera results

Table 4.2 lists ratios of the carrier to sideband power and the sideband to sideband for various states as measured by the phase camera. To make understanding of these easier, let us assume for the moment that the modes are perfectly matched, and that the only difference is the overall power of the carrier and the individual sidebands. With this assumption, predictions can be made based on the recycling gains and the relative amplitudes of the carrier and sidebands at the input to the interferometer.

In the cold state, the interferometer had known sideband recycling gain, G_{sb} , of roughly 13 in power. The highest G_{sb} achieved in a thermally tuned state was 26 [5]. With good spatial overlap of the upper and lower sidebands, the NSPOB signal should be proportional to G_{sb} . If we take into account the modulation depth $\Gamma = 0.4$ and a carrier gain G_{cr} of 50 we would predict a ratio of roughly 88 for the carrier power relative to the an individual sideband's power for the cold, untuned state. For a thermally tuned state with $G_{sb} = 26$, we would expect the ratio of roughly 44.

The phase camera shows a mean ratio of 95 for the no TCS state and a ratio of 53 for the good TCS state between the carrier and sideband power. Comparing the actual ratio of of NSPOB of $395/220 = 1.8$, and the ratio of the mean value of the power measured by the sidebands in the two states of $95/53 = 1.8$, we find good agreement. Similarly we expect $297/220 = 1.4$ and $500/220 = 2.3$ based on the NSPOB values in the Reversed TCS and Power up states relative to the No TCS state. The phase camera finds $95/49 = 1.9$ and $95/40 = 2.4$ for the Reversed TCS and Power up states relative to the No TCS state. These comparisons are summarized in Table 4.1. Here the the Reversed TCS state is in very poor agreement with the NSPOB signal, which suggests poor spatial matching of the sidebands might be the dominant effect.

As noted in Section 2.2.2, the AS_I signal can be generated either by an amplitude difference between the sidebands or higher order carrier modes beating with higher order sideband modes. Table 4.2 lists the ratio of the power in the upper sideband to the ratio

Table 4.1: NSPOB and sideband power ratios

	No TCS	Tuned TCS	Reversed TCS	Power up with TCS
NSPOB ratio (relative to No TCS case)	1	1.8	1.4	2.3
Ratio of Sideband Power (relative to No TCS case)	1	1.8	1.9	2.4

of the power in the lower sideband for the various states, which should be compared to the AS_I_Corr signal on the same table. In the Tuned TCS case, where the AS_I_Corr signal has effectively been nulled, the upper and lower sidebands have an average mismatch smaller than the uncertainty of the overall measurement. Effectively, they have been balanced in power successfully by the TCS system.

In the No TCS state the upper sideband is around 9% larger than the lower sideband, which accounts for the non-zero AS_I_Corr signal. The Power up state also has a similar amount of mismatch with, but in the opposite direction with more lower sideband than upper sideband. In both cases, approximately 5% to 15% mismatch required the same correction on the AS photodiode.

The Reversed TCS state interestingly has a very close match in overall power in the sidebands, but required over twice the correction as the No TCS and Power up states. This implies a spatial mismatch is dominating the AS_I_Corr signal. An example set of images from the Reversed TCS state is present in Figure 4-5. If we compare with example images from the other three states, shown in Figures 4-6, 4-7 and 4-8, we see the spatial form of the sidebands are severely distorted, with a corresponding distortion in the carrier on sideband 24.5 MHz signal. Certainly in the No TCS and Tuned TCS states, the sidebands have a much better spatial distribution matching the extent of the carrier.

The Reversed TCS case in Figure 4-5 clearly shows a mismatch between the spatial extent of the sidebands and the carrier. This affects the carrier on sideband beat resulting in only a small, non-symmetric portion of the carrier contributing to the AS photodiode signal. The plot of the difference of the amplitudes between the two sidebands clearly shows a mismatch, which is large when compared to the other heating cases. So even though the total power in the sidebands are similar, the difference in spatial structure results in poor beating and less NSPOB signal. There is a small difference in the position of the peaks as

Table 4.2: Heating states and phase camera power ratios

	No TCS	Tuned TCS	Reversed TCS	Power up with TCS
NSPOB (counts)	220	395	297	500
AS_I_Corr (counts)	-200	0	-450	-200
Carrier/USB	90 ± 2	52 ± 2	50 ± 2	42 ± 2
Carrier/LSB	99 ± 4	54 ± 2	49 ± 3	38 ± 2
USB/LSB	1.09 ± 0.05	1.03 ± 0.04	0.98 ± 0.05	0.9 ± 0.05

measured at the phase camera, roughly 10% of the waist size of the sidebands in this case, where the waist size is the radius where the amplitude of the sideband falls by $1/e$.

The No TCS case in Figure 4-6 has the sidebands matching the carrier in spatial extent, although their overall power and structure are different from each other. These differences can lead to AS_I signal when the sidebands are beat against the carrier. The sideband on carrier signal, effectively seen at the AS port photodiode, is the same size as the carrier signal, with a null band due to the difference in the curvature of the sidebands and carrier signals on the phase camera photodiode. The asymmetry in the carrier on sideband signal is due to the structure in the sidebands.

The Tuned TCS case in Figure 4-7 shares many similarities with the No TCS case, except that the sidebands look more Gaussian, are better balanced in power and shape, and the carrier on sideband signal is somewhat cleaner, with less overall power present. Even though the central peak is slightly higher than in the carrier on sideband image than in the No TCS case, there is less “junk” in the outer sections, reducing the overall power. All these changes indicate less AS_I signal should be seen by the AS photodiode and that the AS_I_Corr signal required to balance it should decrease relative to the No TCS case, which was in fact the observed effect.

The Power up with TCS case in Figure 4-8 shows effects similar to the Reversed TCS case, namely smaller spatial extent for the sidebands, although not to the same degree. This is because the TCS servo, which was designed to null the AS_I_Corr signal, was not fully tuned at the time this data was taken, and resulted in less compensation than was required. The unusual shape of the carrier on sideband signal is related to the squashed nature of the sidebands, although it is not precisely understood where the source of the asymmetry relative to the carrier is coming from. Due to time constraints before the start

of LIGO's 5th science run, data collection from different points of the interferometer were never completed, limiting our ability to understand the source of these differences. One possibility is uneven absorption and heating in the ITMs by the main interferometer laser. Comparing the sidebands we see the lower sideband is spatially smaller than upper sideband, but with nearly twice the peak power at its center. Such amplitude mismatches as noted earlier are a clear source of AS-I which needs to be corrected.

Taken together the images from the various heating states show a trend of shrinking sideband spatial extent with increased heating of the ITMs, whether by the main interferometer laser, or the TCS system operating in such a way as to add heat. They also show significant differences between the sidebands in spatial shape, especially in the high power state. However, they also show with a properly tuned TCS system, these differences in the sidebands can be reduced, balancing their overall and peak power, producing a better match between each other and the carrier, which corresponds with a drop in the necessary AS-I correction and allowing for higher power operation with the same AS port photodiodes and electronics.

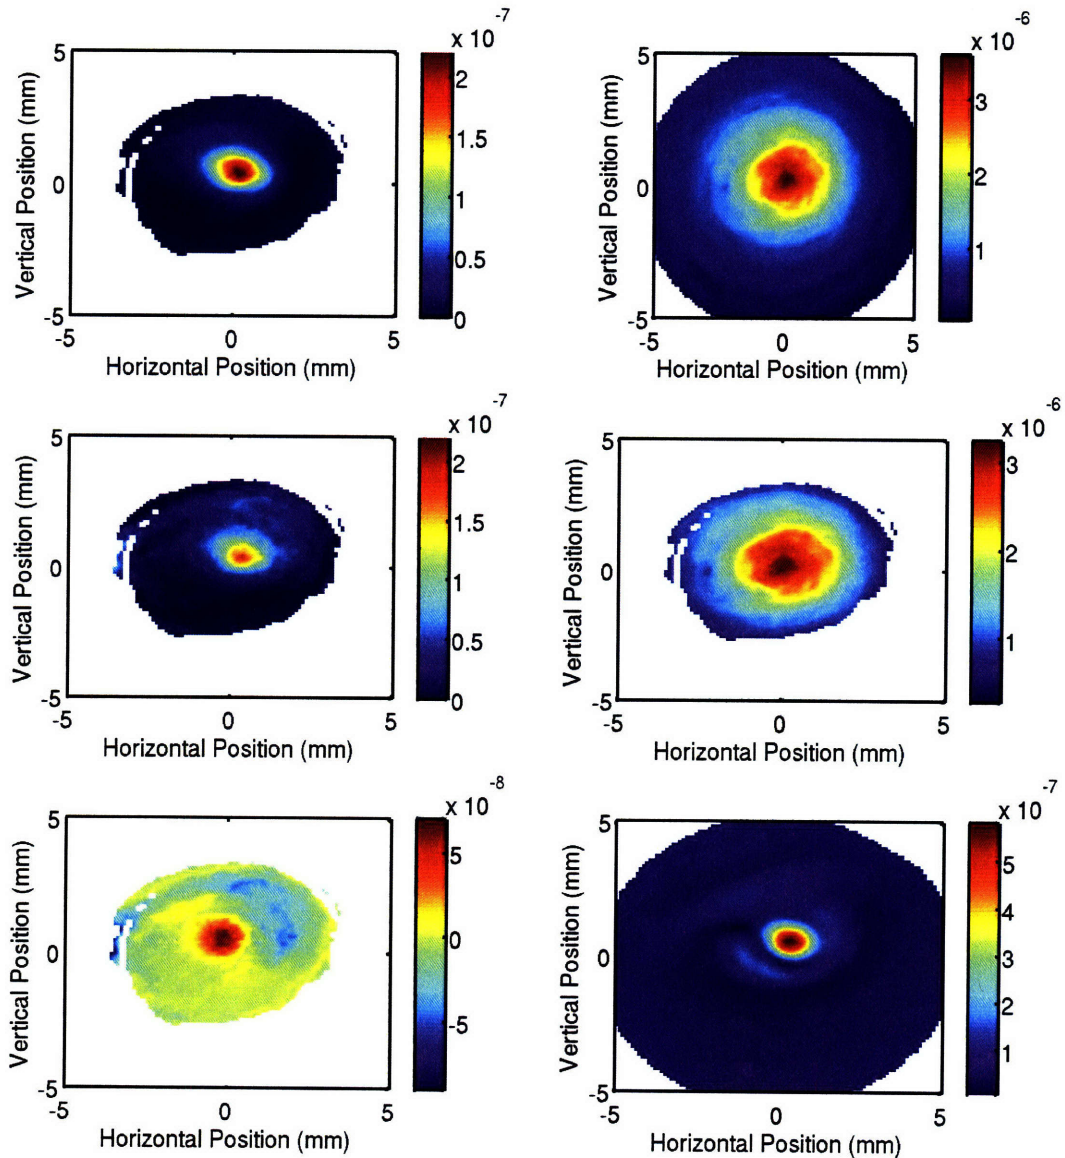


Figure 4-5: The above images are from the Reversed TCS case. The color scales are power in watts. The axes are position in millimeters, measured at the position of the phase camera photodiode. The top right plot shows the power from DC signal. The top left plot shows the power in the upper sideband, demodulated at 50.5 MHz. The middle left plot shows the lower sideband, demodulated at 99.5 MHz. The middle right plot shows the power in the carrier, demodulated at 75 MHz. The lower left plots shows the subtraction of the lower sideband from the upper sideband. The lower right plot shows the carrier on sideband beat, demodulated at 49 MHz. The upper sideband, lower sideband, sideband difference, and carrier images have had the effects of the reference beam removed.

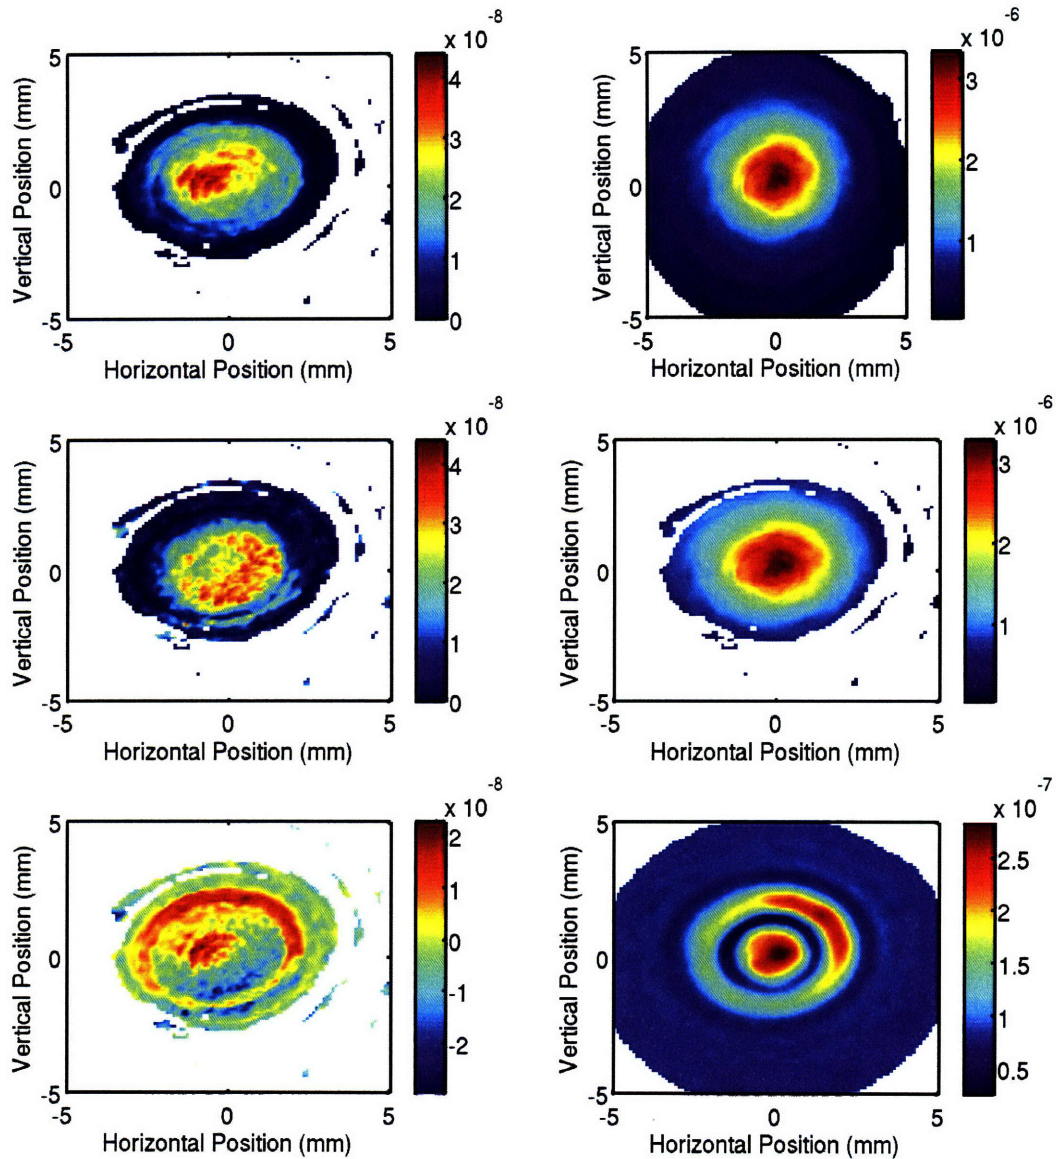


Figure 4-6: The above images are from the No TCS case. The color scales are power in watts. The axes are position in millimeters, measured at the position of the phase camera photodiode. The top right plot shows the power from DC signal. The top left plot shows the power in the upper sideband, demodulated at 50.5 MHz. The middle left plot shows the lower sideband, demodulated at 99.5 MHz. The middle right plot shows the power in the carrier, demodulated at 75 MHz. The lower left plots shows the subtraction of the lower sideband from the upper sideband. The lower right plot shows the carrier on sideband beat, demodulated at 49 MHz. The upper sideband, lower sideband, sideband difference, and carrier images have had the effects of the reference beam removed.

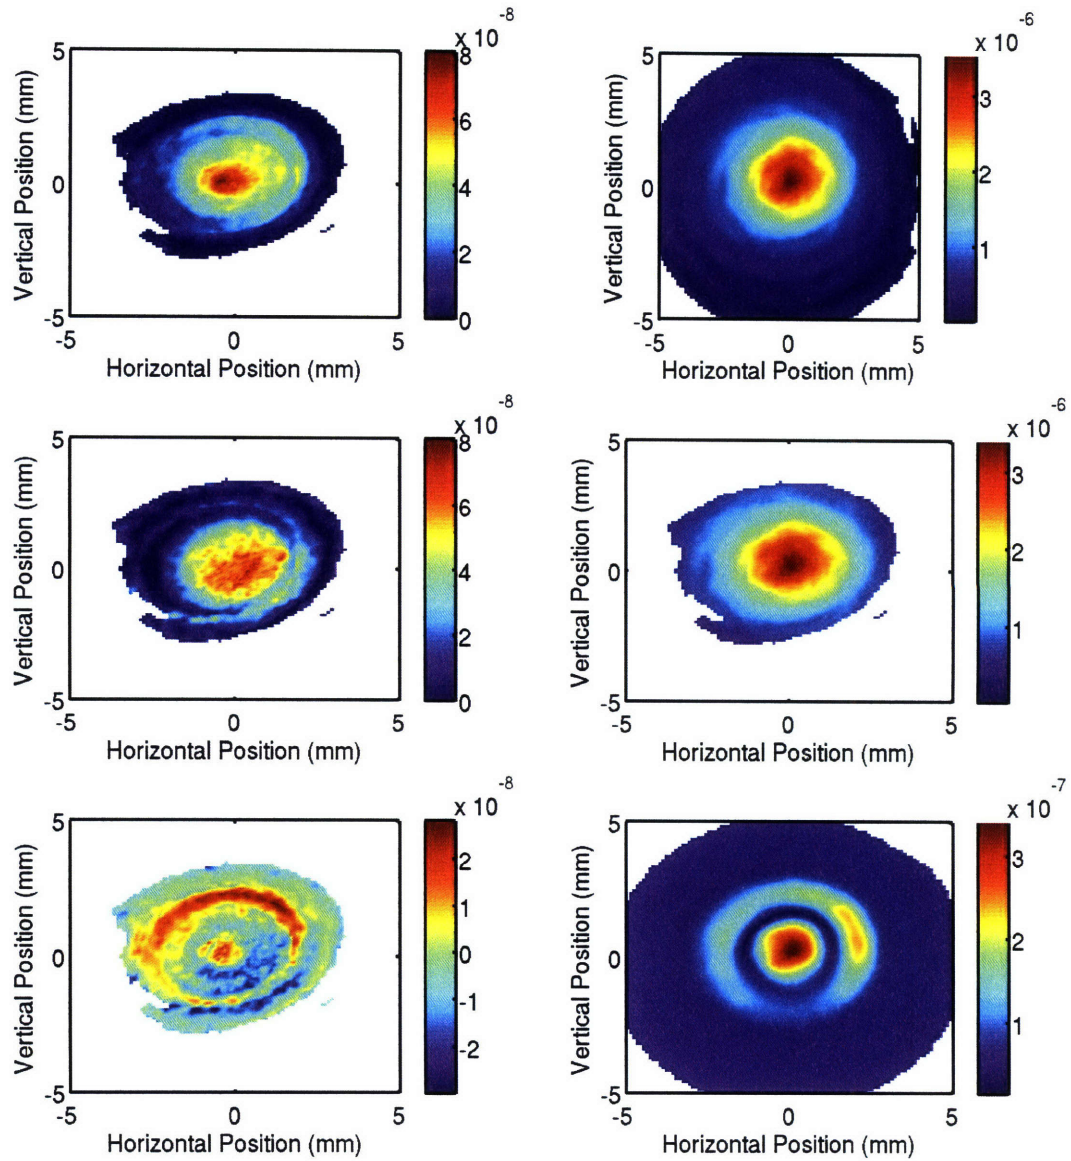


Figure 4-7: The above images are from the Tuned TCS case. The color scales are power in watts. The axes are position in millimeters, measured at the position of the phase camera photodiode. The top right plot shows the power from DC signal. The top left plot shows the power in the upper sideband, demodulated at 50.5 MHz. The middle left plot shows the lower sideband, demodulated at 99.5 MHz. The middle right plot shows the power in the carrier, demodulated at 75 MHz. The lower left plots shows the subtraction of the lower sideband from the upper sideband. The lower right plot shows the carrier on sideband beat, demodulated at 49 MHz. The upper sideband, lower sideband, sideband difference, and carrier images have had the effects of the reference beam removed.

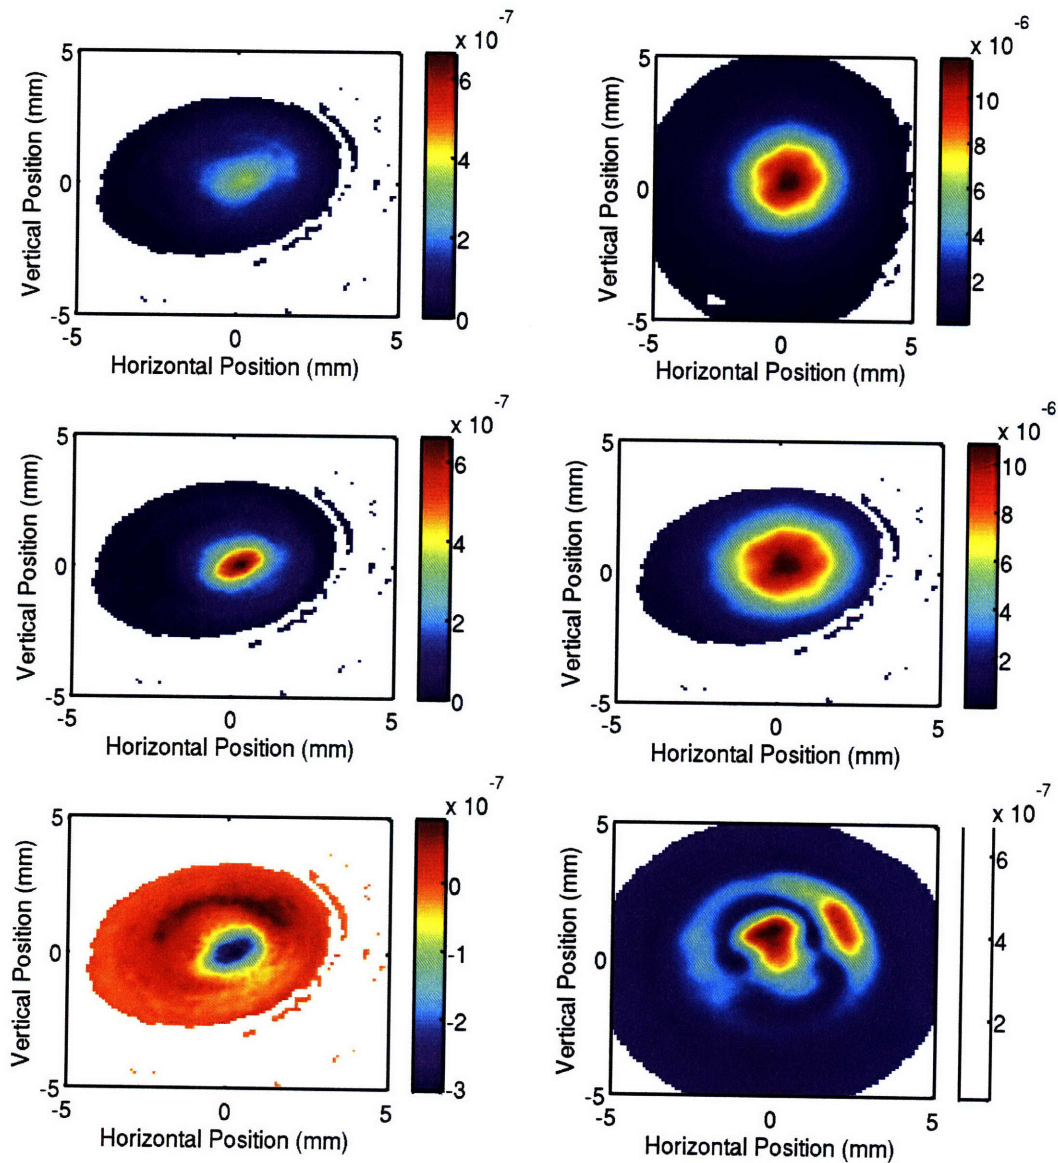


Figure 4-8: The above images are from the Power up with TCS case. The color scales are power in watts. The axes are position in millimeters, measured at the position of the phase camera photodiode. The top right plot shows the power from DC signal. The top left plot shows the power in the upper sideband, demodulated at 50.5 MHz. The middle left plot shows the lower sideband, demodulated at 99.5 MHz. The middle right plot shows the power in the carrier, demodulated at 75 MHz. The lower left plots shows the subtraction of the lower sideband from the upper sideband. The lower right plot shows the carrier on sideband beat, demodulated at 49 MHz. The upper sideband, lower sideband, sideband difference, and carrier images have had the effects of the reference beam removed.

Part III

Targeted continuous gravitational wave searches

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Neutron stars, spindown models, and estimating upper limits of sources

This chapter describes some basics of neutron stars and discusses some basic models used to estimate potential gravitational wave signals. When looking at sources with missing information, such as neutron stars with unknown spin period and spin down, it is useful to make estimates for the best possible cases to determine if its worth putting computation time into the search. There are several simple but very informative calculations one can go through in this regard.

5.1 Neutron stars and pulsars

Neutron stars are the end point of stellar evolution for intermediate mass stars, those starting with a mass of roughly 6 to 15 M_{\odot} [37]. Lighter stars will become white dwarfs while more massive stars will end as black holes. As the nuclear fuel of a star becomes exhausted, a core grows within an expanding outer shell, the total gravitational collapse of the core being prevented by electron degeneracy pressure. If the mass of the core grows greater than 1.4 M_{\odot} , the electron degeneracy pressure is no longer great enough to prevent further collapse of the core. It begins to rapidly collapse until stopped by neutron degeneracy pressure. However, by this time the collapse has released a large amount of gravitational

potential energy in a few seconds, producing a supernova. Most of the original mass is blown off during the supernova, but the remainder of the star still has a mass on the order of $1 M_{\odot}$.

Neutron stars for which there are measured masses typically are found near $1.35 M_{\odot}$. The theoretical upper limit on the mass of a neutron star is somewhere around $3 M_{\odot}$, and is dependant which equations of state are used, which relate the density, pressure and temperature inside the star. Using the more plausible equations of state the radius of a typical neutron star is $\sim 10 - 11$ km. Modeling a neutron star as a sphere with uniform density allows us to estimate the moment of inertia, I , as 10^{38} kg m². This moment of inertia more of a scale rather than an exact number, because of the differences between possible equations of state and other particular details of any given neutron star. After comparing a range of equations of states, Cook, Shapiro and Teukolsky found a range of 0.5×10^{38} to 8×10^{38} kg m² for the moment of inertia [15].

Pulsars are observed from the regular electromagnetic pulses that they emit. These pulses are due to the magnetic field of the generating star, whose axis in general is not aligned with the rotation axis. As the star rotates, the magnetic field produces magnetic dipole radiation, which would be highly beamed out from the magnetic poles and sweep across space. If that beam intersects with a detector on Earth, it will be observed as a short pulse, whose period is directly related to the period of rotation of the star. Periods of pulsars vary from tens of seconds to milliseconds. The necessary rotation rates to produce pulses with these periods means pulsars must be neutron stars, as other less compact stars can't rotate that fast and remain together. Observations of the power and frequency of pulsar emission provides insight into the strength of the magnetic field necessary to produce them. The estimated field strengths vary from 10^{12} Gauss down to 10^8 Gauss, depending on the pulsar. These dipole field strengths are consistent with the collapse of a normal star with a polar field of order 100 gauss and the flux being conserved in the collapsing stellar material.

The strongest polar fields of order 10^{12} Gauss occur in young pulsars, with spin periods of order tens of milliseconds. The majority of pulsars have fields of order 10^{10} Gauss and spin periods of order seconds. There is a third population, referred to as millisecond pulsars because of their spin periods of order 1 millisecond. These millisecond pulsars are the oldest pulsars and have magnetic fields as low as 10^8 Gauss [24]. In general, the spin period of the pulsar pulses decrease with time (and thus decreasing frequency), with a measured rate of

changes in the period over the period, \dot{P}/P , between 10^{-18} to 10^{-12} s $^{-1}$. However, in the case of millisecond pulsars, they instead seem to have been spun up (and in some cases still being spun up) by accretion from a low mass companion star. This is supported by the fact that well over half of millisecond pulsars are found in binaries. Millisecond pulsars appear to be a smaller population of the primary body of pulsars, found through out the galaxy and especially in globular clusters, with a longer lifespan but much lower re-birth rate [37].

There is an estimated active population of 10^5 to 10^6 pulsars in our galaxy, although only about 1500 have been observed due to selection effects and strength of the signals [37]. Most of these observed pulsars are within a layer 1 kpc thick centered on the galactic plane and within a radial distance of about 10 kpc of the center of the galaxy. Most have high velocities, presumably from the initial supernova explosion which formed them, and are moving away from the plane at around 200 km s $^{-1}$. This is consistent with an origin of within 100 pc of the galactic plane and their lifetime of roughly 10 million years for their active stage. Using the current estimated numbers of pulsars, their lifetime of 10 million years, and assuming a roughly constant birth rate would imply a population on the order of 10^9 neutron stars (out of the 10^{11} total stars) in our galaxy [37].

5.2 Estimated deformations of neutron stars

It is important to consider the ellipticities of neutron stars as this plays directly into the strength of any gravitational wave reaching us and influences heavily what we can detect. Both the maximum and most likely ellipticities of neutron stars are highly uncertain [2]. An upper limit estimate on the maximal deformation supported by the rigidity of a neutron star's crust is [51]

$$\epsilon_{max} \approx 5 \times 10^{-7} \left(\frac{\sigma}{10^{-2}} \right) \quad (5.1)$$

where σ is the breaking strain of the solid crust. The numerical coefficient is small mostly because the shear modulus of the inner crust, where most of the crust's mass is located, is smaller than the pressure by a factor of 10^3 . Shear modulus is the shear stress over the shear strain and as it goes to zero, the crust becomes a liquid which can't support non-axial deformations. Thus, the small shear modulus results in a small maximum ellipticity. Equation (5.1) uses the fiducial breaking strain of 10^{-2} since that is approximately the best

terrestrial alloys have. However, for a perfect crystal with no defects σ could be as high as 10^{-1} , or several orders of magnitude smaller for amorphous solids or a crystal with many defects [32].

There have been suggestions that exotic and speculative star types could support larger maximum ϵ values. The highest ellipticity model is that of a solid strange-quark star which has [42]

$$\epsilon_{max} \approx 4 \times 10^{-4} \left(\frac{\sigma}{10^{-2}} \right). \quad (5.2)$$

The reason this maximum ϵ value is so high is due to the high shear modulus, which for some strange star models can almost be as large as the pressure. Another speculative but more robust model is the hybrid star. Such a star would consist of a normal neutron star outside a solid core of mixed quark and baryon matter, extending from the center to just below the crust. Such hybrid star models have a maximum ellipticity of [42]

$$\epsilon_{max} \approx 9 \times 10^{-6} \left(\frac{\sigma}{10^{-2}} \right). \quad (5.3)$$

However, the maximum supportable ellipticity is not the same as being able to achieve the maximum ellipticity in practice. We can consider a few mechanisms that might produce ellipticity in a neutron star. In the case of young pulsars, as they spin down it is likely stresses build up, possibly because of reduced centrifugal force due to slowing rotation, until the crust is forced to change shape and crack. Such cracking could produce deformations or “mountains” on the surface. It is unclear how long it would take gravity to smooth out these deformation after cracking, which could possibly leave long lived distortions [14].

Strong internal magnetic fields can also produce ellipticities in neutron stars. Such fields if not aligned with the rotation axis, could produce non-axisymmetric deformations due to magnetic tension, effectively pulling on the matter of the star. A strong dipolar field not aligned with the rotation axis in a star with a superconducting interior, which has expelled all the magnetic field lines from the superconducting region, could lead to deformations of order 10^{-6} [11]. This deformation would result in gravitational wave emission at both the rotation frequency and twice the rotation frequency.

Torodial magnetic fields could also generate ellipticity. Differential rotation immediately after the core collapse which formed the neutron star can produce large internal toroidal

magnetic fields. The magnetic fields lines trapped with in the different layers are effectively wrapped into a toroidal shape as they rotate at different speeds around the star. Dissipation within the star would tend to drive the symmetry axis of this toroidal field toward the star's equator, the orientation which maximizes ellipticity. For toroidal magnetic fields on the order of 10^{15} Gauss, an ellipticity of order $\sim 10^{-6}$ would occur [16].

5.3 Energy considerations

Neutron stars when born will in general have some spin angular momentum. This angular momentum acts as an energy reservoir which is used to power various phenomenon including radio pulsation, ejection of particles from the surface, and gravitational waves. Each of these activities cause the star to spin down over time which will reduce the frequency of any emitted gravitational waves. This evolution of the spin period turns out to be large enough that it needs to be accounted for during any gravitational wave searches of reasonable length. It also gives insight into how the parameter space of ellipticity, frequency, and emitted gravitational wave strength is populated by isolated neutron stars. Neutron stars with companions can undergo spin up via accretion and will tend to an equilibrium state between the spin up and the spin down mechanisms, thus the following arguments do not apply in such cases. In this section we will derive equations describing the spin down and then apply them to the case of gravitational radiation.

Let us assume an isolated neutron star has a moment of inertia I and angular frequency Ω so that we can write the kinetic energy as

$$E_{kin} = \frac{1}{2}I\Omega^2 . \tag{5.4}$$

If we take the time derivative of the energy we have

$$\dot{E}_{kin} = I\Omega\dot{\Omega} . \tag{5.5}$$

If we assume that all of the energy loss is due to a single emission mechanism, we can consider a general relationship between the rate of change in the angular frequency, $\dot{\Omega}$, and the angular frequency, Ω , of the form

$$\dot{\Omega} = K\Omega^n \quad (5.6)$$

where K is the torque function which contains all the physics of the torque causing the spin down and n is the braking index. The braking index has a value of 3 for spin down due to dipolar magnetic radiation and a value of 5 for spin down due to gravitational radiation. This can be seen from the fact that the power emitted in dipolar magnetic radiation goes as the 4th power of frequency while the power in gravitational quadrupole radiation goes as the 6th power of frequency. Eq. (5.5) relates this power emission to Ω and $\dot{\Omega}$ to get n .

By taking the time derivative of Eq. (5.6) we can solve explicitly for the braking index in terms of observables Ω , $\dot{\Omega}$, and $\ddot{\Omega}$. The braking index is then just [43]

$$n = \frac{\ddot{\Omega}\Omega}{\dot{\Omega}^2}. \quad (5.7)$$

If such values are known, by observation of radio pulsations for example, it gives us insight into what mechanisms might be causing the spin down.

In the case of unknown frequency and spin down rate, we can find a relation between the frequency and spin down if we know the approximate age of the neutron star and assume the ellipticity is constant. By using Eq. (5.6) we can solve for Ω as a function of time. Moving all the factors of Ω and $\dot{\Omega}$ to one side, we can integrate with respect to time with the boundaries set at zero, the time of initial angular frequency Ω_0 , and the current time t_{age} , which corresponds to angular frequency Ω_t . This is written as

$$\int_0^{t_{age}} K dt = \int_{\Omega_0}^{\Omega_t} \frac{d\Omega}{\Omega^n} \quad (5.8)$$

and after performing the integration we get

$$Kt_{age} = \frac{-1}{(n-1)\Omega_t^{(n-1)}} \left(1 - \frac{\Omega_t^{(n-1)}}{\Omega_0^{(n-1)}} \right). \quad (5.9)$$

If we assume that the current rotation frequency is small compared to the initial rotation frequency, $\Omega_t \ll \Omega_0$, we can write a simple relation between the current frequency, the current rate of frequency change, and the approximate time since the neutron star was born (also known as the spin down age). This approximation is good to 1% in the case of gravitational radiation, the $n = 5$ case, when $\Omega_t/\Omega_0 < 0.3$. Using this approximation and

the fact that $K = \dot{\Omega}/\Omega^n$, from Eq. (5.6), we solve for t_{age} to get

$$t_{age} = \frac{1}{n-1} \left| \frac{\Omega_t}{\dot{\Omega}_t} \right| = \frac{1}{n-1} \left| \frac{f}{\dot{f}} \right|. \quad (5.10)$$

We can now apply the above derivations to the specific case of gravitational radiation. We know that the gravitational wave luminosity is equal to [40]

$$L_{gw} = \frac{1}{5} \langle \ddot{I}_{uv} \ddot{I}^{uv} \rangle. \quad (5.11)$$

We already know \ddot{I}_{uv} for a non-axisymmetric neutron star with some ellipticity ϵ from Eq. (1.13) so we can substitute it in and simply get

$$L_{gw} = \frac{32}{5} \Omega^6 I^2 \epsilon^2. \quad (5.12)$$

If we make the assumption that spin down is due only to gravitational wave emission we can simply set the energy loss, \dot{E}_{kin} equal to the gravitational luminosity, L_{gw} . This is a best case assumption as then all the energy leaving the star is in a form we can detect. If other energy loss processes are occurring this will increase the energy loss and spin down, resulting in slower spinning stars and weaker signals. Having made this assumption, we can calculate an upper limit on the gravitational wave strain at Earth if we know the distance, frequency and spin down rate, such as for the Crab pulsar [43]. This is done in Section 5.6.

In the case of unknown frequency and spin down rate, we can use Eq. (5.13) to limit the range of parameter space that needs to be considered if the age of the neutron star is known. As can be seen by comparing Eqs. (5.5), (5.6), and (5.12), the braking index n is equal to 5 for gravitational radiation. This lets us rewrite Eq. (5.13) as

$$t_{age} = \frac{1}{4} \left| \frac{\Omega_t}{\dot{\Omega}_t} \right| = \frac{1}{4} \left| \frac{f}{\dot{f}} \right| \quad (5.13)$$

for the gravitational case. This relationship is useful in setting up searches for gravitational waves from isolated neutron stars with partially known parameters. If we know an approximate age we can estimate the largest possible strain reaching a detector here at earth across all frequencies and what ellipticity of the neutron star makes it possible.

5.4 Differences between GW and EM pulses

For targeted searches with radio pulse data, such as the Crab pulsar, it is important to consider physical situations in which the gravitational wave frequency might differ from the electromagnetic frequency or multiples of it. If the frequencies are not exactly matched, then a wider search may be necessary to cover the parameter space sufficiently, so as not to miss the true frequency and spin down of the gravitational radiation.

5.4.1 Free precession

The case of steady rotation, when the principle axis and the inertial axis I_3 are aligned was introduced in the Chapter 1. This results in gravitational radiation at twice the rotation frequency, $\Omega_{GW} = 2\Omega$. In this case, a hot spot which is producing the radio pulses will come around once per rotation of the pulsar, producing a pulse frequency of $\Omega_{EM} = \Omega$. In this simple case, we only need to look at exactly twice the electromagnetic pulse frequency.

If the rotation axis of the Crab pulsar is not aligned with the principle axis, it could precess like a top [56] [30]. Let us consider the slightly more complex situation of biaxial free precession. In this case, $I_1 = I_2 < I_3$. We define a constant angle θ between the body's third principle axis x_3 and the total angular momentum vector \vec{J} axis of rotation, which has an inclination angle i relative to the direction of the observer as shown in Figure 5-1. We then have two rotation periods described by

$$\dot{\phi} = \frac{J}{I_1} \quad (5.14)$$

and

$$\dot{\psi} = -\frac{I_3 - I_1}{I_1} \dot{\phi} \quad (5.15)$$

where $\dot{\phi}$ represents the inertial space x_3 axis angular velocity and $\dot{\psi}$ represents the body frame precessional angular velocity.

In this case gravitational waves can be calculated to have the form

$$h_+ = \frac{2I_1 \dot{\phi}^2 \epsilon \sin(\theta)}{r} \left[(1 + \cos^2(i)) \sin(\theta) \cos(2\dot{\phi}t) + \cos(i) \sin(i) \cos(\theta) \cos(\dot{\phi}t) \right], \quad (5.16)$$




Figure 5-1: Diagram of a neutron star with a bump on one side. The angle θ is the angle between the I_3 moment of inertia and the angular momentum vector. The red circle represents the location where the EM pulse emission is being generated. The angle λ is the angle between angular momentum vector and the point on the surface where the EM emission occurs. The angle i is the inclination angle, the angle between the angular momentum vector and the vector towards the detector.

and

$$h_x = \frac{2I_1\dot{\phi}^2\epsilon \sin(\theta)}{r} \left[\cos(i) \sin(\theta) \sin(2\dot{\phi}t) + \sin(i) \cos(\theta) \sin(\dot{\phi}t) \right] \quad (5.17)$$

with $\epsilon = (I_3 - I_1)/I_1$ and r being the distance to the source.

However, the electromagnetic pulses have a more complicated frequency in this case. Let us define the angle between the pulsar beam and the x_3 axis as λ . If in general $\lambda \neq 0$ then there will be modulation of the electromagnetic pulse frequency.

When the angle θ is larger than the angle λ , the pulse would be seen when the x_3 pulsar axis passes close enough to the observer's line of sight. During the body's precession time the observer would pass through the pulsar radiation beam from many different directions.

For the precession to be invisible, the pulsar beam would have to both be nearly axisymmetric and could not have any observable linear polarization. Any net linear polarization would rotate through 360° during a precession time, since the orientation of the electromagnetic emission source would rotate completely around the body axis. The electromagnetic pulse frequency is just $\dot{\phi}$. This can be seen from Figure 5-2. The rotation about the star's I_3 axis primarily changes the orientation of the source while the rotation about the x_3 inertial frame axis brings the pulse around like a light house. For this case, we can write

$$\theta > \lambda \Rightarrow \langle \Omega_{EM} \rangle = \dot{\phi}. \quad (5.18)$$

The angle λ being larger than the angle θ is more likely to be occurring without producing observable effects in the radio timing. In this case the mean electromagnetic pulse frequency is the sum of precession angular velocity, $\dot{\psi}$, and the inertia space x_3 angular velocity, $\dot{\phi}$. This can be seen from Figure 5-3, where the electromagnetic source point is being rotated about the star by both the inertial space angular velocity and the angular velocity about the I_3 axis. The observer always passes through the pulsar beam from approximately the same direction, so no significant changes in pulse profile or polarization would be expected. A simple knife beam model of the pulsar radiation pattern gives the result for small θ that during a precession time pulses arrive early and late by a phase of up to $\theta/\tan(\lambda)$, with sinusoidally varying phase shift. We can write the time averaged electromagnetic pulse frequency for this case as

$$\theta < \lambda \Rightarrow \langle \Omega_{EM} \rangle = \dot{\phi} + \dot{\psi}. \quad (5.19)$$

By comparison, the gravitational wave emission would be observed at $\dot{\phi}$ and $2\dot{\phi}$.

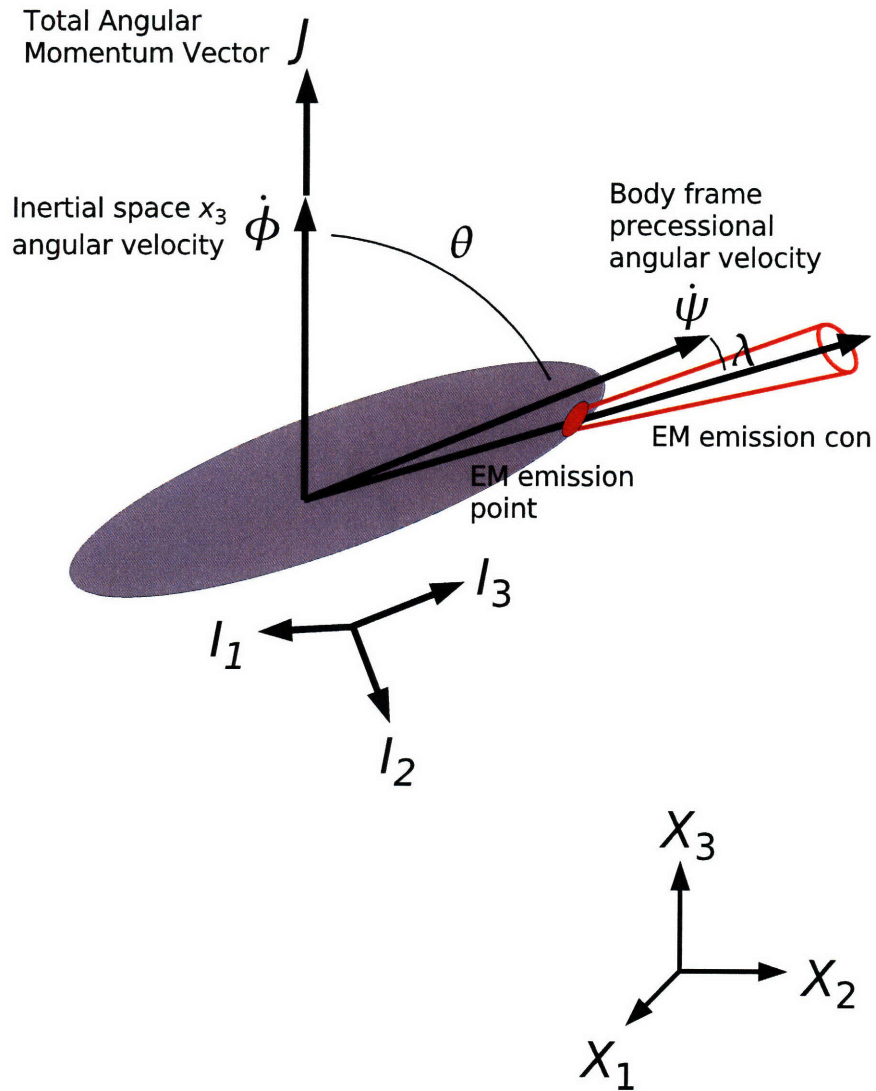


Figure 5-2: Exaggerated diagram of a biaxially symmetric neutron star with $\lambda < \theta$. The electromagnetic emission pulse frequency is determined by the angular velocity $\dot{\phi}$, as the star rotates around the inertial space axis x_3 . The rotation about the I_3 moment of inertia axis of the star will cause small variations in that pulse frequency, but which average out over many precession cycles. The rotation about the I_3 axis will also result in a changing orientation, which if linear polarization is present, will result in that polarization rotating through 360° over the course of a single precession.

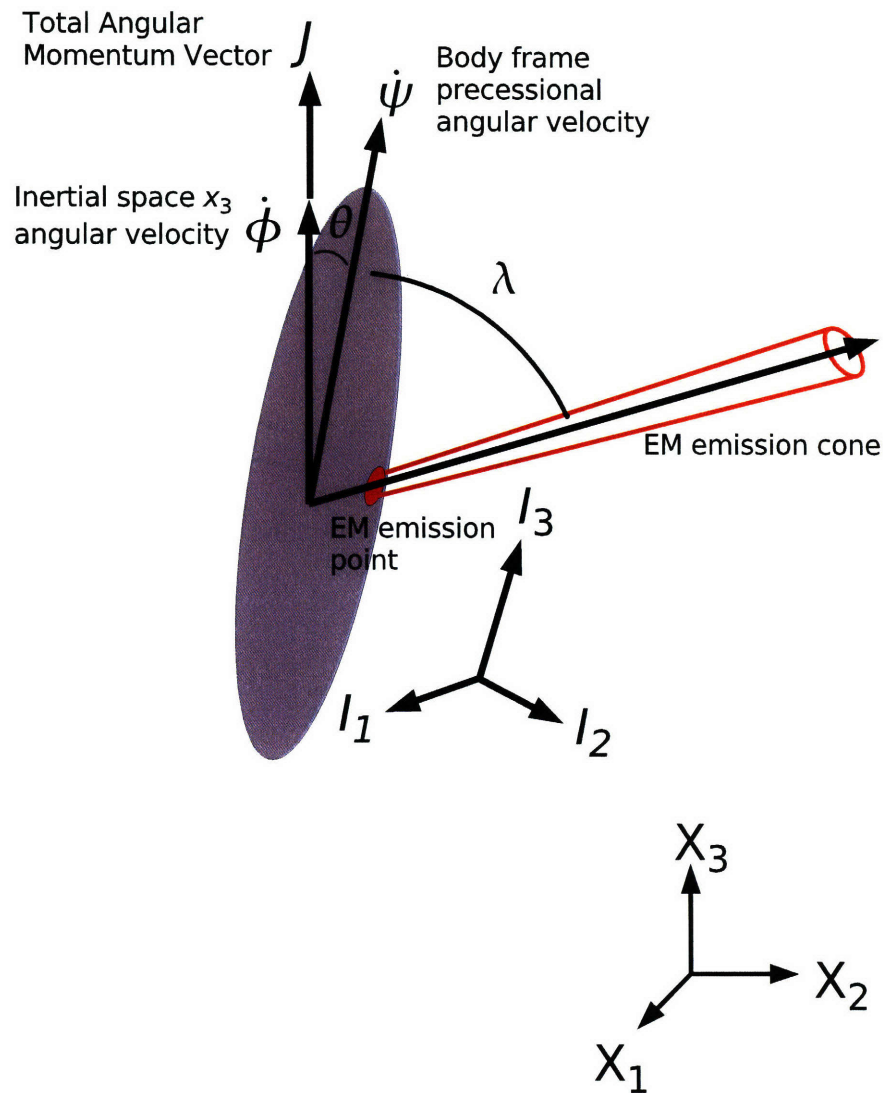


Figure 5-3: Exaggerated diagram of a biaxially symmetric neutron star with $\theta < \lambda$. The inertial space x_3 angular velocity, $\dot{\phi}$ and the body axis angular velocity, $\dot{\psi}$, on average add together to determine the total angular velocity of the electromagnetic wave source. This is because the rotation vectors are close and the electromagnetic source never passes in between the x_3 and I_3 axes. The rotations are in the same direction in this case, where as in the case shown in Figure 5-2, the body axis angular velocity adds as often as it subtracts from the overall angular velocity, which when averaged over time doesn't modify the $\dot{\phi}$ rotation rate. This also means the electromagnetic source is always in approximately the same orientation when facing an observer. It can not flip over, relative to this diagram, by passing in between the x_3 and I_3 axes.

In the more complicated case of triaxial precession, where $I_1 < I_2 < I_3$, one needs to make some assumptions to be able to determine approximate waveforms. If we assume $\theta \ll 1$, which is just a small wobble angle, $(I_3 - I_1)/I_1 \ll 1$, which is small oblateness, and $(I_2 - I_1)/I_1 \ll (I_3 - I_1)/I_1$, which means the non-axisymmetry is small compared to the oblateness, then the wave form can be approximated by [55]

$$h \sim (I_2 - I_1) \cos(2(\dot{\phi} + \dot{\psi})t) + \sim (I_3 - I_1)\theta \cos(\dot{\phi}t) + \sim (I_3 - I_1)\theta^2 \cos(2\dot{\phi}t). \quad (5.20)$$

Each term in this equation is an approximation with different prefactors being neglected, hence the somewhat unusual usage of \sim , where $x + \sim y$ means x plus terms which look like y . From Eq. (5.20) we can see that we have gravitational wave emission at the frequencies of $2(\dot{\phi} + \dot{\psi})$, $\dot{\phi}$ and $2\dot{\phi}$. For electromagnetic pulses we expect the same as in Eq. (5.19) and Eq. (5.18).

5.4.2 Two component model

Some neutron stars are observed to have sudden increases in pulsation frequency, which are usually called glitches [37]. These spin ups generally decay away on various time scales. One possible explanation for this is that pulsars consist of two components, such as a crust component and a core component, rotating at different rates. For the Crab and a few other young pulsars approximately a fraction of 10^{-5} of the spin-down is reversed during the glitch. The decay time varies on the order of days to months. We can use this information to produce an estimate of the maximum frequency difference between the two components and thus between the electromagnetic and gravitational frequencies, since the electromagnetic pulse frequency will be tied to the crust while the gravitational frequency could be tied to the core.

One way to approach this information is with a coupling-torque model. First for the crust we define its rotation frequency f_1 and its moment of inertia I_1 . For the core we similarly define its rotation frequency f_2 and its moment of inertia I_2 . If we allow the spin down torque to act on just one component, say the crust, and couple the two components with a friction-like torque we can write [29]

$$I_1 \dot{f}_1 = -\alpha(f_1 - f_2) - T \quad (5.21)$$

and

$$I_2 \dot{f}_2 = \alpha(f_1 - f_2) \quad (5.22)$$

where α parameterizes the strength of the coupling torque and $T > 0$ is the spin down torque. The spin down torque T effectively is $I\dot{\Omega}$ from section 5.3, with the exact evolution of $\dot{\Omega}$ dependent on whether the slow down is due to electromagnetic or gravitational radiation or perhaps some other mechanism. One can show that the steady state lag, Δf , is then

$$\Delta f = f_1 - f_2 = -\frac{T\tau_c}{I_1} \quad (5.23)$$

where

$$\tau_c = \frac{1}{\alpha} \left(\frac{1}{I_1} + \frac{1}{I_2} \right)^{-1} \quad (5.24)$$

and is the coupling timescale on which transients are damped out. This could be expected to come into play after a glitch occurs.

We now use the results of Section 5.3 to estimate a spin down age. This is the timescale on which the pulsar slows down. We begin by using Eq. (5.13) and noting that the torque T is $T = 2\pi(\dot{f}_1 I_1 + \dot{f}_2 I_2)$. We define a new f such that

$$f = \frac{I_1 f_1 + I_2 f_2}{I_1 + I_2}. \quad (5.25)$$

We then note that the torque T is $T = 2\pi(\dot{f}_1 I_1 + \dot{f}_2 I_2) = 2\pi\dot{f}(I_1 + I_2)$. We solve for \dot{f} and substitute into Eq. (5.13). This then lets us write the steady state spin down timescale as

$$\tau_{ss} = \frac{1}{n-1} \frac{2\pi f(I_1 + I_2)}{T}. \quad (5.26)$$

This is just the timescale on which the pulsar takes to spin down to its current frequency of f from some sufficiently fast initial frequency, as described in section 5.3.

We now use Eq. (5.26) to substitute for the torque in Eq. (5.23) and rewrite it as

$$\Delta f = -\frac{1}{n-1} \frac{2\pi f \tau_c}{\tau_{ss}} \frac{I_1 + I_2}{I_1}. \quad (5.27)$$

This relates the lag between the two components to two observed timescales, the glitch recovery timescale and the spin down age. The steady state spin down age can be approximated by the age of the neutron star, which in the case of the Crab pulsar is 10^3 years. Since we observe the neutron star returning to its original spin down rate on a timescale of weeks or months, it is believed that this is also the timescale associated with coupling of the two components.

We identify the torque T with the electromagnetic torque. There would be a gravitational wave torque, but we assume the electromagnetic dominates. Given that glitches in the Crab pulsar always spin it up, this seems like a reasonable assumption. If we included the effects of a gravitational wave torque in the above equations, it would simply make the lag between the two frequencies smaller.

Assuming I_1 is of the order or larger compared to I_2 makes $(I_1 + I_2)/I_1$ of order 1, which is reasonable since we are associating component 1 with the crust [37]. In addition as the difference between the two components is small compared to frequency of the crust's rotation, we can use f_1 as a good approximation to the frequency f . Using all of the above facts, plus the Crab pulsar's braking index of $n = 2.5$, we find an upper estimate of Δf of 3×10^{-3} for the Crab pulsar, which is in fact small compared to its 29.77 Hz pulse frequency.

5.5 Application to RX J1856.5-3754

In the specific case of RX J1856.5-3754, we did not have a known frequency or spin down values. However we did have a reasonable estimate of the age, somewhere between 10^5 and 10^6 years old, and its distance of about 120 pc. By equating the age of the star with its spin down age, and assuming a typical moment of inertia, $I = 10^{38}$ kg m², we can solve for epsilon as a function of the neutron star's possible current frequency, to determine if its consistent with reasonable neutron star models. Similarly, we can calculate the best case strain at any given current frequency to determine if its detectable within the LIGO band.

To begin with, we replace $\dot{\Omega}$ in Eq. (5.6) with $(-1/4)(\Omega_t/t_{age})$. We then solve for ϵ and get

$$\epsilon = \left(\frac{5}{128} \frac{1}{t_{age} I \Omega^4} \right)^{\frac{1}{2}} \propto \Omega^{-2}. \quad (5.28)$$

Since ϵ goes as Ω^{-2} this will cancel with the normal frequency dependance to the 2nd power

, so that the final maximum strain is not dependant on the frequency. Using Eqs. (5.28) and (1.15) we can write the h_+ and h_\times strains as a function of Ω to get

$$|h_+| = \frac{1}{R} \left(\frac{5}{2} \frac{I}{t_{age}} \right)^{\frac{1}{2}} (1 + \cos^2(i)) \quad (5.29)$$

and

$$|h_\times| = \frac{1}{R} \left(\frac{5}{2} \frac{I}{t_{age}} \right)^{\frac{1}{2}} (2 \cos(i)) . \quad (5.30)$$

In the case of pure gravitational spin down, the strain is only a function of age, irregardless of the current spin frequency. In Figure 5-4 the ellipticity is plotted versus potential frequency at three different possible ages for RX J1856.5-3754, along with the expected strain at each age. In these plots, it is assumed that the neutron star has a typical 10^{38} kg m² moment of inertia and the signal has been averaged over possible inclination angles. Also factors of G , the gravitational constant, and c , the speed of light, were reintroduced to get the units into more traditional form. Similarly in Figure 5-5 we plot the 1st spin down value, which is the first derivative of frequency with respect to time, versus current frequency. In the next chapter, we will look at the detectability of these strains.




Figure 5-4: The top plot shows the expected ellipticity if all of the spin down of RX J1856.5-3754 was due to gravitational radiation for three different possible ages. The lower plot shows the best case strain of the gravitational waves reaching Earth. Other spin down mechanisms at work in addition to gravitational radiation would move the ellipticity lines down, requiring less ellipticity but at the same time, reducing the strain reaching the Earth.

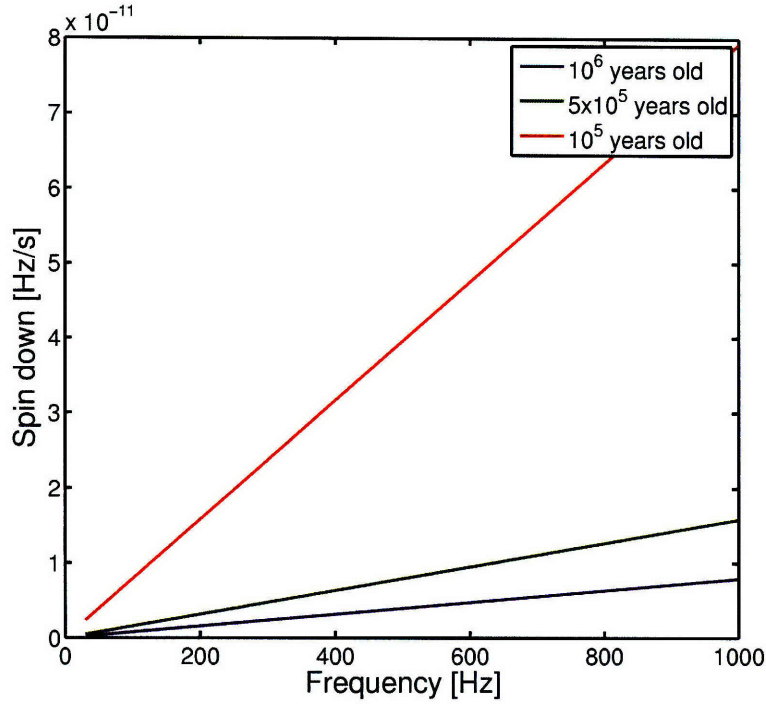


Figure 5-5: The expected 1st spin down value if all the spin down was due to gravitational radiation as a function of current frequency. The younger ages have higher spindown since the neutron star will have had less time to spin down from its initial frequency. To reach the same current frequency as an older age estimate, but in less time, it must be spinning down faster.

5.6 Application to the Crab pulsar

The Crab has been observed enough by radio telescopes that we know its frequency f , its 1st spin down parameter, which is df/dt , and also its 2nd spin down parameter, which is d^2f/dt^2 . The Crab has a frequency of 29.77 Hz, a 1st spin down of $-3.73 \times 10^{-10} \text{ Hz s}^{-1}$, and a 2nd spin down of $1.24 \times 10^{-20} \text{ Hz s}^{-2}$, taken from the Jodrell Pulsar Timing Bank near the beginning of LIGO's fifth science run [28]. So far we have been working in units where $c = G = 1$, but it is easy to convert to more conventional units for the following calculations. If we assume the Crab pulsar has a typical moment of inertia of $I = 10^{38} \text{ kg m}^2$ we can calculate the necessary ellipticity ϵ needed to cause the spin down seen.

$$\dot{\Omega} = -3.7294045 \times 10^{-15} \text{ Hz/s} = -\frac{32}{5} \frac{G}{c^5} I^2 \epsilon^2 \Omega^5 \quad (5.31)$$

Solving for ϵ gives

$$\epsilon = \left(\frac{-5 c^5 \dot{\Omega}}{32 G \Omega^5 I} \right)^{\frac{1}{2}} \quad (5.32)$$

and a value of 7.5×10^{-4} , again assuming a standard $I = 10^{38} \text{ kg m}^2$ [43]. It then becomes a simple matter to calculate the strain using Eq. (1.15) here at Earth if we know the Crab's distance, which is about 2 kpc. This provides an upper limit on the strain of 1.4×10^{-24} . It is interesting to note that while we have been assuming a moment of inertia $I = 10^{38} \text{ kg m}^2$ there have been suggestions that the Crab's moment of inertia could be up to three times larger than this typical value [7] [34]. This would increase the necessary ellipticity to explain the spin down by a factor of $I^{1/2}$, down to 4.3×10^{-4} , and increase the upper limit on the strain by the same factor, up to 2.4×10^{-24} .

We also have information for the 2nd spin down value and can use it to calculate the braking index, n . When we calculate its value from Eq. (5.7) we come up with something close to 2.5, rather than the 5 of gravitational emission, or even the 3 of dipolar magnetic radiation. This indicates that some additional energy loss mechanisms must be occurring.

In 2000 Palomba published a paper using the braking index of 2.5, along with the age of the Crab, to derive upper limits on the ellipticity and gravitational wave emission. He began by assuming the total spin down, $\dot{\Omega}$ was due to two components

$$\dot{\Omega} = \Omega_{GW} + \Omega_{other} \quad (5.33)$$

where Ω_{GW} is due to gravitational radiation and has a braking index of $n_{GW} = 5$ while the Ω_{other} component is due to some unknown emission and which has a braking index n_{other} which is allowed to vary below 3.

He then defined a ratio

$$Y(\Omega) = \frac{\dot{\Omega}_{GW}}{\dot{\Omega}_{other}} = \frac{K_{GW}}{K_{other}} \Omega^{(5-n_{other})} . \quad (5.34)$$

By solving for $\dot{\Omega}_{other}$ in terms of the ratio $Y(\Omega)$ and substituting into Eq. (5.33) he wrote

$$\dot{\Omega} = K_{GW} \Omega^5 \left(\frac{1 + Y(\Omega)}{Y(\Omega)} \right) \quad (5.35)$$

which can be solved numerically to produce a set of solutions for Ω . This set of solutions

depend on the ellipticity, which is contained in K_{GW} , on n_{other} and on the initial angular velocity Ω_0 . By using the known age of the Crab and choosing the solution with the smallest n_{other} consistent with an overall $n = 2.5$, he was able to arrive at an upper limit on ellipticity of 3×10^{-4} . This corresponds to an upper limit on h_0 of 5.5×10^{-25} [43].

Chapter 6

Search methods

This chapter describes several continuous gravitational wave search methods and discusses their merits in directed searches.

6.1 Search codes and constraints

When looking for a continuous wave gravitational source there are several limiting factors which determine the final sensitivity one can achieve. The fixed constraints include the quality and quantity of data available, and also the amount of computation power available. These are combined with the size of the parameter space, which includes position in the sky and frequency parameters of the source, to determine the final sensitivity. Larger searches in parameter space necessarily turn up more outliers due to noise, which in turn degrades any upper limits that might be made from them. Thus, any continuous wave search must weigh these different constraints to determine the optimal method of search to be used.

One can use coherent or incoherent methods to sift the data for a weak, but consistently present source. A coherent search will create a signal template with the known parameters of the source and convolve it with the data. This will produce the most sensitive search since one is looking for the signal with an exact match and utilizes all the data available [19] [27]. An incoherent search will break the data up into shorter time segments and Fourier transform them with the fast Fourier transform algorithm (FFT) and then add up the power at individual frequencies from each data segment. These shorter transformed data chunks are usually referred to as short Fourier transforms (SFTs) [39]. Since phase information is lost between segments, the sensitivity is necessarily worse, but gains in computational speed.

6.2 Detector response function

In order to estimate our final sensitivity, it is necessary to understand how an incoming gravitational wave interacts with our detector. We will follow the derivation by Jaranowski, Krolak and Schutz in [27]. This will then factor in to our calculation of overall sensitivity of the LIGO detectors to RX J1856.5-3754 and the Crab pulsar.

The dimensionless detector's response function h is defined as the difference between the wave induced relative length changes of the two interferometer arms and can be computed from the following formula [27]

$$h(t) = \frac{1}{2} \hat{n}_1 \cdot [\tilde{h}_{ab}(t) \hat{n}_1] - \frac{1}{2} \hat{n}_2 \cdot [\tilde{h}_{ab}(t) \hat{n}_2] \quad (6.1)$$

where \hat{n}_1 and \hat{n}_2 are the unit vectors parallel to the interferometer arms. The order of arms is defined such that $\hat{n}_1 \times \hat{n}_2$ points outwards from the surface of the earth. \tilde{h}_{ab} is the 3-dimensional matrix of the spatial metric perturbation produced by the wave in the proper reference frame of the detector, or in other words, just the spatial components of $h_{\mu\nu}$ as defined in Eq. (1.5) but transformed to the detectors frame. We can compute \tilde{h}_{ab} by

$$\tilde{h}_{ab} = M(t) h_{ab}(t) M(t)^T \quad (6.2)$$

where M is the 3-dimensional orthogonal matrix of transformation from the wave Cartesian coordinates, written (x_w, y_w, z_w) , to the Cartesian coordinates in the detector's proper reference frame, written (x_d, y_d, z_d) , and T indicates transposition of the matrix. If we choose the direction of the gravitational wave propagation as the $+z_w$ direction we can write in the wave's frame that

$$h_{ab}(t) = \begin{pmatrix} h_+(t) & h_\times(t) & 0 \\ h_\times(t) & -h_+(t) & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (6.3)$$

where we have absorbed all the time dependence and distance factors into $h_+(t)$ and $h_\times(t)$, which represent the two independent wave polarizations. If we consider Eqs. (6.1) through (6.3) it becomes apparent that the detector's response function h is a linear combination of the functions $h_+(t)$ and $h_\times(t)$. This lets us write

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t) \quad (6.4)$$

where F_+ and F_\times are called the beam-pattern functions.

The Earth's rotation rotates the detectors along with it causing the beam-patterns F_+ and F_\times to be periodic functions of time with a period equal to one sidereal day. Ideally we want to determine the time dependence of F_+ and F_\times and also their dependence on the source's position and orientation, in terms of right ascension α , declination δ , and the wave's polarization ψ .




Figure 6-1: Diagram of a Interferometer located on the Earth's surface. Ω_r is the Earth's rotation rate, λ is the latitude, and ϕ_r is a reference phase which defines the position of the Earth in its diurnal motion at $t = 0$. The inset shows a pair of interferometer arms as if looked down on from above. The angle ζ is the angle between the arms (90° in the case of LIGO) while γ is the angle between East and the arm bisector.

We can write the matrix M as

$$M = M_3 M_2 M_1^T \quad (6.5)$$

where M_1 is the transformation matrix from the wave frame coordinates to the celestial frame coordinates, M_2 is the transformation matrix from celestial coordinates to the cardinal coordinates and M_3 is the transformation matrix from cardinal coordinates to the detector proper reference frame coordinates. The definition of these coordinate systems are as follows: In the celestial sphere coordinates the z axis coincides with the Earth's rotation axis and points towards the North pole, and the x and y axes lie in the Earth's equatorial plane, with the x axis pointing towards the vernal point. In cardinal coordinates the (x, y) plane is tangent to the surface of the Earth at the detector's location with the x axis aligned in the North-South direction and the y axis aligned in the West-East direction, and the z cardinal axis along the Earth's radius pointing outwards from the surface of the Earth. In detector coordinates the z axis coincides with the z axis of cardinal coordinate and the x axis is along the first interferometer arm. With these definitions we can write M_1 , M_2 , and M_3 as follows

$$M_1 = \begin{pmatrix} \sin \alpha \cos \psi - \cos \alpha \sin \delta \sin \psi & -\cos \alpha \cos \psi - \sin \alpha \sin \delta \sin \psi & \cos \delta \sin \psi \\ -\sin \alpha \sin \psi - \cos \alpha \sin \delta \cos \psi & \cos \alpha \sin \psi - \sin \alpha \sin \delta \cos \psi & \cos \delta \cos \psi \\ -\cos \alpha \cos \delta & -\sin \alpha \cos \delta & -\sin \delta \end{pmatrix} \quad (6.6)$$

$$M_2 = \begin{pmatrix} \sin \lambda \cos(\phi_r + \Omega_r t) & \sin \lambda \sin(\phi_r + \Omega_r t) & -\cos \lambda \\ -\sin(\phi_r + \Omega_r t) & \cos(\phi_r + \Omega_r t) & 0 \\ \cos \lambda \cos(\phi_r + \Omega_r t) & \cos \lambda \sin(\phi_r + \Omega_r t) & \sin \lambda \end{pmatrix} \quad (6.7)$$

$$M_3 = \begin{pmatrix} -\sin(\gamma + \zeta/2) & \cos(\gamma + \zeta/2) & 0 \\ -\cos(\gamma + \zeta/2) & -\sin(\gamma + \zeta/2) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (6.8)$$

In the above equations λ is the latitude of the detector's site, Ω_r is the rotational angular velocity of the Earth, ϕ_r is a deterministic phase which defines the position of the Earth in its diurnal motion at $t = 0$, such that $\phi_r + \Omega_r t$ coincides with the local sidereal time of the detector's site, γ is the angle measured counter-clockwise from East to the bisector of the interferometer arms, and ζ is the angle between the interferometer arms. These angles are shown in Figure 6-1 for reference. The vectors \hat{n}_1 and \hat{n}_2 in the detector's reference frame

are $\hat{n}_1 = (1, 0, 0)$ and $\hat{n}_2 = (\cos(\zeta), \sin(\zeta), 0)$.

By using Eqs. (6.1)-(6.8) and the relevant definitions we can algebraically determine the following expressions for $F_+(t)$ and $F_\times(t)$:

$$F_+(t) = \sin(\zeta)[a(t) \cos(2\psi) + b(t) \sin(2\psi)] \quad (6.9)$$

and

$$F_\times(t) = \sin(\zeta)[b(t) \cos(2\psi) - a(t) \sin(2\psi)] \quad (6.10)$$

where

$$\begin{aligned} a(t) = & \frac{1}{16} \sin(2\gamma)(3 - \cos(2\lambda))(3 - \cos(2[\alpha - \phi_r - \Omega_r t])) \\ & - \frac{1}{4} \cos(2\gamma) \sin(\lambda)(3 - \cos(2\delta)) \sin(2[\alpha - \phi_r - \Omega_r t]) \\ & + \frac{1}{4} \sin(2\gamma) \sin(2\lambda) \sin(2\delta) \cos(\alpha - \phi_r - \Omega_r t) \\ & - \frac{1}{2} \cos(2\gamma) \cos(\lambda) \sin(2\delta) \sin(\alpha - \phi_r - \Omega_r t) \\ & + \frac{3}{4} \sin(2\gamma) \cos^2(\lambda) \cos^2(\delta) \end{aligned} \quad (6.11)$$

and

$$\begin{aligned} b(t) = & \cos(2\gamma) \sin(\lambda) \sin(\delta) \cos(2[\alpha - \phi_r - \Omega_r t]) \\ & + \frac{1}{4} \sin(2\gamma)(3 - \cos(2\lambda)) \sin(\delta) \sin(2[\alpha - \phi_r - \Omega_r t]) \\ & + \cos(2\gamma) \cos(\lambda) \cos(\delta) \cos(\alpha - \phi_r - \Omega_r t) \\ & + \frac{1}{2} \sin(2\gamma) \sin(2\lambda) \cos(\delta) \sin(\alpha - \phi_r - \Omega_r t) . \end{aligned} \quad (6.12)$$

With the explicit equations for the $F_+(t)$ and $F_\times(t)$ in hand, it is possible to look at the response of the detectors to an average source. The averages of $F_+^2(t)$ and $F_\times^2(t)$ are both equal to $\frac{1}{5}$, when averaged over right ascension, declination, polarization, and integrating for sufficiently long periods (longer than a few days).

6.3 Signal form in terms of Doppler and amplitude parameters

We can write a general, elliptically polarized gravitational wave in the reference frame of the source as [47]

$$\begin{aligned} h_+(\tau) &= A_+ \cos(\Phi(\tau)) \\ h_\times(\tau) &= A_\times \sin(\Phi(\tau)) \end{aligned} \quad (6.13)$$

where A_+ and A_\times are amplitudes containing the detector's response function that depends on the wave's polarization and source inclination angle and τ is the time at which the wave left the source. If we assume a quasi-monochromatic signal with slowly-varying intrinsic frequency $f(\tau)$, the signal phase $\Phi(\tau)$ can be Taylor expanded as

$$\Phi(\tau) = \phi_0 + \phi(\tau), \quad \phi(\tau) = 2\pi \sum_{k=0}^s \frac{f^{(k)}(\tau_{ref})}{(k+1)!} \Delta\tau^{k+1} \quad (6.14)$$

where $\Delta\tau \equiv \tau - \tau_{ref}$ and τ_{ref} is the reference time at which the initial phase ϕ_0 and the $s+1$ spin parameters $f^{(k)} \equiv d^k f(\tau)$ are defined, i.e. the 1st spin down parameter, the 2nd spin down parameter, etc.

From an Earth based detector's frame the signal amplitude of an isolated neutron star is modulated by the time dependent antenna patterns $F_+(t)$ and $F_\times(t)$ and also Doppler-modulated by the relative motion of the detector with respect to the source. We can write a relation between the detector arrival time t of a wave-front that left the source at time $\tau(t)$. If we consider a neutron star at sky position \hat{n} , defined in equatorial coordinates as $\hat{n} = (\cos \delta \cos \alpha, \cos \delta \sin \alpha, \sin \delta)$ in terms of right ascension α and declination δ , and at a distance $d(t)$ then we can write

$$\tau(t; \hat{n}) = t + \Delta_{\odot}(t; \hat{n}) - \frac{d(t)}{c} \quad (6.15)$$

where

$$\Delta_{\odot}(t; \hat{n}) = \frac{\mathbf{r}(t) \cdot \hat{\mathbf{n}}}{c} + \Delta_{E\odot}(t; \hat{n}) + \Delta_{S.}(t; \hat{n}) \quad (6.16)$$

and where $\mathbf{r}(\mathbf{t})$ is the vector from the solar system barycenter to the detector on Earth. Here $\Delta_{S\odot}$ is the general-relativistic Shapiro delay caused by the gravitational field of the sun and $\Delta_{E\odot}$ is delay caused by the gravitational redshift and time dilation at Earth [49]. Both of these effects are much smaller than the $\frac{\mathbf{r}(\mathbf{t}) \cdot \hat{\mathbf{n}}}{c}$ Roemer-delay term, which is simply the difference in the propagation time of the wave to reach the solar system barycenter and the detector. [47]. In the case of searching for a signal, the time dependent distance $d(t)$ is typically treated as constant and neglected, which involves no loss of generality since any effects can be absorbed into the frequency evolution of the source. We can now write the measured strain signal $h(t)$ as

$$\begin{aligned} h(t; \mathbf{A}, \lambda) &= F_+(t; \hat{n}, \psi) A_+ \cos(\phi_0 + \phi(t; \lambda)) \\ &\quad + F_\times(t; \hat{n}, \psi) A_\times \sin(\phi_0 + \phi(t; \lambda)) \end{aligned} \quad (6.17)$$

where we define the set of four ‘‘amplitude parameters’’ as $\mathbf{A} \equiv \{\mathbf{A}_+, \mathbf{A}_\times, \psi, \phi_0\}$ and we define a set of ‘‘Doppler parameters’’ $\lambda \equiv \{\mathbf{f}^{(\mathbf{k})}, \hat{\mathbf{n}}\}$.

Using 6.9 and 6.10 we can separate $h(t)$ ’s dependence on the amplitude and Doppler parameters and write

$$h(t; \mathbf{A}, \lambda) = \sum_{\mu=1}^4 \mathbf{H}^\mu \mathbf{h}_\mu(t; \lambda) \quad (6.18)$$

in terms of four basis waveforms

$$\begin{aligned} h_1(t; \lambda) &= a(t; \hat{n}) \cos[\phi(t; \lambda)] \\ h_2(t; \lambda) &= b(t; \hat{n}) \cos[\phi(t; \lambda)] \\ h_3(t; \lambda) &= a(t; \hat{n}) \sin[\phi(t; \lambda)] \\ h_4(t; \lambda) &= b(t; \hat{n}) \sin[\phi(t; \lambda)] \end{aligned} \quad (6.19)$$

and the amplitude vector H^μ is defined as

$$\begin{aligned}
H^1 &= A_+ \cos(\phi_0) \cos(2\psi) - A_\times \sin(\phi_0) \sin(2\psi) \\
H^2 &= A_+ \cos(\phi_0) \sin(2\psi) + A_\times \sin(\phi_0) \cos(2\psi) \\
H^3 &= -A_+ \sin(\phi_0) \cos(2\psi) - A_\times \cos(\phi_0) \sin(2\psi) \\
H^4 &= -A_+ \sin(\phi_0) \sin(2\psi) + A_\times \cos(\phi_0) \sin(2\psi) .
\end{aligned} \tag{6.20}$$

In the case of a non-precessing triaxial neutron star, as described in Section 1.2, the amplitudes A_+ and A_\times can be written as

$$A_+ = \frac{1}{2}h_0(1 + \cos^2(i)), \quad A_\times = h_0 \cos(i), \tag{6.21}$$

where h_0 is the overall amplitude. From Section 1.2, it is $64\pi^2 f^2 \epsilon I_{33}/R$. In more traditional units and using $f_{GW} = 2f$ (i.e. looking at a frequency twice the rotation rate) we can rewrite this as

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{\epsilon I_{33} f_{GW}^2}{R}. \tag{6.22}$$

It should be noted that for other types of continuous wave generation, different relations hold between the rotation frequency and the gravitational wave frequency. For example, free precession emits near the rotation frequency and twice the rotation frequency. However, the above representations of A_+ and A_\times can be used without loss of generality since the mapping between $\{h_0, \cos(i)\}$ and $\{A_+, A_\times\}$ is one to one.

6.4 The F statistic: a coherent search method

The use of the F statistic as a search method was developed and first presented by Jaranowski, Krolak, and Schutz in [27]. The search code was developed by the LIGO Continuous Wave pulsar group [46]. I will present a short summary here of the method.

One method to extract that signal is the method of maximum likelihood detection. A key part of this method is the likelihood function Λ , a kind of probability distribution that expresses the likelihood that our data is described by a particular model or set of parameters. By finding the signal parameters which maximize Λ , by taking the derivative

of Λ with respect to those parameters and setting it equal to zero to find the extrema, we effectively find the best match in the data to our possible signal. The signal parameters which maximize Λ are called the maximum likelihood estimators. The maximum of Λ determines the probability of detection of the signal. If the maximum of Λ exceeds a certain threshold determined by the false alarm probability we can afford, we say the signal is detected.

We begin by making the assumption that the noise in the detector is an additive, stationary, Gaussian and zero-mean process. Then the data x we get from the detector when a signal h is present is

$$x(t) = n(t) + h(t) \quad (6.23)$$

where $n(t)$ is our noise. The log likelihood function has the form

$$\log \Lambda = (x|h) - \frac{1}{2}(h|h) \quad (6.24)$$

where the scalar product $(\bullet|\bullet)$ is defined by

$$(x|y) \equiv 4\Re \int_0^\infty \frac{\tilde{x}(f)\tilde{y}^*(f)df}{S_h(f)} \quad (6.25)$$

where \sim denotes a Fourier transform, $*$ denotes complex conjugation, and S_h is the one sided spectral density of the detector noise.

The gravitational wave signal consists of a narrowband component around the frequency f , twice the rotational frequency. We can assume that over the bandwidth of the signal $S_h(f)$ is nearly constant, stationary, and equal to $S_h(f_0)$ where f_0 is the frequency of the signal h at $t = 0$. With this assumption the detector noise term can be factored out of the inner product integral. By using the definition of the Fourier transformation, we can change the frequency domain intergral of Eq. (6.25) into a time domain integral, a more convenient form for periodic source analysis. Thus the above scalar product can be approximated by

$$(x|h_1) \cong \frac{2}{S_h(f_0)} \int_{-T_0/2}^{T_0/2} x(t)h(t)dt \quad (6.26)$$

where T_0 is the observation time starting at $t = -T_0/2$. We introduce a new scalar product,

$$(x||y) \equiv \frac{2}{T_0} \int_{-T_0/2}^{T_0/2} x(t)y(t)dt , \quad (6.27)$$

without the detector noise term and normalized by the observation time. We instead include these terms explicitly in the log likelihood function. The log likelihood function for this signal with this new scalar product is approximately given by

$$\log \Lambda \cong \frac{T_0}{S_h(f_0)} \left[(x||h) - \frac{1}{2}(h||h) \right] . \quad (6.28)$$

The maximum likelihood estimators can be found by maximizing the following normalized log likelihood function

$$\log \Lambda'_1 = (x||h) - \frac{1}{2}(h||h) . \quad (6.29)$$

The normalized log likelihood function has had its explicit dependence on the spectral density of the noise in the detector removed.

The signal h depends linearly on the four amplitudes H^μ . These in turn depend on the parameters h_0, θ, ψ, i , and ϕ_0 and are independent. In a search where not all of these parameters are known, we must find the maximum of $\log \Lambda$ as a function of the unknown parameters.

If we insert Eq. (6.18) into Eq. (6.29) we can write

$$\log \Lambda'(x; \mathbf{A}, \lambda) = \mathbf{A}^\mu \mathbf{x}_\mu - \frac{1}{2} \mathbf{A}^\mu \mathbf{A}^\nu \mathbf{M}_{\mu\nu} \quad (6.30)$$

where

$$x_\mu \equiv (x||h_{/mu}) \quad (6.31)$$

and

$$M_{\mu\nu}(\lambda) \equiv (\mathbf{h}_\mu||\mathbf{h}_\nu) . \quad (6.32)$$

With these we can now maximize $\log \Lambda$ over A^μ to obtain the maximum likelihood estimators A_{ML}^μ from the data $x(t)$ by

$$\frac{\partial \log \Lambda'}{\partial A^\mu} = 0 \implies A_{ML}^\mu = M^{\mu\nu} x_\nu \quad (6.33)$$

where $M^{\mu\alpha}M_{\alpha\nu} = \delta_\nu^\mu$. Placing these maximum likelihood estimators back into Eq. (6.30) we obtain a new partially maximized detection statistic

$$2F(x; \lambda) = \mathbf{x}_\mu \mathbf{M}^{\mu\nu} \mathbf{x}_\nu \quad (6.34)$$

which is called the F-statistic. Since we have maximized over the amplitude parameters, the F-statistic only depends on the Doppler parameters λ . In the presence of a signal, the expectation value of the F-statistic with perfectly matched Doppler parameters, $\lambda = \lambda_{\text{sig}}$, is given by $4 + \rho^2$, where ρ is the optimal signal to noise ratio. It is shown in [27] that

$$\rho = \sqrt{(h|h)} \quad (6.35)$$

and that this scales as $\frac{h_0}{\sqrt{S_h}} \sqrt{T_{\text{obs}}}$ for a single detector. One can show, again in [27] that $2F$ is a random variable with a χ^2 distribution with 4 degrees of freedom and a non-centrality parameter ρ^2 . The probability density function of this is

$$p_1(F) = \frac{1}{2} e^{-(F+\rho^2)/2} \sqrt{\frac{F}{\rho^2}} I_1(\sqrt{F\rho^2}) \quad (6.36)$$

where I_1 is the modified Bessel function of the first kind of order one. In the absence of signal, when $\rho = 0$, this reduces to a central χ^2 distribution

$$p_0(2F) = \frac{F}{2} e^{-F} . \quad (6.37)$$

With these probability distributions we can compute the false alarm and false dismissal probabilities for a threshold $2F_c$ and a signal to noise ρ . The false alarm probability is simply the integration 6.37 from the threshold $2F_c$ to infinity. In a Frequentist framework it represents the limiting fraction of $2F$ values which exceed the threshold, given an infinite number of identical trials on Gaussian, stationary, white noise. We can perform this integration to get

$$P_{FA}(2F_c) = \int_{2F_c}^{\infty} p_0(2F) d(2F) = \frac{(1 + F_c)}{2} e^{-F_c} . \quad (6.38)$$

For example, if we choose a 1% false alarm rate for a single template search, the necessary cutoff is $2F_c \approx 13.3$ [46]. The false dismissal rate is the integration of 6.36 from negative

infinity to the threshold $2F_c$ for the given signal to noise ρ . In general this needs to be integrated numerically. As an example, if we use the cutoff derived from a 1% false alarm for a single template search and require a false dismissal rate of 10%, we find a $\rho \approx 4.5$.

6.5 Coherent vs incoherent templates and unknown parameters

When determining what type of search to perform, it is important to consider how the computation costs scale and the relative gains in sensitivity for that cost. The computation cost is most easily expressed in the number of templates needed to cover a given parameter space for a given type of search. Searching each template, which defines the sky position and frequency evolution of the source to be looked at, takes the same amount of time, and thus the computation scales directly with the number of templates.

We also need to consider the sensitivity scaling of the search, which is a function of the coherence time, T_{coh} , which is effectively the length of time over which we integrate, keeping track of all the parameters of the signal including both phase and power. In the case of a fully coherent search, this T_{coh} is equal to the total amount of data analyzed, T_{total} . In the case of an incoherent search, where short segments of data are converted via FFT and then combined ignoring phase, T_{coh} is the length of the individual segments. In that case, $T_{total} = NT_{coh}$, where N is the number of segments.

As we integrate, the mean power in the noise goes as $1/T_{coh}$. The signal to noise in terms of power therefore goes as T_{coh} . Note that the strain we want to detect, h_0 , is an amplitude rather than a power. The detectable strain is therefore proportional to $1/\sqrt{T_{coh}}$, remembering that a smaller signal is easier to detect with higher signal to noise. So for a fully coherent search which simply uses $T_{coh} = T_{total}$ the detectable h_0 goes as $1/\sqrt{T_{total}}$. In the case of an incoherent search of N segments of fixed length T_{coh} , the individual segments have a mean noise power of $1/T_{coh}$, but when added incoherently the power of the noise scales as $1/(\sqrt{N}T_{coh})$. The signal to noise of the incoherent search therefore goes as $\sqrt{N}T_{coh} = \sqrt{T_{total}T_{coh}}$. The detectable strain then scales as $1/(T_{coh}T_{total})^{1/4}$. The coherent search is therefore more sensitive than the incoherent search. However, depending on the size of the parameter space, the coherent search may require orders of magnitude more computation time.

In order to calculate the number of templates, we begin with defining the resolution of the search in terms of frequency, which sets the scale of how far templates can be in the parameter space before losing information. Then we must consider that for a sufficiently large \dot{f} , or other higher order derivative, the frequency of the source will move by more than the resolution over the length of a search. That makes it necessary to cover the 1st spin down space with templates. Similarly, how well the source location is known needs to be compared to the resolution in sky position.

The resolution in frequency for either the coherent or the incoherent search, Δf , which determines the maximum spacing in frequency of the templates is $\frac{1}{T_{coh}}$. With this spacing, a signal exactly matched to one template would show up with half the power in an adjacently placed template. Thus the number of frequency points needed to reasonably cover a search from a minimum frequency f_{min} and a maximum frequency f_{max} is just $\frac{f_{max}-f_{min}}{\Delta f}$. Most blind continuous wave LIGO searches use a minimum frequency of 50 Hz, motivated by the performance of the detectors at low frequency, and a maximum frequency of 1 kHz, which is motivated by the fastest pulsars seen.

Next, we need to consider how a signal might evolve over the course of the measurement. We start from Eq. (6.14) and derive the Taylor expansion of the frequency as a function of time at the solar system barycenter (SSB) by writing

$$\frac{d\Phi}{dT} = f = f_0 + \frac{\dot{f}}{1!}(T - T_0) + \frac{\ddot{f}}{2!}(T - T_0)^2 + \dots \quad (6.39)$$

where Φ is the phase of the signal, T is the time at the SSB, and T_0 is the start time at the SSB. We can translate this to the detector frame of reference here on Earth by using Eqs. (6.15) and (6.16) to write

$$\frac{d\Phi}{dt} = \frac{d\Phi}{dT} \frac{dT}{dt} = f_{\oplus} = \left(1 + \frac{\vec{v} \cdot \hat{n}}{c}\right) \left(f_0 + \frac{\dot{f}}{1!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0\right) + \frac{\ddot{f}}{2!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0\right)^2 + \dots\right) \quad (6.40)$$

where \vec{v} is the earth's velocity relative to the SSB, \hat{n} is the unit vector pointing from the Earth to the source, \vec{r} is the vector from the Earth to the SSB, and t is the time at the detector on Earth. The above treatment neglects the daily rotation that the detector will under go on the surface of the earth, but the exact same methodology can be applied to

include those effects. This also neglects the small general relativity corrections.

So with this expansion we now consider grid points with differing \dot{f} which are necessary to map the spin down parameter space. It is important to note that we consider each parameter of the parameter space separately. We hold all the other parameters fixed while calculating the spacing of that single parameter. We begin by considering the m^{th} grid spacing in \dot{f} space while holding the frequency, f_0 , the higher order derivatives such as \ddot{f} , and the sky position all constant. This can be written as

$$f(m) = \left(1 + \frac{\vec{v} \cdot \hat{n}}{c}\right) \left(f_0 + \frac{\dot{f} + m \Delta \dot{f}}{1!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0 \right) + \frac{\ddot{f}}{2!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0 \right)^2 + \dots \right) \quad (6.41)$$

where \dot{f} is the first templates frequency time derivative and $\Delta \dot{f}$ is the spacing between templates in \dot{f} space. The difference between the m^{th} template and $m^{\text{th}} + 1$ cannot be larger than the resolution Δf . If we choose a spacing such that they are equal then we have

$$\Delta f = f_1(m+1) - f_1(m) = \left(1 + \frac{\vec{v} \cdot \hat{n}}{c}\right) \Delta \dot{f} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - t_0 - \frac{\vec{r}_0 \cdot \hat{n}}{c} \right) \quad (6.42)$$

The maximum value for $\frac{\vec{r} \cdot \hat{n}}{c}$ is ~ 480 seconds and thus the difference could be at most twice this at ~ 960 seconds. This compared to the days or months of most LIGO observation runs, defined as $T_{run} = t - t_0$, will be quite small and can be neglected. Similarly, the Earth's orbital velocity relative to the SSB is small when compared to the speed of light, so $1 + \frac{\vec{v} \cdot \hat{n}}{c}$ can be approximated as 1. Using $\Delta f = \frac{1}{T_{coh}}$ we get the maximum spacing in \dot{f} space. Simply

$$\Delta \dot{f} \cong \frac{\Delta f}{T_{run}} = \frac{1}{T_{coh} T_{run}} \quad (6.43)$$

However the boundaries to search within in \dot{f} space are not quite as clear as in the f space case. The maximum \dot{f} we might be expected to see from a star spinning down only to gravitational radiation is discussed in 5. However, for the general case of \dot{f} proportional to n^{th} power of frequency is $\left(\frac{1}{n-1}\right) \left| \frac{\Omega(t)}{\dot{\Omega}(t)} \right|$. Thus a reasonable upper estimate is to drop the prefix $\frac{1}{n-1}$ and set \dot{f}_{max} equal to $\frac{f_{max}}{t_{age}}$ where t_{age} is spin down age below which we will not look for sources. We then a number of 1st spin down points equal to $\frac{f_{max}}{t_{age}} T_{coh} T_{run}$.

For sufficiently long runs or sufficiently young targets a 2nd spin down parameter will be necessary (and in principle even higher orders but for the searches considered here the 2nd spin down is the highest order necessary). Following a similar procedure as for the 1st spin down case, it can be shown that the spacing is $\Delta \ddot{f} = \frac{2}{T_{coh} T_{run}^2}$ and an upper limit on the necessary number of points is $\frac{f_{max}}{t_{age}^2} \frac{T_{coh} T_{run}^2}{2}$.

Lastly we need to consider the resolution in terms to of points on the sky. We once again go back to Eq. (6.40), noting we can represent $\vec{v} \cdot \hat{n}$ with $|v| \cos \theta$ where θ is simply the angle between the Earth's velocity and the vector pointing from Earth to the source. Let us now consider the m^{th} template, but with the spacing in the θ rather than \dot{f} , while holding the frequency parameters constant. We can write

$$f(m) = \left(1 + \frac{|v| \cos(\theta + m \Delta \theta)}{c}\right) \left(f_0 + \frac{\dot{f}}{1!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0\right) + \frac{\ddot{f}}{2!} \left(t + \frac{\vec{r} \cdot \hat{n}}{c} - T_0\right)^2 + \dots\right) \quad (6.44)$$

where $\Delta \theta$ is our spacing in angle. We again set equal the different between the m^{th} and $m^{th} + 1$ points equal to Δf , neglect $\frac{\vec{r} \cdot \hat{n}}{c}$, but also assume $\sin(\Delta \theta) \cong \Delta \theta$ to get

$$\Delta f = \frac{|v| \sin(\theta) \Delta \theta}{c} \left(f_0 + \dot{f} T_{run} + \dots\right). \quad (6.45)$$

Solving for $\Delta \theta$ and using f_{max} for the largest contributing term of $f_0 + \dot{f} T_{run} + \dots$ we find

$$\Delta \theta_{min} = \frac{c}{v_{max}} \frac{1}{f_{max} T_{coh}}. \quad (6.46)$$

We can now look at the total number of templates necessary for coherent and incoherent searches. In the case of incoherent searches, since the T_{coh} is fixed, the templates only scale as T_{run}^3 , one factor of T_{run} from the 1st spin down and two factors of T_{run} from the 2nd spin down. However in the case of coherent searches, it scales as T_{run}^8 since $T_{coh} \cong T_{run}$. A single template using either an incoherent or coherent search code takes approximately $1 \mu s$ per 1800 second SFT when run on a single node of a cluster, such as those at the Hanford LIGO site or the Caltech cluster. Given that the time to run a single template is proportional to the amount of data in both coherent and incoherent searches, the total computational cost scales as T_{run}^4 for incoherent searches while coherent searches scale as T_{run}^9 . In directed

searches, if the position of the target is known to better than $\Delta\theta$ then the coherent search scales as T_{run}^7 .

If the parameter space can be kept sufficiently small so that the coherent search is computationally possible, then because its sensitivity scales as the square root of the amount of data available, compared to the fourth root of the amount of data available for incoherent searches, it is the better choice. For the searches presented here, it was possible to keep the computation time to days or weeks on a cluster, which is what prompted the use of the coherent F statistic code over an incoherent method..

6.6 Exact grid spacing

We have just discussed the scaling of the number of templates with the search parameter space, but we began with the statement that δF is simply $1/T_{coh}$. We need to consider how to quantify how good that spacing is, and whether we need to increase or decrease it from this nominal value. The overall scaling will remain the same, but it affects the exact number of templates we choose.

For the case of the F-statistic coherent search method, we begin to consider this by defining the mismatch μ , the difference between the F value exactly at the signal, which we will call F_{match} and its value at the different position where our search grid told us to place a template, which we will call F_{search} . The small offset in template space we will label dx^μ such that $x^0 = f, x^1 = \alpha, x^2 = \delta, x^3 = \dot{f}, x^4 = \ddot{f}$. We can then write

$$m \equiv \frac{F_{match} - F_{search}}{F_{match}} \quad (6.47)$$

In the local limit as the grid spacings dx^μ goes to 0 and neglecting effects due to amplitude parameters, the mismatch in Eq. (6.47) can approximately be written as

$$m = g_{\mu\nu} dx^\mu dx^\nu \quad (6.48)$$

The metric components were worked out by Ian Jones, Ben Owen, and David Whitbeck, and later tested by Reinhard Prix [45]. To do so they used an epicyclic (or Ptolemaic) motion of the detectors to account the effects of the Earth's motion on the incoming signal. This was done so that an analytic expression could be used to derive the metric components

explicitly and reduce the amount of numerical computation during searches. We can write the metric components of interest (f, \dot{f}, \ddot{f}) using the shorthand $f_k = d^k f / dt^k$ as

$$\begin{aligned} g_{ff} &= \frac{(\pi T)^2}{3} \\ g_{ff_k} &= \frac{(2\pi)^2 T^{k+2} f}{(k+1)(k+3)} \end{aligned} \quad (6.49)$$

$$g_{f_j f_k} = \frac{(2\pi f T)^2 T^{j+k}}{(j+1)(k+1)(j+k+3)} . \quad (6.50)$$

We can create a new metric with the frequency projected out, using the analytic formula

$$\gamma_{\mu\nu} = \left(g_{\mu\nu} - \frac{g_{f\mu} g_{f\nu}}{g_{ff}} \right) . \quad (6.51)$$

Explicitly substituting in 6.49 into 6.51 we can solve for the diagonal metric components, which we then can use to define our spacings in the individual subspaces of the overall parameter space. We get

$$\gamma_{\mu\nu} = \frac{(2\pi)^2 T^{j+k+2}}{(j+2)(k+2)(3+j+k)} \quad (6.52)$$

and the spacings are just

$$\begin{aligned} df &= \sqrt{\frac{m}{g_{ff}}} \\ d\dot{f} &= \sqrt{\frac{m}{\gamma_{f\dot{f}}}} \end{aligned} \quad (6.53)$$

$$d\ddot{f} = \sqrt{\frac{m}{\gamma_{f\ddot{f}}}} . \quad (6.54)$$

There is a slight discrepancy in the definitions between the code and this spacing. For example the code uses $f(t) = f + \dot{f}t + \ddot{f}t^2/2 + \ddot{\ddot{f}}t^3/6...$ when calculating the final frequency to look at, while the metric does not take into account the additional factor of 2 or 6 for the higher order spin downs. Effectively a prefix of $n!$ should be added to the spacings where n

is order of the derivative associated with the spacing.

6.7 Interferometer sensitivities

The overall sensitivity we can detect is highly dependent on the sensitivity of the interferometers used since their noise hides the actual gravitational wave signals. One method of estimating $S_h(f)$, the single-sided power spectral density, is to take the SFTs, which are simply 1800 second long Fourier transforms, averaging the absolute magnitude at each frequency and normalizing by the 1800 seconds. Thus

$$S_h(f) = \frac{1}{N} \sum_{i=1}^N \frac{|SFT_N(f)|^2}{T_{SFT}} \quad (6.55)$$

where $|SFT_N(f)|$ is the absolute value of the N^{th} SFT at frequency f and T_{SFT} is the length of time used to make each SFT. Typically, the noise is reported as the $\sqrt{S_h}$.

6.8 Estimating limits of detection

We now have developed sufficient tools in the chapter to be able to make estimates for the smallest signal a single template or a multiple template search could find.

We begin with the optimal signal to noise ratio ρ . Using Eq. (6.35) and the same reasoning leading up to Eq. (6.26) we can write

$$\rho = \sqrt{\langle h|h \rangle} \cong \left[\frac{2}{S_h(f_0)} \int_{-T_0/2}^{T_0/2} [h(t)]^2 dt \right]^{1/2} \quad (6.56)$$

where f_0 is our gravitational wave frequency, typically twice the rotation frequency. We then use Eqs. (6.17) and (6.21) to re-write this while dropping the cos and sin terms which oscillate about the frequency of the gravitational wave so that

$$\rho^2 \cong \left[\frac{1}{4} (1 + \cos^2(i))^2 \int_{-T_0/2}^{T_0/2} F_+^2 dt + \cos^2(i) \int_{-T_0/2}^{T_0/2} F_\times^2 \right] \frac{h_0^2}{S_h(f_0)}. \quad (6.57)$$

For an all sky survey, to find the average optimal signal to noise ratio, we could average this equation over all right ascensions, declination, polarizations and inclination angles which would give

$$\langle \rho^2 \rangle_{\alpha,\delta,\psi,i} \cong \frac{4}{25} \frac{h_0^2 T_0}{S_h(f_0)}. \quad (6.58)$$

Lastly, we can then solve this equation for h_0 in terms of the observation time, the noise in the detector, and the required signal to noise ratio for 1% false alarm and 10% false dismissal rates to get

$$h_0 \cong \sqrt{\frac{25\rho^2}{4} \frac{S_h(f_0)}{T_0}} \cong 11.4 \sqrt{\frac{S_h(f_0)}{T_0}} \quad (6.59)$$

where we have used the $\rho \cong 4.5$ for a single template search from section 6.4.

The more interesting case is that for the targeted search, where we average over all right ascensions, polarizations and inclinations while fixing the declination of our source, for either the Crab pulsar or RX J1856.5-3754. In this case, results depend on the detector's location, whether Hanford, Washington or Livingston, Louisiana.

For the Hanford detectors and the Crab pulsar we obtain

$$\langle \rho^2 \rangle_{\alpha,\psi,i} \cong 0.1418 \frac{h_0^2 T_0}{S_h(f_0)} \quad (6.60)$$

while for the Livingston detector we get

$$\langle \rho^2 \rangle_{\alpha,\psi,i} \cong 0.1670 \frac{h_0^2 T_0}{S_h(f_0)}. \quad (6.61)$$

In the case of RX J1856.5-3754 and the Hanford detectors we have

$$\langle \rho^2 \rangle_{\alpha,\psi,i} \cong 0.2127 \frac{h_0^2 T_0}{S_h(f_0)} \quad (6.62)$$

and for the Livingston detector we have

$$\langle \rho^2 \rangle_{\alpha,\psi,i} \cong 0.1951 \frac{h_0^2 T_0}{S_h(f_0)}. \quad (6.63)$$

Thus, for a single template search, the final detectable signal strength for the Crab would be roughly $h_0 \cong 11.6 \sqrt{\frac{S_h(f_0)}{T_0}}$ while for RX J1856.5-3754 it would be about $h_0 \cong 10.1 \sqrt{\frac{S_h(f_0)}{T_0}}$.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 7

Searches and results

This chapter describes the use of directed search of RX J1856.5-3754 as a test to compare the versions 1 and 2 of the ComputeFStatistic code. Version 2 of the ComputeFStatistic code is then used for a search of the Crab pulsar over a physically motivated parameter space. *These results are presently under review, and should, therefore, be considered preliminary.*

7.1 RX J1856.5-3754 directed search

The ComputeFStatistic code, whose basic principles were described in Chapter 7.2.1, has been used in several previous searches [2]. At the time a search for RX J1856.5-3754 was under consideration, a newer multi-interferometer search code was becoming available, which we call the version 2 code. The fundamental difference between these codes was the version 2 code took the data from all three interferometers simultaneously, to calculate a single $2F$ value, while the original code could only deal with a single interferometer at a time. As validation test of this newer code, we ran both the original search code and the version 2 code and compared their results.

7.1.1 Parameter space chosen

The search parameter space was chosen to extend from 50 Hz to 250 Hz, using 13 days of data (624 SFTs). This was based on the estimated computation time of ~ 2 weeks on the LIGO Hanford computer cluster and a total of 7×10^{10} templates to be searched. It included the frequency band of highest sensitivity of the interferometer, around 130 Hz to 150 Hz. In that band it was estimated the search would be able to place better direct upper limits on

the gravitational radiation than the theoretical upper limit calculations from Section 5.5, assuming RX J1856.5-3754 was 10^5 years old and spinning down purely to gravitational radiation. In addition, the 1st spin down (dF/dt) parameter space was selected so that it covered ages of 10^5 years and older, by using Eq. (5.13). This meant the 1st spin down parameter space covered increased linearly with the frequency searched.

7.1.2 Search results

For a search using 7.1×10^{10} templates, and assuming Gaussian white noise, we would expect a 1% false alarm $2F$ cutoff value of 66. The overall search would have an expected largest $2F$ value of 58. The calculation methods are described in Section 7.2.1.

The largest $2F$ values for every 0.25 Hz band for H1 and L1 for the version 1 search code are shown in Figures 7-1 and 7-2, and for the combined search with the version 2 code in Figure 7-3. In each case, the floor of $2F$ of approximately 40 is expected based on the number of templates in each 0.25 Hz search band. That estimate is based on the probability density function for the F statistic when only Gaussian noise is present. The slight trend upwards is due to the fact that the number of templates increases at higher frequencies, so as to cover a sufficient 1st spin down space.

By itself H1 has three large templates above what might be expected from Gaussian noise at the 1% false alarm level. These features are associated with known lines present in the instrument itself during the S4 run. They correspond to 60 Hz line noise, a 16 Hz comb (which is a collection of narrow lines at all multiples of 16 Hz caused by noisy electronics) with a particularly large spike at 112 Hz, and the leakage from one of the hardware injected test pulsars near 194 Hz. It is not clear why the 7th harmonic of the 16 Hz comb was so large relative to the background noise, when compared to the other harmonics. L1 individually has a very complex “forest” of events, due to known instrument lines. In this case, the easiest way to reduce the number of lines was the coincident step, which is described in the next section.




Figure 7-1: Plot of the largest $2F$ values every 0.25 Hz for the H1 search with the version 1 code. The expected value for any particular 0.25 Hz band is around 40, with slightly higher values expected at higher frequencies because there are more 1st spin down values covered and thus more templates searched per 0.25 Hz band. The line at 60 Hz is due to 60 Hz power line noise, the line at 112 Hz is due to a 16 Hz comb, and the line at 194 Hz is due to a hardware injected pulsar.




Figure 7-2: Plot of the largest $2F$ values every 0.25 Hz for the L1 search with the version 1 code. The forest of lines are due to many narrow and strong instrument lines, such as a 37 Hz comb, 60 Hz power lines, and other combs. As in the H1 case, the expected value is roughly 40, increasing at higher frequencies as the number of templates increases per 0.25 Hz band.

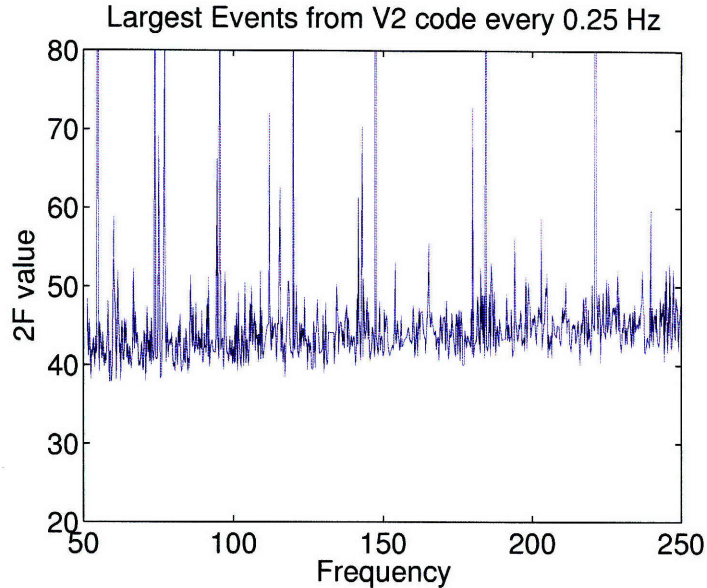


Figure 7-3: Plot of the largest $2F$ values every 0.25 Hz for the combined H1 and L1 search with the version 2 code. The forest of lines we see is mostly due to the instrumental lines present in L1. The expected value is roughly 40, increasing at higher frequencies due to additional templates searched 0.25 Hz band.

7.1.3 Version 1 coincidence step

To improve the upper limit results of two individual interferometer searches, one can set a coincidence window in frequency space and find the largest coincident events. This step resulted in a 5–10% increase in sensitivity with a coincidence window of 10^{-5} Hz, approximately 10 bins to either side in frequency space. The mean $2F$ value for the largest H1 templates every 1 Hz before requiring coincidence was 43.6, but dropped to 39.4 for events in coincidence with L1. Similarly, the mean $2F$ value for the largest L1 templates every 1 Hz was 62.6 before the coincidence step and 56.5 after. Figure 7-4 show the largest $2F$ templates remaining after the coincidence step. The coincident largest templates still have a few remaining large events. Near 60 Hz is the usual 60 Hz line noise. The lines at 74 and 111 Hz correspond to a 37 Hz comb present in the L1 data. The 194 Hz line was the injected test pulsar leakage.




Figure 7-4: Plot of the largest $2F$ values surviving the coincidence requirement, every 0.25 Hz. There are several lines surviving the coincidence step such as the 60 Hz power line which is present in both interferometers, lines corresponding to a 16 comb, and also with the hardware pulsar injection at 194 Hz.

7.1.4 Upper limit injections

Both versions of the code used Monte Carlo injections to produce upper limits in 1 Hz wide search bands. These injections are done by choosing an h_0 injection strength, randomizing the signal parameters, and injecting a fake pulsar signal into the data. This is then searched for with the same search method as the full search, but restricted to a small volume of parameter space surrounding the injection. If the resulting search produces a value above the cutoff, it is considered found. After enough injections (approximately 4,000) are performed, the confidence of the upper limit can be calculated from the total number of injections found divided by the total number of injections. The injection sets are repeated with different h_0 values until the h_0 value which produced a confidence of 0.95 was found in each 1 Hz band.

In the version 1 case, the cutoff was determined by the sum of the log of the largest coincident $2F$ values in the 1 Hz band. This was compared to the sum of the log of the $2F$ values from an identical injection into both H1 and L1's data. In the version 2 case, the cutoff was simply the largest $2F$ value found within the 1 Hz band.

7.1.5 Final comparison

The version 1 and 2 upper limits for each 1 Hz sub-band are shown in Figure 7-5. For reference, the spectrum of the H1 and L1 interferometers at a time of good sensitivity are also shown in Figure 7-5, scaled by a factor of 100. For a single interferometer the expected upper limits would be $37\sqrt{S_h/T_{obs}}$. This was calculated by taking into account the declination of RX J1856.5-3754, the detector response functions, the 1% false alarm cutoff and a 95% confidence for the upper limits. The method of this calculation is detailed in Section 7.2.3.

The ratio of the version 1 upper limits to the version 2 upper limits are shown in Figure 7-6. An improvement on the order of $\sqrt{2}$ is expected from the version 1 to the version 2 code since it uses both sets of data to compute the final $2F$ value, effectively integrating the noise floor for twice as long. The exceptions where the version 1 code does better occur in the presence of extremely strong instrument lines present in only one of the interferometers. For example, near 147 Hz and 221 Hz the version 2 code produces unusually high upper limits. This is due to a comb of narrow lines at multiples of 37 Hz present in L1. Near these two points, the L1 data dominated the $2F$ calculation, and thus resulted in unusually high cutoff values. The version 1 coincidence search did not see these as H1 did not have any large coincident templates. These narrow lines in L1 were eventually removed when an oscillator was replaced on March 10, 2005.

What this tells us is that when using the version 2 code, we need to check for lines in individual interferometers to ensure we are not seeing very loud events which are in reality restricted to only one interferometer. With that caveat however, the version 2 code performs better than the coincidence method with the version 1 code. These lessons were ported over to the next targeted search performed, that of the Crab pulsar.




Figure 7-5: Plot of the upper limits for the version 1 and version 2 searches. Plotted for reference are the strain of the H1 and L1 interferometers scaled down by a factor of 100. We find that a majority of the very large upper limit values correspond to line features seen in the interferometers. The unusually high upper limits at 148 Hz and 221 are due to the extreme $2F$ values found in the search for those bands, which in turn translates into a higher upper limit. These large values were due to a 37 Hz comb of lines in L1.




Figure 7-6: Plot of the ratio of the version 1 upper limits to the version 2 upper limits. The few exceptional points where the version 2 code does worse than the version 1 code were in locations where a very strong and narrow line feature determined the largest $2F$ value found. This in turn resulted in larger upper limits, since the largest $2F$ value determines how large the injections need to be to be counted for upper limits. The fact that these particular lines were not present in H1 meant that the coincidence method was better at those particular points. In cases where large features are present in both interferometers, such as near multiples of 60 Hz, the version 2 code does significantly better.

7.2 Crab wide parameter search

A time domain search has already been carried out using LIGO data targeted on the exact radio ephemeris available for the Crab pulsar [1]. However, as noted in chapter 5, there are several physical mechanisms by which the electromagnetic pulse frequency could differ from the gravitational wave frequency. At a frequency difference of 3×10^{-3} Hz between the electromagnetic pulse and gravitational wave frequencies, this search would be roughly 10^5 times more sensitive than the exact single template search, simply because of the resolution of the search for 9 months of data is about 4×10^{-8} Hz. Thus, the wide parameter search presented here on the Crab pulsar is complimentary to that search effort.

7.2.1 Chosen search parameter space and results

Chapters 5 and 7.2.1 describe the competing constraints in determining the final parameter space to search, namely the computational cost of the search, the expected sensitivity, and physical motivations.

If we consider the free precession models described in Eqs. (5.16) through (5.20) we find a maximum difference between $2f_{EM}$ and f_{GW} of order $2\epsilon f_{EM}$. For ϵ we can consider the discussion from Section 5.2 and take the largest reasonable values of order 10^{-5} for hybrid stars. These combined would give a band to search of order 6×10^{-4} given the Crab's radio pulse frequency of 29.777 Hz. The two component model using Eq. (5.27) and the Crab's parameters also gave an estimate on the order of 6×10^{-3} . These both have an overall form of

$$f_{GW} = 2f_{EM}(1 + \delta) \quad (7.1)$$

where δ is corresponding to the small variation these models induce in the gravitational wave frequency, and is of the order 10^{-4} for these models. The associated search space is $\Delta f_{GW} = \delta 2f_{EM}$. One way to estimate the variation in 1^{st} and 2^{nd} spin downs is to realize that

$$\dot{f}_{GW} = 2(\dot{F}_{EM}(1 + \delta) + f_{EM}\dot{\delta}). \quad (7.2)$$

If we assume that f_{EM}/\dot{f}_{EM} is of the order of $\delta/\dot{\delta}$ or smaller, then we are left with a largest

1st spin down search band of

$$\Delta \dot{f}_{GW} \sim 3 \dot{f}_{EM} \delta. \quad (7.3)$$

The above assumption corresponds to assuming the gravitational wave frequency spins down at a rate close to the electromagnetic wave frequency. In particular, for the case of free precession we are assuming that the geometry of the precession, essentially the wobble angle, evolves only on the spin-down timescale. For the two component model this would correspond to the steady state lag that occurs after all the glitch transients have died out.

If we take another derivative with respect to time we get

$$\ddot{f}_{GW} = 2 \ddot{f}_{EM} (1 + \delta) + 4 \dot{f}_{EM} \dot{\delta} + 2 f_{EM} \ddot{\delta}. \quad (7.4)$$

If we are willing to assume that $\dot{f}_{EM}/\ddot{f}_{EM}$ and $\dot{\delta}/\ddot{\delta}$ are of the same order as f_{EM}/\dot{f}_{EM} , we can determine a 2nd spin down band of

$$\Delta \ddot{f}_{GW} = 8 \ddot{f}_{EM} \delta. \quad (7.5)$$

This is equivalent to assuming the rate of change of the gravitational wave frequency spin down is close to that of the rate of change of the electromagnetic wave frequency. Plugging in the Crab frequency parameters from Table 7.1 into $\dot{f}_{EM}/\ddot{f}_{EM}$ and f_{EM}/\dot{f}_{EM} , we do find that they are explicitly of the same order, roughly 3×10^{10} to 8×10^{10} .

Given $\delta \approx 10^{-4}$ and the Crab pulsar's frequency parameters, this suggests choosing a parameter space 6×10^{-3} Hz wide in frequency space, 1.2×10^{-13} Hz wide in 1st spin down space, and 1×10^{-23} wide in 2nd spin down space, all centered on twice the radio pulse frequency and frequency evolution of the Crab pulsar. The relevant Crab parameters for the start of the S5 run are detailed in Table 7.1 [28]. A single sky point was all that was necessary due to the well known position of the Crab pulsar.

At the time of this writing, we have access to a majority of the data from LIGO's fifth science run, which has been running since November 4, 2005. However, a large glitch occurred in the Crab pulsar on August 23, 2006 and its uncertain what effect this would have on the gravitational wave frequency evolution. That glitch is thus a natural stopping point in data consideration. This leaves approximately nine months of data available to be

searched. From that nine month data stretch we created 8798 SFTs from H1, 6649 SFTs from L1, and 9877 from H2.

Using the discussion from Chapter 7.2.1, we find this parameter space requires roughly 8.1×10^6 templates when using all the available SFTs. This is with a mismatch parameter of 0.15, which injections show causes a loss of roughly 5% in sensitivity from an exactly matched template. The full search on all three LIGO interferometers with the necessary Monte Carlo upper limit injections takes approximately one day on 100 nodes on the LIGO Caltech computing cluster.

As for final sensitivity, this is affected by both the number of templates and the region of frequency space needed from the LIGO interferometers. For Gaussian noise, the final expected sensitivity scales slowly with templates, as shown in Figure 7-7. However as we increase the parameter space we also add in more noise from the surrounding frequency bins, especially that caused by the 60 Hz lines. This is apparent in the case of L1 which is discussed in the following sections.




Figure 7-7: Plot of the relative sensitivity of a multi-template search relative to a single template search as a function of the number of templates

Table 7.1: Crab parameters at the start of S5

Frequency	29.7772 Hz
1 st Spin down	-3.7292×10^{-10} Hz/s
2 nd Spin down	1.2426×10^{-20} Hz/s ²
Right Ascension	1.4597 radians
Declination	0.3842 radians

Expected and actual largest templates

Given our parameter space with 8.1×10^6 templates, we can estimate the expected largest $2F$ value that will be found, and also a 1% false alarm cutoff value, assuming Gaussian noise. Starting from the noise probability density function, p_0 in Eq. (6.37), we can calculate the probability P_{below} that a single template is at or below a cutoff $2F$ of $2F_c$. That is

$$P_{below}(F_c) = \int_0^{F_c} p_0(F) dF = \frac{1}{2} - \frac{(1 + F_c)}{2} e^{-F_c}. \quad (7.6)$$

To find the probability of N templates being below $2F_c$ we simply raise $P_{below}(F_c)$ to the N^{th} power. To find the probability of one or more templates being above the cutoff we only need to calculate

$$P_{FA} = 1 - P_{below}^N(F_c). \quad (7.7)$$

We can then numerically solve for the $2F_c$ such that P_{FA} is equal to 0.01. We find that $2F_c = 46$ for $N = 8.1 \times 10^6$. This is effectively a 3σ level of confidence requirement on a positive detection result.

We can also calculate the expected largest $2F$ value found. In this case we want $N - 1$ templates being below our expected $2F_e$, our expected $2F$ value, and one template being between $2F_e$ and $2F_e + dF$. We can write the probability distribution for $2F_e$ in N trials as

$$p_N(2F_e) = \frac{N}{2} F_e e^{-F_e} (1 - P_{below}(F_e))^{N-1}. \quad (7.8)$$

We can then compute the expected maximum $2F_e$ numerically in N trials as

$$\int_0^\infty 2F_e p_N(2F_e) d(2F_e) \quad (7.9)$$

which for $N = 8.1 \times 10^6$ is 39.

When we look at the actual result from the full search we find it to be very consistent with Gaussian noise with no signal present. We find an maximum $2F$ value of 38.27 which is below our threshold of detection and very close to our expected value of 39.

7.2.2 Upper limits

Monte Carlo injections

With our largest $2F$ values being below a 1% false alarm cutoff, we now turn to setting upper limits on the maximum gravitational wave strain coming from the Crab pulsar. This is done by Monte Carlo injections. The following methodology is applied.

Choose an initial injection strength, h_0 , that we think is near the minimum detectable level. We fix right ascension and declination to that of the Crab pulsar. Then randomly choose polarization, inclination and initial phase from all possible values. Randomly choose frequency, 1st spin down, and 2nd spin down within the parameter space of the search. With the signal now defined, inject it into the actual data via software. Search for this injected signal. To save time we search only a small space surrounding the signal, including ± 2 template spacings in frequency, 1st spin down, and 2nd spin down. If the largest $2F$ value found in this injection search is greater than the largest $2F$ value found in the actual search, we declare the signal found. We then repeat until we have reached our desired number of injections, at which point we calculate our confidence of that injection run. The confidence is simply the number of injections declared to be “found” divided by the total number of injections made. If the confidence of the injection run is over 0.95 we decrease our injection strength and repeat the run. If the confidence is less than 0.95, we increase the injection strength. We also decrease our step size if we cross the 0.95 confidence level between injection runs. This eventually settles on the injection level which is found 95% of the time. In this way we find the 95% confidence upper limit value.

It is important to note how many injections are sufficient for a run. The number of injections are sufficient if over multiple injection runs at the same strength, the confidence level does not vary too much, say of the order 1%. Having performed a test run with 36000 injections and breaking it into 9 sets of 4000, we find that the confidence value of each small set stayed within 95.3% and 96.2% when the confidence value from all the injections was 95.6%. This is shown in Figure 7-8. This suggests that 4,000 injections are sufficient to

determine the confidence level to within 1%.




Figure 7-8: Plot of the confidence for 9 sets of 4000 injections compared to the confidence for all 36000 injections.

Upper limit result

Having performed the upper limit injections on the search, the upper limit on the strain coming from the Crab pulsar was found with $h_0 = 1.6 \times 10^{-24}$ at a confidence of 95.5% with 20000 injections. This upper limit holds over the parameter space of the search.

7.2.3 Checks on the search result and upper limits

Several checks were performed to determine if the search and upper limits were performed correctly. First, data from each individual interferometer (i.e. H1, H2 and L1) was run through the entire analysis pipeline by itself. These were then compared to a set of software generated Gaussian white noise also run through the analysis pipeline.

SFT noise

Our cuts on the data set used to generate our SFT segments were minimal. The only data removed was from periods out of lock, had bad or non-existent calibration lines, was during

periods of high winds over 30 mph, or within 30 seconds of a lock loss. This means that when we look at individual SFTs it is no surprise that within our frequency band of interest, there are periods where the noise was significantly louder than others. Figure 7-9 shows the mean noise in a frequency band from 59.52 to 59.58 Hz in terms of the one sided power spectral density. This is the band size the search code considered when calculating $S_h^{1/2}$.




Figure 7-9: Plot of spectral density of strain in the 59.52 to 59.58 Hz band versus SFT (in chronological order) for H1 (top), H2 (middle), and L1 (bottom). The many lines coming out of the noise floor for each interferometer shows the non-stationarity of the noise over the course of the S5 run. In order to get a good estimate of the noise floor a noise weighting scheme is necessary in order to deal with this non-stationarity.

These large variations are dealt with in the search code by applying a weight factor based on a running median with a window size of 50 bins. Effectively, for a given data segment, the noisier the frequencies surrounding the one we are interested in the less weight we give to that portion of the data. Since each SFT corresponds to 1800 seconds of data, each bin is $1/1800 = 5.55 \times 10^{-4}$ Hz wide. The median value of these bins is effectively set as the noise floor for the frequency under consideration. In the analysis the calculation is weight by $1/w(f)$ where $w_\alpha(f) = \langle S_{\alpha,k}(f) \rangle_k$ where $\langle \rangle_k$ denotes the median over k bins and $S_{\alpha,k}$ denotes the noise of k bins of the α^{th} SFT. The weighting minimizes the impact excessively noisy segments have on the search.

In order to estimate the effective noise under this weighting scheme, we need to perform a similar weighting on the SFTs. This is done by calculating the mean value of all the bins within a window size of 50 bins and associating that value with the central frequency. We denote this as $\langle SFT_i(f) \rangle_\alpha$ indicating the i^{th} SFT and an average of α bins. We then calculate $\langle 1 / \langle SFT_i(f) \rangle_\alpha \rangle_i^{-1}$ where in this case $\langle \rangle_i$ denotes averaging over all SFTs. This is effectively estimates the double sided power spectral density. Finally we convert to the one sided power spectral density and take the square root to get it into the usual form of $S_h^{1/2}$. This method produces the plot in Figure 7-10.

In addition to running on data from the three LIGO interferometers, the above was also carried out on a Gaussian white noise data set, which was treated as if being generated at Hanford and using the same number of SFTs as H1. This data set was created with a $S_h^{1/2}$ value of 4×10^{-22} with the code *lalapps_Makefakedata*. When looked at in the same way as above, it produces Figures 7-11 and 7-12. From Figure 7-11 we see that the SFTs are of the correct strain and in Figure 7-12 shows that our noise is effectively white, varying by less than 0.1% over the frequency band of interest.




Figure 7-10: Estimate of spectral density of strain for the all three LIGO interferometers in the frequency band of interest for the Crab pulsar. H1 has the best noise performance, followed by L1 and then H2. However, L1 has a large trend upwards with increasing frequency indicating the presence of an instrumental line feature.




Figure 7-11: Plot of spectral density of strain versus a Gaussian white noise data set SFT (in chronological order). This particular data set has been generated with an expected 4×10^{-22} spectral density of strain.




Figure 7-12: Estimate of spectral density of strain versus frequency for the entire Gaussian noise fake data set. This shows the estimation method is working properly, since the estimated strain shown here is within 0.5% of the expected value of 4×10^{-22} .

Individual interferometer search results

The individual interferometers were run through the search pipeline and none of them produced a template with a $2F$ value beyond the 1% false alarm cutoff of 46 for 8.1×10^6 templates. The largest $2F$ value for H1 was 43.67, for H2 was 37.02, and for L1 was 30.43. The largest $2F$ value for the Gaussian Noise data set was 35.83. These can be compared with the overall search's largest $2F$ value of 38.27 and also the expected largest $2F$ value assuming 8.1×10^6 templates of 39.

Using Eq. (6.37) we can estimate the expected number of templates with a given $2F$ value assuming Gaussian noise and normalizing by the total number of templates in the search. This can be plotted directly against the actual number of templates falling within a band of $2F$ values. Because the interferometers spanned slightly different time periods due to their duty cycle, the number of templates used by each individual search was slightly different. The H1 search used 7.9×10^6 templates, the H2 search used 8.1×10^6 templates, and the L1 search used 7.3×10^6 templates. The Gaussian Noise test was based on SFTs with the same time stamps as the H1 SFTs and thus also used 7.9×10^6 templates. The Figure 7-13 shows the expected plotted against the actual in each of these cases. Note that the full search does not record templates with a $2F$ value below 20 for data storage reasons.

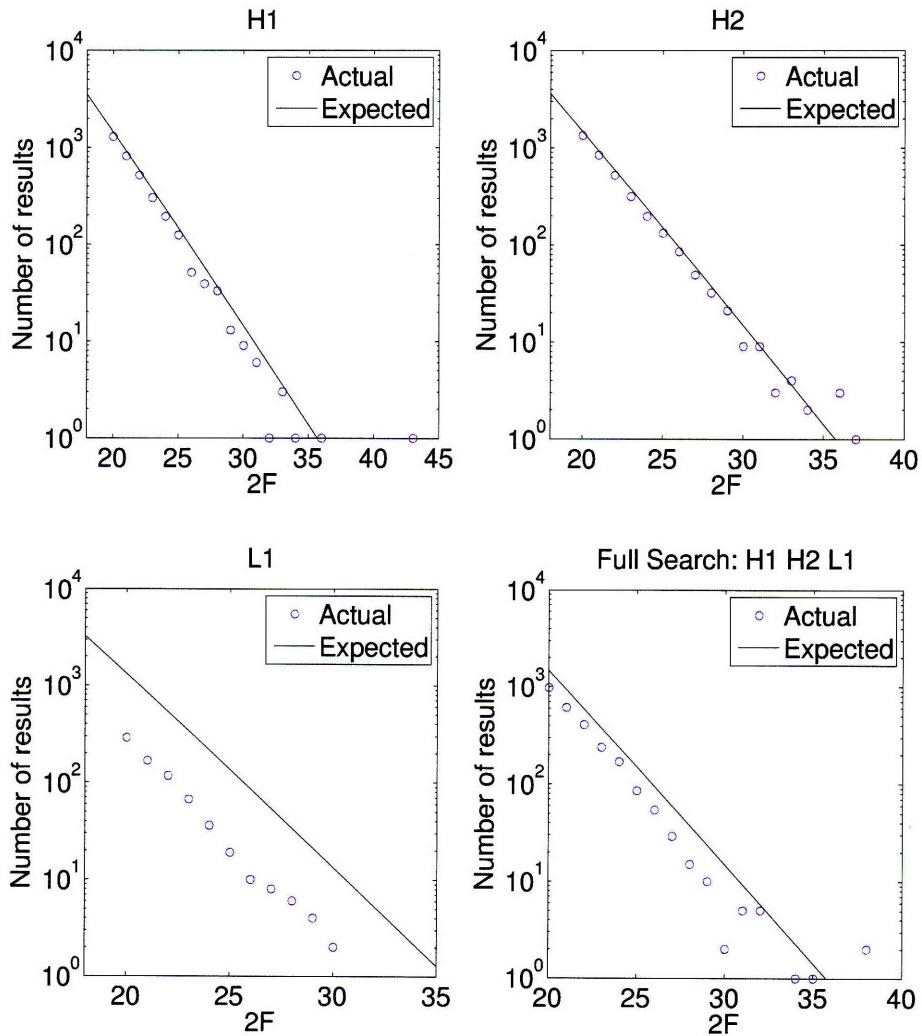


Figure 7-13: For H1 (upper left), H2 (upper right), L1 (lower left) and the full multi-interferometer search (lower right) are shown plots of the number of templates returned by the search with $2F$ values within a band versus the expected number. In the case of L1 there are far fewer large templates than one would expect primarily due to non-random features in the L1 spectrum. The full search being below expected is mainly due to this disturbance in L1.

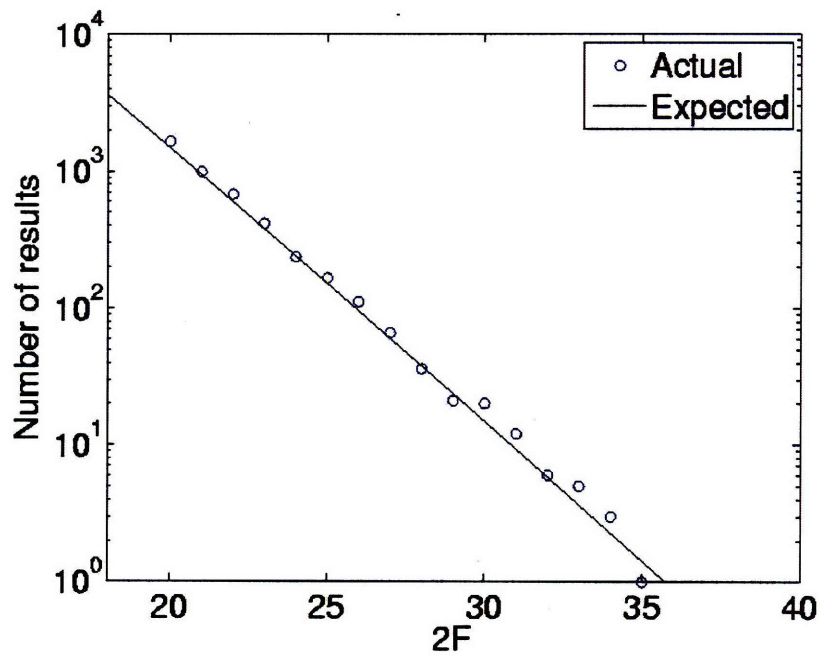


Figure 7-14: Plot of the number of templates returned by a search on a Gaussian noise fake data set with $2F$ values within a certain band versus the expected number. The fact that the Gaussian noise matches the expected value implies that deviation from the expected in Figure 7-13 for L1 is due to the non-Gaussian nature of the noise rather than a bias in the search pipeline.

The first thing to notice in these Figures is the actual interferometer results tend to be below the expected value, while the Gaussian noise data set is at the expected level, most strikingly apparent in the case of L1. Since the Gaussian noise data does not show any bias, this must be an effect of the non-Gaussian and non-stationary nature of the interferometer noise. By comparing the extreme case of L1 to the H1 case we can gain some insight.

We first go and look at the surrounding noise, as shown in Figure 7-15, and compare it with the somewhat cleaner case of H1 in Figure 7-16. There is the obvious large 60 Hz peak in both H1 and L1 caused by 60 Hz line noise. In addition, there appears to be more structure in the L1 case. If we go back to the specific band of interest in Figure 7-10 we note that the L1 noise increases by about 15% over the band relative to its mean value, where as in H1 and H2 the variation is less than 3% of its mean value. Suggesting that the L1 search is more on top of 60 Hz “wings” and associated noise than the respective searches in H1 and H2, but all are obviously non-white to some degree. In this case, our assumptions of Gaussian stationary white noise is poor for L1 hence leading to a break down of our predictions based on Eqs. (6.36) and (6.37). The final upper limits, discussed in Section 7.2.3, are not affected by more than a few percent because we use the largest $2F$ value found in the search as the cutoff, rather than a theoretically calculated cutoff. The direction of the bias comes from the fact that a non-random instrument line brings up the effective noise floor without affecting the Gaussian distribution of noise on top of it. Since the F statistic is proportional to the square of the optimal signal to noise, this leads to a corresponding drop in the F statistic values found.




Figure 7-15: Plot of the weighted noise in L1 from 59 to 60 Hz. Distinct steps are visible in the noise, indicating some non-random instrumental effect is raising the noise floor.




Figure 7-16: Plot of the weighted noise in H1 from 59 to 60 Hz. The really only visible line is the 60 Hz line due to power line noise, while around 59.5 Hz the noise seems effectively flat.

2F distribution in parameter space

We can also take a look at the distribution of the $2F$ values within the parameter space searched, to check if it looks like a random distribution as opposed to clumping near an obvious signal or line. These are shown in Figures 7-17 and 7-18. The templates returned with large $2F$ do look uniformly distributed except possibly in the case of L1, which as already noted in Section 7.2.3 seems to have a broad feature in the noise.

If a strong narrow line were present, it would manifest itself as a diagonal line of large (redder in the plot) values from the upper left trailing to the lower right. This is because of the negative spin down values are being applied to the templates starting frequency over the course of the run to determine what frequency to look at. Larger spin down values (in an absolute sense) results in lower frequencies being searched for that particular template, meaning a narrow line feature will be seen by templates with high frequencies and large spin downs or low frequencies with small spin downs. This suggests that there is no narrow line feature in the searched frequency band causing unusually large results.

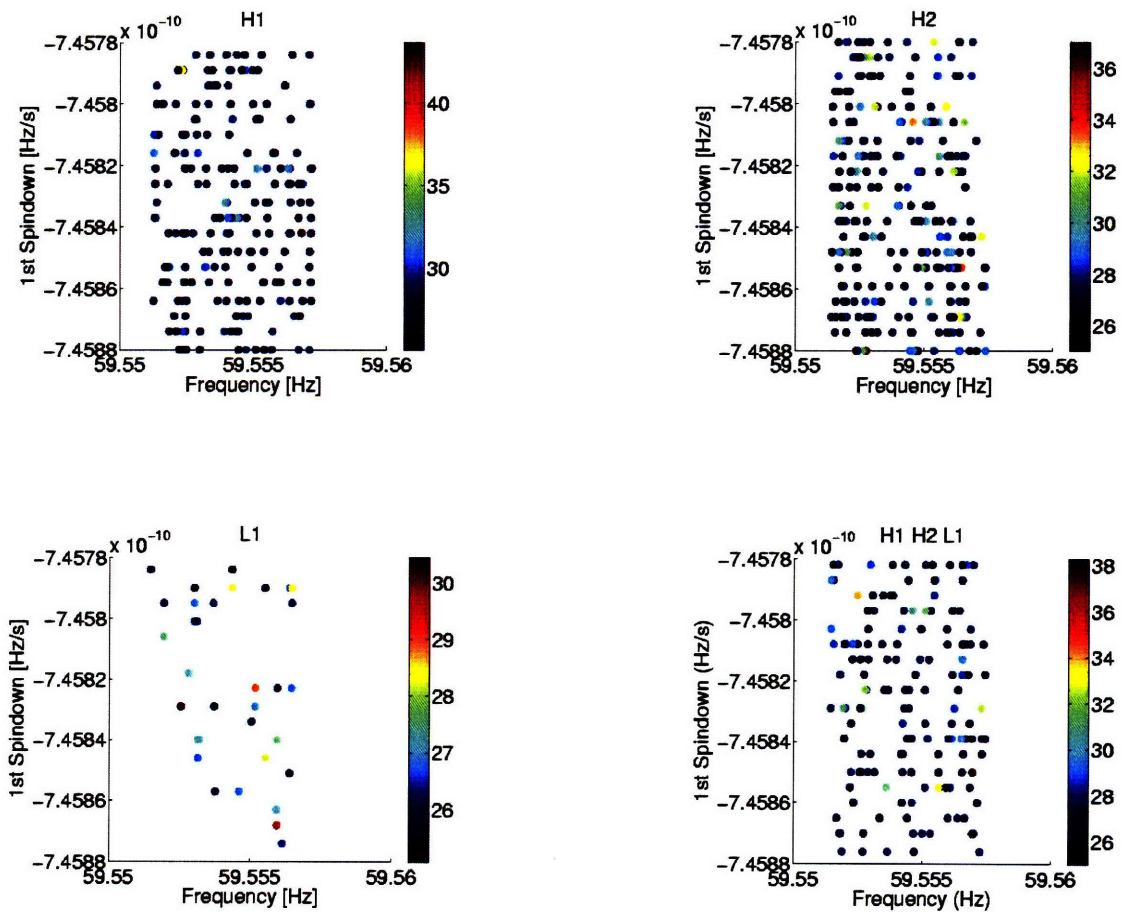


Figure 7-17: For H1 (upper left), H2 (upper right), L1 (lower left) and the full multi-interferometer search (lower right) are shown plots of the distribution of $2F$ values greater than 25 in the parameter space searched. The color signifies the $2F$ value, red is larger and blue is smaller. For H1, H2 and the multi-interferometer search the distributions look uniform, while L1 looks different due to a combination of fewer results to work with and the broad noise feature present in the L1 spectrum. No strong narrow line features appear to be present in any of the interferometers.




Figure 7-18: Plot of the distribution of $2F$ values greater than 25 for a Gaussian noise fake data set. The color signifies the $2F$ value, red is larger and blue is smaller. This is useful as a comparison to the Figures in 7-17 since this fake data set was created without any line or broad spectrum features.

Expected individual interferometer upper limits

Given the maximum $2F$ values found for each individual interferometer and the Gaussian noise test case, one can use Eq. (6.36) to estimate the injection strength necessary to have 95% of the injections exceed the respective maximum $2F$ values. To do so also requires taking into account the detector response functions since the injections have varying polarizations and inclination angles which non-trivially affects the effective detectability of the injection. Lastly, Eq. (6.36) applies to templates perfectly matched to the signal, which is not the case in our search or Monte Carlo injections.

In order to determine the effect of perfectly matching templates, an additional Monte Carlo injection set was performed on the H1 search but with the follow up searches each placed exactly at the injection. The 95% confidence level for strain was found to be $2.12 \times 10^{-24} \text{ Hz}^{-1/2}$, compared to the actual H1 95% confidence level of $2.25 \times 10^{-24} \text{ Hz}^{-1/2}$. This suggests that the grid spacing worsens the upper limit by making it larger by roughly 5%, which should be taken into account when using Eq. (6.36) to make predictions.

We can use Eq. (6.57) along with the detector response functions with the known positions of the Hanford and Livingston detectors, shown in Table 7.2 [4], and the fixed right ascension and declination of the Crab pulsar, to write the following

$$\langle \rho^2 \rangle = A(\psi, i) \frac{h_0^2 T}{S_h} \quad (7.10)$$

where $A(\psi, i)$ encodes the detector response function dependence on just the polarization angle ψ and inclination angle i .

The function $A(\psi, i)$ can be calculated from the detector response functions [27] by fixing the declination of the source and averaging over the right ascension. The function $A(\psi, i)$ for the Hanford and Livingston locations are plotted in Figures 7-19 and 7-20. It effectively represents the changing response of the interferometers for different possible polarizations of the wave and for the different possible inclination angles of the source relative to the interferometers. By considering this function over the entire injection space and using the calculated ρ in Eq. (6.36) we can calculate the overall expected 95% confidence level, or at any other desired confidence level. It should be noted because of the 95% confidence level requirement, the lowest 5% of values from Figures 7-19 and 7-20 determine the effective detection threshold, as opposed to an average over the whole parameter space. In any

case, these expected values can then be directly compared to upper limit injections at and around the 95% confidence level derived from actually performing the injections. These are shown in Figures 7-21. As seen in the plots, the difference between the expected and actual confidence values are less than 5%.

The good agreement between the actual and expected values shows that the analysis pipeline is working and that the results are consistent with being generated by mostly Gaussian noise. The upper limits seem to be reasonably robust in the face of the noise features in L1 since we used the largest $2F$ value found of 30 which effectively factors in the real distribution of the noise into the upper limits. So while the noise in L1 is not perfectly Gaussian, actual signals above the noise floor are still detected without problems.

Table 7.2: Positions of Hanford and Livingston detectors

Detector	Latitude (degree)	Longitude (degree)
Hanford	46.45 N	119.41 W
Livingston	30.56 N	90.77 W




Figure 7-19: Plot of the function $A(\psi, i)$ for Hanford. Effectively, the physical response of the interferometer as a function of the incoming wave's polarization and source inclination, with the sky position fixed to that of the Crab pulsar. Redder colors correspond to higher values of $A(\psi, i)$.




Figure 7-20: Plot of the function $A(\psi, i)$ for Livingston. Effectively the physical response of the interferometer as a function of the incoming wave's polarization and source inclination, with the sky position fixed to that of the Crab pulsar. Redder colors correspond to higher values of $A(\psi, i)$.

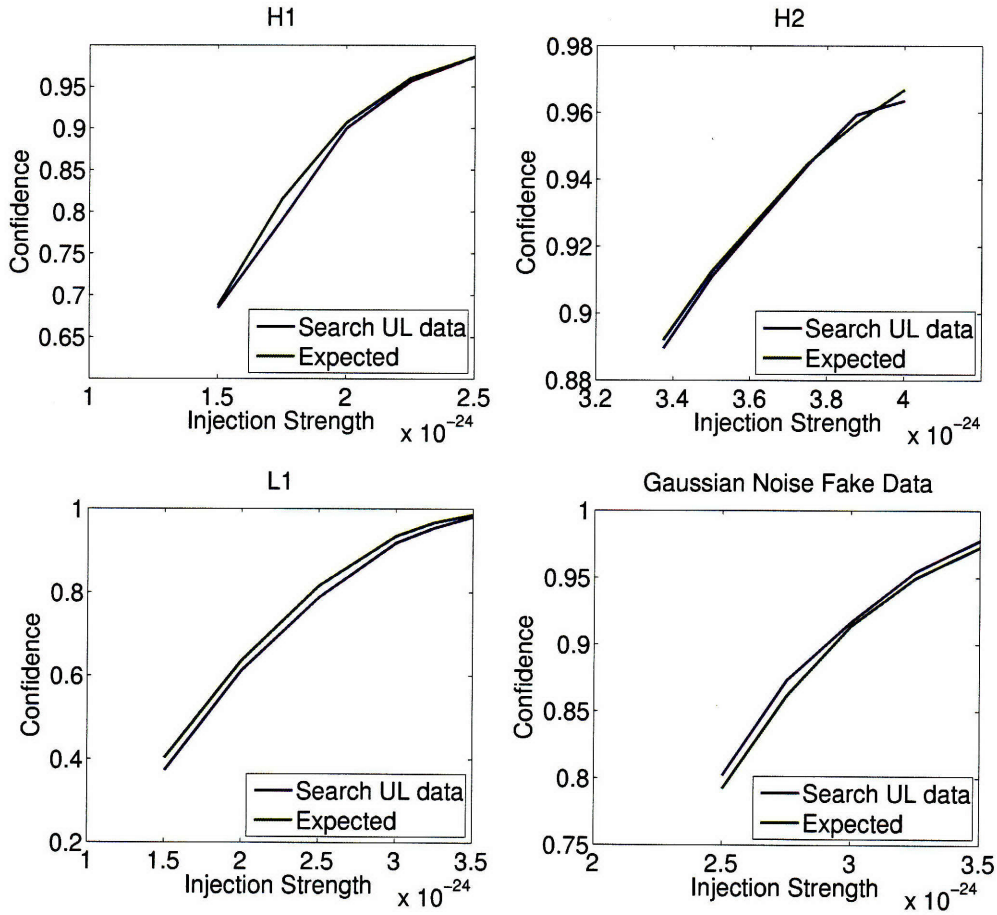


Figure 7-21: For H1 (upper left), H2 (upper right), L1 (lower left) and a search on a set of Gaussian noise fake data (lower right) are shown plots of the expected confidence versus injection strength of the Monte-Carlo injections. The position of the curves are determined by the Crab pulsar's sky position, the function $A(\psi, i)$, the noise in the individual interferometers, along with the largest $2F$ value found in the individual searches. The shape of the curves are determined by Eq. (6.36). The difference between the expected and actual values is in general less than the uncertainty in the upper limit result for runs of 4000 injections, which gives us confidence in the upper limit values produced by the search pipeline.

7.2.4 Constraining energy loss in the Crab pulsar

With a strain upper limit of 1.6×10^{-24} that we have confidence in, we can now turn to the question of what we can say about the Crab pulsar.

Equation 6.22 relates our h_0 upper limit to the moment of inertia, ellipticity, distance and spin frequency of the Crab pulsar. We know the Crab pulsar is located $2kpc$ distant, and we know the radio frequency is 29.777 Hz. If we make the assumption that the Crab pulsar has the a typical neutron star's moment of inertia of 10^{45} g cm², we can calculate the maximum ellipticity the Crab pulsar could have. In this case, we calculate $\epsilon_{max} = 8.5 \times 10^{-4}$. This corresponds to a gravitational luminosity of 4.6×10^{38} erg/s. These values are larger than the indirectly calculated values of $\epsilon_{max} = 3 \times 10^{-4}$ and 1.6×10^{38} erg/s estimated by Palomba [43].

However, there is a good deal of uncertainty in the moment of inertia of the Crab pulsar, and as noted earlier it might be up to a factor 3 larger than the canonical value [7] [34]. For a value of I of 1.59×10^{38} kg m² or higher we begin beating the canonical upper limit on the Crab pulsar's gravitational wave emission. For a value of I of 3×10^{38} kg m² we can constrain the energy loss due to gravitational waves within the search parameter space to less than 53% of the total spin down energy loss of the Crab pulsar. The relationship between ellipticity and the moment of inertia is plotted in Figure 7-22. The ratio of the energy loss due to gravitational radiation versus the total spin down energy loss is plotted in Figure 7-23.




Figure 7-22: Plot of the upper limits generated by this search on ellipticity as a function moment of inertia compared to that of the canonical upper limit calculation. The area above the green line is effectively excluded by theoretical calculations, while the area above the blue line has been excluded by direct measurement in this search.




Figure 7-23: Plot of the ratio of gravitational energy loss upper limit to total spin down energy loss as a function of moment of inertia. Ratio values greater than 1 indicate we are not constraining the Crab pulsar's energy budget of its spin down at that moment of inertia.

Figure 7-22 effectively shows the excluded regions of ellipticity and moment of inertia parameter space. The area above the green line is excluded by theoretical calculations, while the area above the blue line is excluded by direct measurement in this search. The wedge to the right of the crossing point is effectively the additional information the search is providing.

In Figure 7-23 values above one are effectively unphysical and thus the search is effectively providing no information in that region. However, for the higher values of the moment of inertia, we are providing a directly measured constraint on a portion of the Crab pulsar's energy budget of the spin down. This direct measurement is interesting, because even though there are estimates of the Crab pulsar's spin down energy budget, the uncertainties involved in those estimates leave a significant portion of the budget available for gravitational wave emission.

We will briefly consider the total spin down energy budget of the Crab pulsar. If we take the Crab pulsar's frequency and 1st spin down values, and let the moment of inertia range from 0.5×10^{38} to 8×10^{38} kg m² [15], the total energy available is $\dot{E} = 4\pi^2 I f \dot{f} = 2 \times 10^{38}$ to 1.6×10^{39} erg/s. This energy loss can be made up of several components in addition to the gravitational wave emission, including the observed electromagnetic emission from the pulsar and the surrounding nebula, and also the nebula's accelerating expansion.

The Crab pulsar's black-body thermal emission is a very minor component, less than $\sim 3 \times 10^{34}$ erg/s, assuming a 10 km radius star [26]. The total bolometric luminosity of the Crab nebula is roughly $1 - 2 \times 10^{38}$ erg/s [17]. Lastly, the energy necessary for the nebula expansion could be between 3×10^{38} and 9×10^{38} erg/s [8] [9]. The uncertainties in the nebula expansion estimates are large due to the uncertainty in the mass of the Crab Nebula itself, which is between 2.8 and 6.4 M_{\odot} [21]. Using these values, we can estimate the total loss to non-gravitational wave mechanisms as being between 4×10^{38} and 10×10^{38} erg/s.

In comparing the possible range of rotational energy loss, 2×10^{38} to 1.6×10^{39} erg/s based on possible moments of inertia, to the non-gravitational wave losses, 4×10^{38} to 10×10^{38} erg/s, we see there is a large variation in the difference, which we might attribute to gravitational radiation. From this comparison, we might estimate that anywhere from a negligible to 75% of the energy loss might be in the form of gravitational radiation. This then makes limits from the search presented here or from the narrow time domain search [44] useful in understanding the overall energy budget.

Part IV

Conclusion

Conclusion

In the first half of this thesis, we describe the LIGO interferometers and the phase camera subsystem. We discuss how the phase camera was used to understand an output mode cleaner placed at the AS port of the interferometer and how it was shown to be working as expected. We also discuss the thermal compensation system, and how it was a necessary component for the interferometers to reach high power (and higher sensitivity) operation. We examine the impact of the thermal compensation system on the AS port light through the use of the phase camera. We compare different states of thermal compensation in order to understand how the light changes to reduce signal in the Iphase of the AS port photodiodes and reduce the need for a correction servo, whose limit had previously determined the maximum laser power LIGO could use.

In the second half of this thesis, we performed a single target search for continuous gravitational wave emission from the nearby star RX J1856.5-3754 using data from LIGO's 4th science run. While a periodicity in RX J1856.5-3754's X-ray emission was found below the LIGO search band, the search itself was still a useful test of a newly developed multi-interferometer F-statistic code and an application of targeted search methods. This experience was directly applicable to a second search looking for continuous gravitational wave emission from the Crab pulsar, searching over a wide parameter space to complement very narrow targeted searches already performed. This parameter space was chosen based on physical motivations, computational constraints, and expected final sensitivity. Using data from the first nine months of LIGO's 5th science run to perform the search, no gravitational wave signal was found. Monte Carlo injections were used to place a Frequentist upper limit on the gravitational wave strain from the Crab pulsar of $h_0 = 1.6 \times 10^{-24}$ at a 95% confidence level. This result is currently under review by the LIGO Scientific Collaboration, and as such should be treated as preliminary.

As LIGO's 5th science run finishes, we can look forward to examining the entire data set, which extends the possible depth of targeted continuous wave searches. In addition a period of upgrades to the LIGO interferometers aimed at improving their sensitivity by a factor of 2 [22] is scheduled to start soon after, extending the reach of gravitational wave searches even further.

Appendix A

Table of acronyms

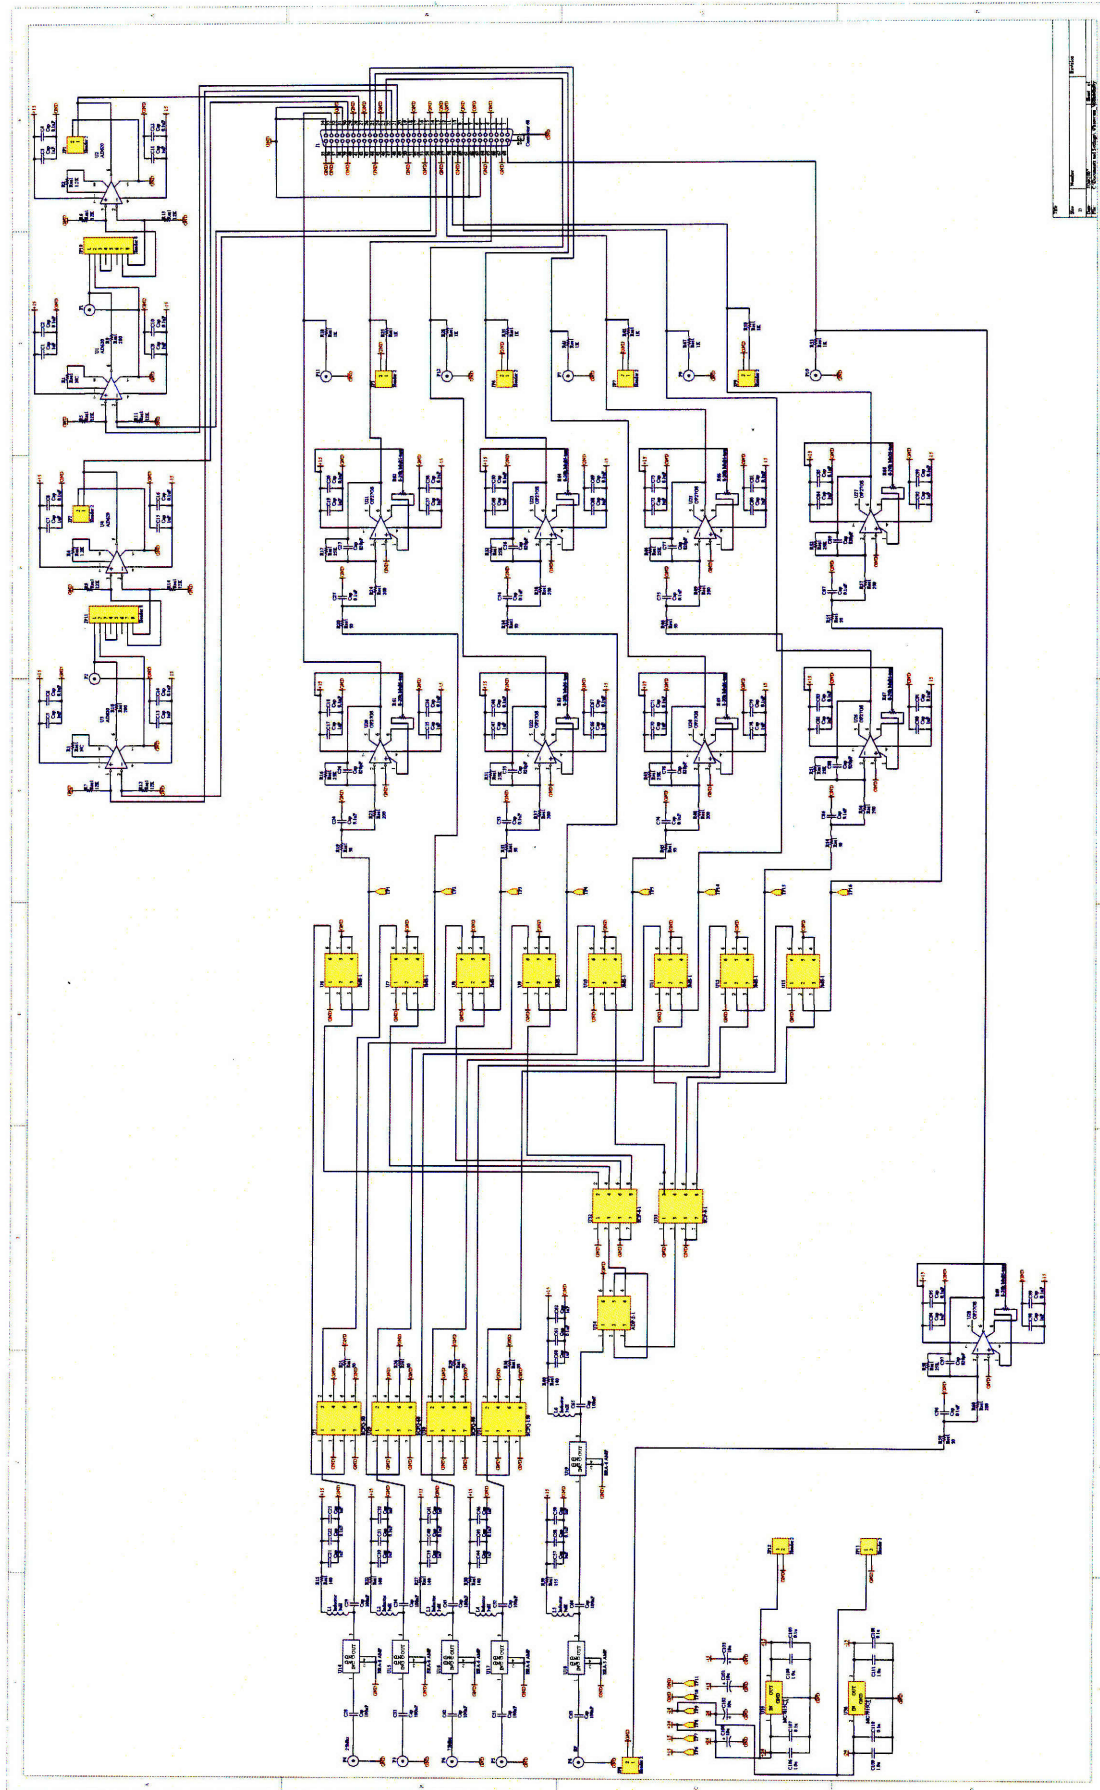
Definitions of acronyms	
ACRONYM	DEFINITION
RM	Recycling Mirror
BS	Beam Splitter
ITMX	Input Test Mass, X-Arm
ITMY	Input Test Mass, Y-Arm
ETMX	End Test Mass, X-Arm
ETMY	End Test Mass, Y-Arm
TCS	Thermal Compensation System
PSL	Pre-Stabilized Laser
PMC	Pre-Mode Cleaner
PC	Pockels Cell
MC	Mode Cleaner
OMC	Output Mode Cleaner
AOM	Acoustic-Optical Modulator
PZT	Piezoelectric Transducer
NSPOB	Normalized Sideband Power Pick-off Beamsplitter
AS_I	Anti-symmetric Port Photodiode In-phase Signal
AS_I.Corr	Anti-symmetric Port Photodiode In-phase Correction
USB	Upper Side Band
LSB	Lower Side Band
Iphase	In-phase
Qphase	Quadrature-phase

Table A.1:

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

Phasecamera demodulation and control board schematic



Bibliography

- [1] B. Abbot, M. Kramer, and A.G. Lyne et al. Upper limits on gravitational wave emission from 78 radio pulsars. *arXiv*, pages gr-qc/0702039, 2007.
- [2] B. Abbot et al. Coherent searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: results from the second LIGO science run. *arXiv*, pages gr-qc/0605028v2, 2006.
- [3] Rana Adhikari. *Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [4] B. Allen. Gravitational wave detector sites. gr-qc 9607075.
- [5] S. Ballmer. *LIGO interferometer operating at design sensitivity with application to gravitational radiometry*. PhD thesis, Massachusetts Institute of Technology, 2006.
- [6] S. Ballmer, V. Frolov, R. Lawrence, W. Kells, G. Moreno, K. Mason, D. Ottaway, M. Smith, C. Vorvick, P. Willems, and M. Zucker. Thermal compensation system description. Technical report, 2005. LIGO-T050064-00-R.
- [7] M. Bejger, T. Bulik, and P. Haensel. Constraints on the dense matter EOS from the measurements of PSR J0737-3039A moment of inertia and PSR J0751+1807 mass. *MNRAS*, 364, 2005.
- [8] M. Bejger and P. Haensel. Moments of inertia for neutron and strange stars: Limits derived for the Crab pulsar. *A&A*, 396, 2002.
- [9] M. Bejger and P. Haensel. Accelerated expansion of the Crab nebula and evaluation of its neutron-star parameters. *A&A*, 405, 2003.

- [10] J. Betzwieser, K. Kawabe, and L. Matone. Study of the output mode cleaner prototype using the phasecamera. Technical report, 2004. LIGO-T040156-00-D.
- [11] S. Bonazzola and E. Gourgoulhon. Gravitational waves from pulsars: emission by the magnetic-field-induced distortion. *A&A*, 312, 1996.
- [12] V. Burwitz, F. Haberl, and R Neuhauser et al. The thermal radiation of the isolated neutron star RXJ 1856.5-3754 observed with Chandra and XMM-Newton. *A&A*, 399, 2003.
- [13] V. Burwitz et al. The Chandra LETGS high resolution x-ray spectrum of the isolated neutron star rxj1856.5-3754. *A&A*, 379, 2001.
- [14] B. Carter, D. Langlois, and D. Sedrakian. Centrifugal bouyancy as a mechanism for neutron star glitches. *A&A*, 361, 2000.
- [15] G. Cook, S. Shapiro, and S. Teukolski. Rapidly rotating neutron stars in general relativity: realistic equations of state. *ApJ*, 424, 1994.
- [16] C. Cutler. Gravitational waves from neutron stars with large toroidal B fields. *Phys. Rev. D*, 66:084025, 2002.
- [17] K. Davidson and R. A. Fesen. Recent developments concerning the Crab nebula. *Ann. Rev. Astron. Astrophys.*, 23, 1985.
- [18] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31, 1983.
- [19] R. Dupuis. *Bayesian searches for gravitational waves from pulsars*. PhD thesis, University of Glasgow, 2004.
- [20] A. Einstein. *The Meaning of Relativity*. Princeton University Press, 5th edition, 1922.
- [21] R. A. Fesen, M. J. Shull, and P. A. Hurdford. An optical study of the circumstellar environment around the Crab nebula. *Astron. J.*, 113, 1997.
- [22] P. Fritschel, R. Adhikari, and R. Weiss. Enhancements to the LIGO S5 detectors. Technical report, 2005. LIGO-T050252-00-I.

- [23] P. Fritschel, R. Bork, G. Gonzalez, N Mavalvala, D. Ouimette, H. Rong, D. Sigg, and M. Zucker. Readout and control of a power-recycled interferometric gravitational-wave antenna. *Applied Optics*, 40, 2001.
- [24] N. Glendenning. *Compact Stars: Nuclear Physics, Particle Physics, and General Relativity*. Springer-Verlag, New York, 2000.
- [25] K. Goda, D. Ottaway, B. Connelly, R. Adhikari, N. Mavalavala, and A. Gretarsson. Frequency-resolving spatiotemporal wave-front sensor. *Optics Letters*, 29, 2004.
- [26] F. R. Harnden, Jr. and F.D. Seward. Einstein observations of the Crab nebula pulsar. *ApJ*, 283, 1984.
- [27] P. Jaranowski, A. Krolak, and B. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars. I. The signal and its detection. *Phys. Rev. D*, 58, 1998.
- [28] Jodrell Bank Crab Pulsar Monthly Ephemeris. <http://www.jb.man.ac.uk/~pulsar/crab.html>.
- [29] D. I. Jones. Frequency selection for targeted GW searches. Discussion during LIGO Continuous Wave Group meeting.
- [30] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *MNRAS*, 331, 2002.
- [31] D.L. Kaplan, M.H. van Kerkwijk, and J. Anderson. The parallax and proper motion of RXJ 1856.5-3754 revisited. *ApJ*, 571, 2002.
- [32] C. Kittel. *Introduction to solid state physics*. Wiley, 8th edition edition, 1995.
- [33] H. Kogelnik and T. Li. Laser beams and resonators. *Applied Optics*, 5, 1966.
- [34] B. D. Lackey, M. Nayyar, and B. J. Owen. Observational constraints on hyperons in neutron stars. *Phys, Rev. D.*, 73, 2006.
- [35] L.D. Landau and E.M. Lifshitz. *The Classical Theory of Fields*. Pergamon Press, 1951.
- [36] R. Lawrence. *Active Wavefront Correction in Laser Interferometric Gravitational Wave Detectors*. PhD thesis, Massachusetts Institute of Technology, 2003.
- [37] A. Lyne and F. Graham-Smith. *Pulsar Astronomy*. Cambridge University Press, Cambridge, 2006.

- [38] A. G. Lyne, F. G. Smith, and R. S. Pritchard. 23 years of crab pulsar rotational history. *MNRAS*, 265, 1993.
- [39] G. Mendell and M. Landry. StackSlide and Hough search SNR and statistics. Technical report, 2005. LIGO-T05003-00.
- [40] C. Misner, K. Thorne, and J. Wheeler. *Gravitation*. W.H. Freeman and company, New York, 1973.
- [41] D. Ottaway, J. Betzwieser, S. Ballmer, S. Waldman, and W. Kells. In situ measurement of absorption in high-power interferometers by using beam diameter measurements. *Optics Lett.*, 31(4), 2006.
- [42] B. J. Owen. Maximum elastic deformations of compact stars with exotic equations of state. *Phys. Rev. Lett.*, 95, 2005.
- [43] C. Palomba. Gravitational radiation from young magnetars: Preliminary results. *A&A*, 367, 2000.
- [44] Pitkin, M. for the LSC. Beating the Crab pulsar spin-down and other upper limits: New results from the LIGO S5 known pulsar search. APS meeting, April 2007. <http://meetings.aps.org/Meeting/APR07/Event/65424>.
- [45] R. Prix. Monte-Carlo tests of ComputeFStatistic for 4D parameter-spaces. Technical report, 2005.
- [46] R. Prix. Notes on ComputeFStatistic_V2. Technical report, 2006.
- [47] Prix, R. for the LSC. The search for gravitational waves from spinning neutron stars. Technical report, 2006. LIGO-P060039-00-Z.
- [48] M. Smith, P. Ottaway, P. Willems, C. Vorvick, and G. Moreno. Heating beam pattern optical design. CO2 laser thermal compensation bench. Technical report, 2004. LIGO-T040057-01-D.
- [49] J. H. Taylor and J. M. Weisberg. Further experimental tests of relativistic gravity using the binary pulsar PSR 1913 + 16. *ApJ*, 345, 1989.

- [50] A. Tiengo and S. Mereghetti. XMM-Newton discovery of 7 s pulsations in the isolated neutron star RXJ 1856.5-3754. *ApJ*, 657, 2007.
- [51] G. Ushomirsky, C. Cutler, and L. Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *MNRAS*, 319, 2000.
- [52] F.M. Walter. The proper motion, parallax, and origin of the isolated neutron star RXJ 185635-3754. *ApJ*, 549, 2001.
- [53] F.M. Walter and J. M. Lattimer. A revised parallax and its implications for RXJ 185635-3754. *ApJ*, 576, 2002.
- [54] F.M. Walter, S.J. Wolk, and R. Neuhaeuser. Discovery of a nearby isolated neutron star. *Nature*, 379, 1996.
- [55] M. Zimmermann. Gravitational waves from rotating and precessing rigid bodies. II. General solutions and computationally useful formulas. *Phys. Rev. D*, 21, 1980.
- [56] M. Zimmermann and E. Szedenits. Gravitational waves from rotating and precessing rigid bodies: Simple models and applications to pulsars. *Phys. Rev. D.*, 20, 1979.