


Directed searches for continuous gravitational waves from spinning neutron stars in binary systems

by
Grant David Meadors

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in The University of Michigan
2014

Doctoral Committee:

Professor John Keith Riles, Chair
Professor Fred C. Adams
Professor Nuria Pilar Calvet
Research Scientist Herold Richard Gustafson
Professor Timothy A. McKay
Professor Stephen C. Rand



© Grant David Meadors 2014
All Rights Reserved

To the tree of Life, which took stardust and evolved into us.
Pro arbore Vitae, ex nube stellarum ad nos evoluit.

ACKNOWLEDGEMENTS

Thanks should go beyond a simple page. Lest I forget, let me reflect on all the people without whom I would not have made it here. To my parents, Erin O'Rourke-Meadors and Gregory David Meadors, and my brother, Patrick Thomas Meadors. Home's name follows you; three decades have we explored – beyond seas, roots grow. My dear grandmother, Florenceann O'Rourke (née Williams), supported my undergraduate studies at Reed College. Aunt Nan & Uncle Bud Williams made book-reading & museum-going fond memories. Hanford gave time to learn my family's wonderful stories. If only my paternal grandparents, Shuler and Jeanne (Brown) Meadors, could be here too. Ethan Obie Romero-Severson is ever a devoted friend in adventures & colleague in mathematics: Team Science for four years and counting!

Keith Riles is a dedicated and conscientious adviser: steady progress is how new science is born – detecting gravitational waves is a tricky task, and his thoughtful attention to details will be part of what makes it possible. Dick Gustafson introduced me to the field in 2005 and has helped ever since; my only regret working with him is that we never flew the glider, although we nearly won a sailboat regatta.

Evan Goetz pioneered this kind of search: Oregon makes good scientists. Vladimir Dergachev loves the lab: *bolshoye spasibo!* Best to my office-mates, Jax Sanders & Santiago Caride, on their way soon! Keita Kawabe is a concerted collaborator and experimenter, attentive and dilligent. Chris Messenger ever-patiently organized the data challenge and has a keen eye for the future. Marco Cavaglia orchestrated the

fun of the World Science Festival exhibition.

Mike Landry encouraged me to come for my sojourn at Hanford and, with Fred Raab, made it a welcoming place. Sheila Dwyer and Sheon Chua, helped me see the squeezed light at the end of the beam tube. Conor Mow-Lowry and Alexander Khalaidovski, were likewise great squeezer friends. Our excellent operators, Cheryl Vorvick, Gerardo Moreno, Jonathan Berliner and Dani Atkinson, were tireless in reviving H1. Sam Waldman kindly helped us replace the OMC. Lisa Barsotti & Matt Evans and Daniel Sigg made our whole quantum optics experiment possible at Hanford, with Nergis Mavalvala at MIT keeping all running smoothly behind the scenes. Greg Mendell got me started on cluster computing and went far to help me learn about programming in science. Dale Ingram helped me lead my first public LIGO tours. Rana Adhikari & Jenne Driggers are great with filters. Ian Harry & Jess McIver have been invaluable in seeing my feedforward work in the larger world of Detector Characterization. Joey Key gave me a unique venue to present that work too! Jeff Kissel knows how to wear a bowtie, and talk about servos, with style. Nicolás Smith-Lefebvre has been there in LIGO since I started nine years ago. May physics interest us all for years to come.

To my physics friends at Michigan, one of the most tight-knit cohorts. Michelle Adan, Matt Bales & Amanda Boomer, Kevin Bergemann, Alex Burgers and Bo Zhang, thank you for all having organized our Get-a-Job meetings to get over the final hurdles. Melinda Morang helped me find little brown birds, and my friend Beverly Lau from Reed helped me appreciate all animals. Friends are ever in diaspora, but I am fortunate to have kept up philosophy with my friend Suzanne Kaufmann over the years. Miriam related myriad facets of life. From Reed, my friends Gillian Woodruff & Jesse Hallett as well as Carrie Thomas & Evan Pierce included me in

their weddings – congratulations on your little ones! – and Dash Vitullo’s perspective (and particle accelerator tour) anchored me.

Reed’s professors in Portland helped launch me to Michigan. Robert Reynolds advised my undergraduate thesis and showed how really to use a telescope. David Griffiths filled many students, including me, with enthusiasm for physics that lasts to this day. Johnny Powell always had sound advice. Joel Franklin helped me appreciate what gravitational waves meant. Tom Wieting almost converted me to mathematics with his inimitable style of geometry. Ginny Hancock has a sound ear for music and life’s many challenges. Directing the Reactor, Stephen Frantz and Rachel Barnett created the critical conditions for practicing science and meeting dear friends: our elements are fissioned now, but that core made us into new isotopes.

Epidemiology compatriots have made Ann Arbor my home: thanks to Meghan Milbrath (follow the bees), Ian Spicknall (thanks for the coffee), Erin Shellman and Bryan Mayer (who read my whole thesis) – everyone in Edit Club is a Doctor now!

Psychology teammates, Emily Bonem, Ali Earl, and Stephanie Carpenter, brought me through a long winter: good luck, loose seals. Even the sky is not the limit for my Michigan Flyers instructors, Anthony Sumner and especially (Amelia C.) Jayne. And to my dear friends Aric Knuth of English and James Leija of the University Musical Society, who revitalized a love of literature and art, and who brought me in one Thanksgiving. Thanks I give now again.

May my coming endeavor at the Albert Einstein Institute in Hannover help us travel through, and shrink, space and time, just like a gravitational wave. Read on – for more still, named within, made this work a reality.

This dissertation bears LIGO Document number DCC-P1400102 ¹.

¹For further questions, please contact the author: gmeadors@umich.edu

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	x
LIST OF TABLES	xix
LIST OF APPENDICES	xx
CHAPTER	
I. Introduction	1
II. Gravitational Waves and LIGO	5
2.1 Cosmic sources of gravitational waves	5
2.1.1 History from General Relativity	10
2.1.2 Contrast with electromagnetic and particle astronomy	12
2.2 General relativity	14
2.2.1 Symmetry and action principles	14
2.2.2 Derivation of field equations	15
2.2.3 Radiation from quadrupoles	20
2.3 Astrophysical estimates	23
2.3.1 Sources: burst, continuous, inspiral and stochastic	23
2.3.2 Continuous waves from neutron stars	26
2.4 Laser Interferometer Gravitational-wave Observatories	28
2.4.1 From Weber bars to interferometry	28
2.4.2 Gravitational wave interferometry in theory	29
2.4.3 Interferometer theory in practice	32
2.4.4 Observatory operation	33
2.5 Detector characterization and development	37
2.5.1 Feedforward filtering	38
2.5.2 Phase camera	41
2.5.3 Advanced observatories and beyond	44
2.5.4 Worldwide network	46
2.6 Summary	47
III. Feedforward: Auxiliary MICH-PRC Subtraction	49
3.1 Introduction	50
3.2 Description of the feedforward method	52
3.2.1 Auxiliary noise coherence at sensitive frequencies	55

3.2.2	Estimating filters	56
3.3	Feedforward in- and out-of-loop methods	61
3.3.1	Manually designed rational filtering in-loop	62
3.3.2	Frequency-domain automated filter design	62
3.3.3	Wiener filters	64
3.3.4	Prospects for near-real-time filtering	64
3.4	Safeguard and veto methods	65
3.4.1	Calibration integrity	65
3.4.2	Runtime safeguards	66
3.4.3	Post-processing safeguards	67
3.5	Feedforward with MICH and PRC channels	69
3.5.1	Filter fitting across science segments	72
3.6	Results of feedforward	75
3.6.1	Post-processing diagnostics	75
3.6.2	Feedforward benefits and potential	75
3.7	Conclusion	83
3.8	Postscript	85
IV. Squeezing: Quantum Vacuum Phase Noise		92
4.1	Squeezing theory	92
4.1.1	Problems with lasers: thermal compensation	93
4.1.2	Quantum shot noise and radiation pressure	94
4.1.3	Squeezing filter cavities	98
4.2	LIGO Hanford Observatory quantum vacuum squeezing	98
4.2.1	Collaboration and contributions	98
4.2.2	Squeezing’s scientific benefit	106
4.3	Squeezing large interferometers: success and prospects	108
V. TwoSpect: Search for a Simulated Scorpius X-1		109
5.1	Neutron stars in binary systems	109
5.1.1	Continuous gravitational waves from neutron stars	110
5.1.2	Binary spin-up and detectable lifetime	111
5.1.3	Detection rate projections	113
5.2	TwoSpect searches	113
5.2.1	Two spectra: a double Fourier transform	114
5.2.2	Inferring neutron stars with companions	117
5.2.3	TwoSpect algorithm detection statistic	117
5.3	Directed TwoSpect’s greater sensitivity	118
5.3.1	Sky maps using exact templates	122
5.4	Scorpius X-1 mock data challenge	122
5.4.1	Fully-templated search for Scorpius X-1	122
5.4.2	Narrow-band heat maps in parameter space	122
5.4.3	Wide-band heat maps in parameter space	126
5.4.4	Revisiting & refining detection criteria	126
5.5	Mock Data Challenge procedure	129
5.5.1	Detection Claims	130
5.5.2	Detection criteria	131
5.5.3	Parameter Estimation	131
5.5.4	Upper Limits and Detection Efficiency	133
5.5.5	$\cos i$ Ambiguity	133
5.6	Summary of the MDC	146
5.6.1	Mock data challenge results	146

5.6.2	Binary search summary beyond the MDC	147
5.7	Plans for improvement	147
5.7.1	Coherent interferometer synthesis	148
5.7.2	Circular & elliptical polarization	149
5.7.3	Orbital phase & beyond	150
5.7.4	Parameter space patterns	150
5.7.5	Relevance to follow-up	151
5.8	Addendum	152
5.8.1	Scorpius X-1 parameters	152
VI. Directed TwoSpect: Scorpius X-1 and XTE J1751-305		154
6.1	Directed TwoSpect	154
6.1.1	Targeted, directed and all-sky search sensitivity	155
6.1.2	Enhancements enabled by directed searching	157
6.2	Quantifying directedness: sensitivity studies in real data	158
6.2.1	Real S6 data: detection efficiency	159
6.2.2	Real S6 data: h_0 recovered vs injected	161
6.3	Scorpius X-1 search using Directed TwoSpect in S6	162
6.3.1	S6: Scorpius X-1 search plan	162
6.3.2	S6: Scorpius X-1 heatmaps	163
6.3.3	S6: Scorpius X-1 upper limits, random polarization	163
6.3.4	S6: Scorpius X-1 outliers	166
6.4	XTE J1751-305 search using Directed TwoSpect in S6	169
6.4.1	S6: XTE J1751-305 background	169
6.4.2	S6: XTE J1751-305 heatmaps	169
6.5	Summary of Directed TwoSpect S6 searches	176
VII. Outreach: Exhibit at the World Science Festival		178
7.1	Prototypes: travelling kiosks and the Ann Arbor Hands-On Museum	179
7.2	World Science Festival interferometer manufacture	181
7.2.1	Laser, optics and display	183
7.2.2	Aluminum baseboard	185
7.2.3	Plexiglass enclosure	186
7.3	Exhibitions: New York City, Portsmouth, Fort Wayne	190
7.3.1	Exhibit overview	191
7.3.2	World Science Festival 2010	192
7.3.3	Portsmouth and Fort Wayne	192
7.4	Future LIGO outreach	194
VIII. Conclusion		195
8.1	Cycles of science	195
8.1.1	Improvements to observatories	195
8.1.2	Understanding instruments, refining data	196
8.1.3	Searching deep-space	197
8.1.4	Reaching out, looking up	197
8.2	Scientific merit: filtering and analysis	198
8.2.1	Feedforward improvement to LIGO data	198
8.2.2	TwoSpect directed search for neutron stars in binary systems	199
8.3	Entering the advanced detector era	199
8.4	Vision of a dark sky	200

APPENDICES	202
BIBLIOGRAPHY	223

LIST OF FIGURES

Figure

1.1	LIGO Hanford Observatory, at top in the plain in this north-looking photograph. Hanford Route 10 passes closest to the corner station joining the arms. Highway 240 goes northwest. The Yakima River flows southwards, joining the Columbia River downstream.	3
1.2	View closer to LIGO Hanford Observatory. Ripples in the plain catch wind. The corner station connects the X and Y arms, 4 kilometers long, which each have a mid- and end-station. The corner station also contains the laser, beam-splitter, and photo-detector of the gravitational wave interferometer.	4
2.1	The Hulse-Taylor binary PSR 1913+16, orbital period change over time consistent with emission of gravitational radiation from its system [172]. Gravitational waves emission depletes orbital energy, causing the binary stars to slowly inspiral into closer orbits. Shrinking pulsar radius R leads to the shorter period T – Kepler’s third law, $T \propto R^{3/2}$ (first stated in <i>Harmonices Mundi</i> [99]). Timing error can be ruled out, because the orbital decay is not linear.	6
2.2	Six GW polarizations are permitted by metric theories of gravity [178]. Here, all are shown propogating in the z -direction (note different axes along left column: the wave travels into the plane of the page for the top three, along the plane of the page for the bottom three). GR permits the top two, (a) and (b); scalar-tensor-vector theory permits also (c); the last three are found in other theories [179]. Conventionally, (a) is labelled h_+ and (b) is h_\times . Shown is the metric perturbation; physically, this corresponds to measured relative motion of test particles. A complete wave cycle is depicted. Off-page (not shown) third- and time-axis are not perturbed. For illustration, the figure shows strain amplitude $h = 0.25$, large enough that the linearized GR model would become inaccurate. This strain is over 20 orders of magnitude larger than the strains LIGO typically expects; LIGO is designed with GR and h_+ and h_\times polarizations in mind.	22
2.3	Hypothetical internal structures of a neutron star. Left: theoretical neutron superfluid vortices and magnetic field lines (citation: Dany Page, http://www.learner.org/courses/physics/visual/visual.html?shortname=a_neutron_star). Right: putative depths of internal layers (citation: http://heasarc.gsfc.nasa.gov/docs/objects/binaries/neutron_star_structure.html).	26

2.4	Enhanced LIGO primary systems and optics. From left, the phase-stablized laser (PSL) produces coherent light, kept at constant frequency by the frequency-stabilization servo (FSS). The pre-mode cleaner excludes non-Gaussian beamshapes, allowing only the TEM_{00} mode to pass. The beam is the phase-modulated using three electro-optical modulators (EOMs) before being further shaped in the input mode cleaner (IMC). The beam passes through the power recycling mirror (RM) and is split at the beam-splitter (BS). Along the X and Y arms, light resonates in the Fabry-Perot cavity formed between the input test mass (ITM) mirror and the end test (mass) ETM mirror, before recombining at the beamsplitting. Any light with the same phase as before the arms is returned to the interferometer by the RM, but if phase-shifted, perhaps by a gravitational wave, it exits through the output mode cleaner (OMC): a photodiode just downstream of the OMC records the signal.	34
2.5	Screenshot of MEDM control panel. MEDM is a Motif Epic and Display Manager for EPICS, the Experimental Physics and Industrial Control System, which lets operators control LIGO. MEDM allows operators to run locking scripts as well as alignments and tests, and to activate and de-activate filters.	35
2.6	Omega scan of an audible broadband glitch. Shortly before the start of Science Run 6, detector characterization took place to identify categories of glitches, such as the ‘gremlin’, and to eliminate them. The burst group analysis pipeline, Omega, generated time-frequency spectrograms that made this identification easier. Glitches are a limiting factor in the identification of rare events and therefore of potential gravitational wave burst signals.	37
2.7	Real-time servo loop diagram. The MICH_CTRL signal can be though to leak into the true DARM signal via a transfer function, H_0. (The letter ‘h’ is typical for transfer functions as well as gravitational wave strain; coincidentally, this is a noise transfer into the channel for displacement, DARM, which is proportional to strain $h(t)$). H_0 leaks into error signal, which is otherwise kept null thanks to the servo cancellation provided by DARM_CTRL times a physical actuation function AC. The measured error signal is DARM_IN1. If desired, an excitation can be supplied via DARM_EXC for a sum of DARM_IN2. This error signal passes though digital filters D to yield the aforementioned control signal DARM_CTRL. The auxiliary length cancellation loop is simply adding H_md to DARM_CTRL in order to subtract out the corruption of H_0.	39
2.8	Real-time work on a LIGO noise filter. This Bode plot, made in Matlab, shows the correction to the existing MICH damping (cancellation) loop needed mid-2010, toward the end of Science Run 6. The correction is small, because the majority of the coupling fits the flat model expected from theory. This transfer function estimate suffices for post-factor correction. However, in order to be incorporated into the control scheme shown in Figure 2.7, measurements of the open loop gain G and actuation function AC are necessary to incorporate the filter correctly into the closed loop response.	40
2.9	Early work on post-facto noise filtering. After testing out MICH damping filters offline, post-facto, the correction factors for AC and G_{closed} were incorporated and the entire filter imported using Foton into the EPICS control system, where it was used real-time from the September equinox of 2010 until the end of Science Run 6.	41
2.10	Optical table layout with Fabry-Perot cavity and Pound-Drever-Hall locking for phase camera. Faraday isolator and polarizing beam-splitter at upper right; piezo-electrically-actuated mirror at center left. Difficulties with stability, despite a plexiglass enclosure and floated table, meant that locks were fractions of a second at most, hampering efforts.	42

3.1	Gravitational wave strain $h(t)$ is derived from differential arm motion (DARM), read-out from a photodiode downstream of the antisymmetric port. An internal reflection off an anti-reflective coating, on either the beam-splitter (BS) or an input test mass (ITM), provides the Michelson (MICH) channel. The DARM read-out channel predominantly measures the small change in differential arm length, $\delta(L_-) \equiv \delta(L_y - L_x)$, while MICH measures that in the Michelson length $\delta(l_-) \equiv \delta(l_y - l_x)$. There is also a small coupling from $\delta(l_-)$ to the DARM channel. To a lesser extent, changes in the length of PRC, which is defined as $\delta(l_+) \equiv \delta(l_y + l_x)/2$ and is measured in quadrature demodulation with respect to the MICH pick-off, also add noise to DARM.	54
3.2	Sample coherence measurements between $h(t)$ and auxiliary control channels for LIGO Hanford Observatory, H1: 2010 March 21. MICH- $h(t)$ coherence on left, PRC- $h(t)$ coherence on right. Statistically significant coherence justifies fitting; in frequency bands, about 80 to 400 Hz, where coherence rose above background levels, the transfer function fit was weighted more heavily. Units of coherence spectral density ($\text{Hz}^{-1/2}$) vs frequency (Hz).	58
3.3	Sample transfer function measurements (amplitude and phase) from LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ on left, PRC- $h(t)$ on right. Transfer function fit in coherent band – note the difference between raw data residual and the ‘pre-processed residual’, which has been smoothed and weighted to emphasize known-coherent bands. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) and phase (degrees) vs frequency (Hz).	59
3.4	Sample Bode plots of fitted ZPK filter functions (amplitude and phase) for multiple 1024 s windows in a science segment, at LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ on left, PRC- $h(t)$ on right. Colors only represent different time windows. The similarity in the high-coherence, 80 to 400 Hz band leads us to conclude that the filter design is fairly stable throughout a science segment. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) and phase (degrees) vs frequency (Hz).	60
3.5	Sample subtracted spectra for one window, representing the applied feedforward corrections for each channel during that window, at LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ correction on left, PRC- $h(t)$ correction (after MICH- $h(t)$ correction is applied) on right. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) vs frequency (Hz).	61
3.6	Calibration line test: before-feedforward mean of the 393.1 Hz line and two neighboring FFT bins was 8.7261×10^{-22} , after was 8.7569×10^{-22} . Feedforward made the calibration line region noisier by 3.1×10^{-24} or 0.35%, suggesting that we correctly apply Hann-windowed feedforward without subtracting true $h(t)$. Moreover, no spectral line combs are observed to either side of the calibration line peak at 393.1 Hz, indicating that the method does not introduce windowing artifacts.	69
3.7	Time-domain plot of diagnostic channels from a burst injection. Colors are illustrative only to the fact that the envelopes of the traces increase after 1.8 s, indicating that the burst injection time is correct in the new data. ‘Before feedforward’ and ‘after feedforward’ traces occult each other in the graph, because they are almost identical. ‘Before feedforward’ is $h(t)$ data; ‘after feedforward’ is $h(t)$ with feedforward subtraction. ‘Injection estimated strain’ is the digital injection as intended to be introduced into strain, but the actual injection is made on the end test mass X (ETMX), so the calibrated ‘Injection estimated ETMX’ is also displayed. Raw ‘DARM_ERR’ and ‘ETMX_EXC_DAQ’ are redundant but reinforce the trend.	70

3.8	Cross-correlation pairwise between $h(t)$ pre-, post-feedforward, and ETMX injection data: the extrema and zero-crossings match. Note both before-feedforward (blue) and after-feedforward (green) strain traces are almost identical and therefore overlap. The strains appear inverted, but in the same way, due to a sign error in the hardware injections at this time. The absence of a time lag shift between before and after indicates that feedforward has not altered the phase of the data, at least for this injection. The equivalence in cross-correlation magnitude indicates that amplitude also is unaffected.	71
3.9	Feedforward subtraction pipeline to read in $h(t)$, MICH, PRC, and write out AMPS-corrected $h(t)$. Data flows schematically from left to right; the MICH- $h(t)$ stage output is used as input for the PRC- $h(t)$ stage, then data is written. Code online: http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/detchar/AMPS/trunk/aletheia.m	72
3.10	Schematic windowing for one LIGO science segment, illustrating windowing after an initial half-window offset. Filters are calculated for windows up to 1024-s, then 50%-overlapping Hann windows merged, giving a corrected measurement of $h(t)$. Code online: http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/detchar/AMPS/trunk/eleutheria.m	74
3.11	Exemplar of a typical case, +1.1 Mpc (5.9% inspiral range) (<i>GPS time 953164819 to 953165839, 2010 March 21</i>). Read ‘DARM’ as $h(t)$, ‘MICH’ as ‘MICH-PRC’. The most benefit is seen in the 80 to 400 Hz band, especially around 150 Hz, where LIGO is most sensitive. The main fundamental limit in this band is thermal suspension noise, but historically auxiliary channel noise has been a major contaminant. Note that the 60 Hz and harmonic lines are not subtracted, although a separate magnetometer servo does reduce their impact. The 340 to 360 Hz violin mode is not likely amenable to feedforward.	76
3.12	Best improvement seen in S6 for H1 $h(t)$, +4.4 Mpc (29% inspiral range) (<i>GPS 955187679 to 955188191, 2010 April 13</i>). Read ‘DARM’ as $h(t)$, ‘MICH’ as ‘MICH-PRC’. Such a loud cross-coupling would be noticed in real-time by the on-site staff. The elevated noise floor is unusual in science mode, but the fact that feedforward corrects it suggests the importance of controlling auxiliary channels to prevent such glitches. The post-feedforward spectrum is comparatively normal for science mode	77
3.13	Harmonic mean, GPS seconds 931.0×10^6 (2009 July 07) to 932.8×10^6 (2009 July 28): (<i>before-after</i>) (L), (<i>before-after</i>)/ <i>before</i> (R); greater than zero is improvement. The mean shows the absolute and relative difference of before and after, between the average of many spectra. Improvement from 80 to 400 Hz is noticeable; at higher frequencies there is degradation, negligible in relative terms, due to high-frequency filter rolloff. Frequencies below 50 Hz should be disregarded; they are usually not searched by LIGO, so spectra for this plot were generated with a high-pass filter at 38 Hz.	78
3.14	Inspiral range vs time for S6 (starting 2009 July 07) before GPS time 9.33e8 (2009 July 30): Each frequency band with reduced noise and consequent enhanced sensitivity contributes to the inspiral range integral, detailed further in Section 4.2.2. LIGO Hanford Observatory, H1 (top) gains 0.23 Mpc; LIGO Livingston Observatory, L1 (bottom) gains 0.84 Mpc. In this first month of S6, L1 saw greater benefit from post-facto feedforward correction; later data from H1 and L1 would improve by fluctuating amounts thanks to better real-time feedforward servos. Although H1 is less improved than L1 here, real-time tunings were made soon after.	79
3.15	Inspiral range vs time for all S6 (2009 July 07 to 2010 October 20): LIGO Hanford Observatory, H1 (top) gains 0.68 Mpc; LIGO Livingston Observatory, L1 (bottom) gains 0.44 Mpc.	80

3.16	Inspirational range <i>fractional gain</i> vs time for S6 (starting 2009 July 07) before GPS time 9.33e8 (2009 July 30): LIGO Hanford Observatory, H1 (top) 1.68% better; LIGO Livingston Observatory, L1 (bottom) 7.00% better. This plot shows relative gain for the same data for we also show absolute gain.	81
3.17	Inspirational range <i>fractional gain</i> vs time for all S6 (2009 July 07 to 2010 October 20): LIGO Hanford Observatory, H1 (top) 4.14% better; LIGO Livingston Observatory, L1 (bottom) 3.60% better.	82
3.18	Screenshot of diagnostic web pages, indexed by window.	83
3.19	DARM power spectral density for first radio-frequency lock of L1 in Advanced LIGO. Note that the spectrum is uncalibrated.	86
3.20	MICH power spectral density, first lock.	86
3.21	PRCL power spectral density, first lock. This corresponds to the PRC channel in Initial and Enhanced LIGO.	87
3.22	SRCL power spectral density, indicating the signal for the signal recycling cavity length.	87
3.23	MICH coherence with DARM.	88
3.24	PRCL coherence with DARM.	89
3.25	SRCL coherence with DARM.	89
3.26	MICH subtraction from DARM, projection, based on frequency domain transfer function. The full AMPS machinery has not yet been applied.	90
3.27	PRCL subtraction from DARM, projection.	90
3.28	SRCL subtraction from DARM, projection.	91
4.1	Diagram of squeezer integration into the LIGO interferometer. Red lines show 1064 nm laser light, green 532 nm; dashed are squeezed (1064 nm) beams. A pump laser, phase-locked to the PSL, generated laser light at 1064 nm, boosted to 532 nm by the second harmonic generator (SHG). This 532 nm beam resonates in the optical parametric oscillator, containing a periodically-poled KTP crystal for non-linear optical effects. The squeezed beam is generated there and exits the OPO, passes through the Faraday isolator and into the interferometer. Simplified and adapted figure [4].	99
4.2	Image by Lisa Barsotti in Hanford eLog. Counterclockwise from lower left: Sheila Dwyer, Lisa Barsotti, Conor Mow-Lowry, Grant Meadors. This photograph shows the uncovered squeezer table with components from the MIT squeezer experiment unpacked at LIGO Hanford Observatory in November 2010. The table sat in a temporary location by HAM6, where it was recommissioned by Dwyer, Mow-Lowry, Sheon Chua, and Alexander Khalaidovksi until H1 could be brought back online for the squeezing experiment in late 2011.	100
4.3	Table legs testing. The squeezer table is raised to its final height by the leg extensions. From top to bottom: table (WISCT10), existing leg extensions (with flanges), new leg extensions (flangeless), triangular high table legs. This assemblage provided the squeezer a serendipitously-stable (as measured by Sheila Dwyer, Sheon Chua, and Robert Schofield) platform at low cost. Photo in temporary location; the actual squeezer table was anchored to these table legs, grouted, by HAM4.	102
4.4	Table legs SolidWorks schematic. This initial design for the table leg extensions on the squeezer table incorporated a flange, which was removed immediately prior to fabrication, replaced with a larger diameter flangeless tube with incorporated tap-holes. Flanges were thought necessary for a flexible alignment initially, they would be prohibitively expensive to machine, and welding would induce unacceptable distortions into the metal. Manufactured by Brockman Machine Works, of Kennewick, Washington. Design in consultation with Daniel Sigg, Lisa Barsotti, Keita Kawabe, Gerardo Moreno, Richard Savage.	103
4.5	Integrand of inspiral range as a function of frequency, with and without squeezing.	106

4.6	Net effect of squeezing on inspiral range integrand. Scientific benefit from squeezing is evident at the few hundred Hz ‘bucket’ where initial LIGO is most sensitive, and although low frequency noise is slightly worse, this in an already noisy spectral band. Enhanced LIGO unambiguously benefitted from squeezing, proving the technique’s efficacy.	107
5.1	After Doppler-shifting the frequencies into the solar system barycenter, TwoSpect analyses begin on this first, time-frequency plane. A simulated signal at 100.015 Hz and $asini = 1.44$ is injected with $h_0 = 4 \times 10^{-21}$ into 10^6 seconds of Gaussian noise at $Sh = 4 \times 10^{-24}$ (the projected minimum Advanced LIGO noise level); the signal period is 68023.8259 seconds, as with Scorpius X-1.	115
5.2	Fourier-transforming along the ‘rows’ (constant frequency bin, variable time bin) generate a second plane, the frequency-frequency plane. The power of each bin in this transform is plotted as a pixel. By aggregating power, this second Fourier transform enhances signal so that matched templates can be applied for a search. The x -axis label is in fact <i>bin number</i> and each value is the 2nd Fourier frequency in Hz times 1800 s.	116
5.3	Exact templates for putative signals weight the pixels in the frequency-frequency plane to generate R statistic for a simulated pulsar (note: not the same as pulsar 40 in the Scorpius X-1 mock data challenge) at 100.015 Hz and $asini = 1.44$. The resulting R values are heatmap-plotted on the modulation depth vs frequency plane.	119
5.4	The Davies algorithm translates R statistic values for exact templates into (single-template) p -values, plotted on the modulation depth vs frequency plane.	120
5.5	All-sky maps, {H1, L1, V1} interferometer analysis from top to bottom, for template tests varying right ascension and declination. Scorpius X-1 mock data challenge pulsar 16 (101x101 templates), showing $\log_{10} p$ -value on a Mollweide projection. Contour lines at 1-radian great-circle distance intervals from the intended injection location of Sco X-1. The results match the intended injection and confirm that the simulation is accurately representing the known sky location of Sco X-1.	123
5.6	Scorpius X-1 Mock Data Challenge (MDC) pulsar 40 {H1}: 5 Hz band. The p -value (single-template, applying Davies’ Method to the R statistic) in is show in this heatmap, peak in red. All templates are plotted on the (frequency, modulation depth) plane. This is a relatively broadband view.	124
5.7	Heatmaps {H1, L1, V1} of 11x11 templates centered around Scorpius X-1 MDC pulsar 8. This is a relatively narrowband view.	125
5.8	Scorpius X-1 MDC pulsar 8 {H1}: 5 Hz band. This heatmap shows 3.6×10^5 templates, 10 to 22 mHz modulation depth, 120-125 Hz frequency. The peak signal at about ($df = 0.019$, $f = 121.9$) Hz.	127
5.9	Scorpius X-1 MDC statistics. These histograms of the R statistic and p -value distribution helped in understanding noise, temporal gap & spectral leakage, as well as establishing a threshold p -value \sim false alarm probability of 1%. These p -values are for single templates, appropriate to the all-sky search but not to a dense templated search with a large trials factor. Here, histograms show statistics in the absence of a signal. The left-skew of the p -values is associated with gaps in the data (as is right-skew, with different gaps). Diurnal bias in barycentering is assumed, but the correlation is not fully understood.	128
5.10	Error in strain estimation versus circular-effective injected strain. Higher injected strain results in higher absolute errors.	132
5.11	Parameter estimation: error in strain and dependence on recovered p -value (top) and frequency (bottom). The strain appears broadly distributed, without any systematic patterns. The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies. .	134

5.12	Parameter estimation: error in frequency and dependence on recovered p -value (top) and frequency (bottom). The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies.	135
5.13	Parameter estimation: $a \sin \iota$ (projected semi-major axis; directly proportional to modulation depth for a given frequency and period) and dependence on recovered p -value (top) and frequency (bottom). The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies.	136
5.14	Open pulsar detection efficiency curve. Because only 50 pulsars were in the open set, this curve is relatively-poorly defined – the binning has been chosen to give the most accurate representation based on the chosen thresholds. The 95% level is approximately about 3×10^{-25} (again, without the corrective factors of 1.74 and 1.11) but is too imprecise to judge using this method.	137
5.15	Detections and upper limit determination. Depending on whether an injection was seen in three, one, or no detector pairs, it was assigned a color-coded circle and plotted in recovered strain versus effective circular strain injected. (There are no injections seen with two detection pairs, because this plot only shows the loudest outlier from each 5 Hz band; if some injection were seen in two and not three pairs, it would mean two distinct coincidences were seen, only one of which would be the loudest). Color-coding red pulsars as non-detected, blue as single pairwise detection, and green as triple pairwise detection, we identified a shelf of non-detected pulsars that was 95% contained by an upper limit about 2.19×10^{-25} . This number, when corrected, yielded the upper limit of $1.74 \times 1.11 \times 2.19 \times 10^{-25} = 4.23 \times 10^{-25}$ for TwoSpect. The unity-slope line is shown to ascertain whether a further empirical rescaling factor was needed (it was: constant 1.11). The zero-slope line is shown to indicate the ninety-five percent confidence upper limit in the absence of detection.	138
5.16	Simulated detection efficiency curve. Because the $\cos \iota$ ambiguity simulation requires a priori model of detection efficiency, we described it simply. Here, no detections were claimed below 1×10^{-25} , all were detected above 3×10^{-25} , and the probability of detection rose uniformly on the intervening interval.	140
5.17	Distribution of 2 million simulated stars, strain between 3×10^{-26} and 3×10^{-24} under a log-uniform distribution, following application of $\cos \iota$ and detection efficiency cuts.	141
5.18	Regression using grid. By binning the simulated stars on the true strain vs detected (recovered) strain plane, an accurate mean slope for the $\cos \iota$ correction was ascertained. It had to be modified downwards by the equivalent of one bin, to 1.74. However suggestive, the $1\text{-}\sigma$ thresholds proved inaccurate, probably due to noise fluctuations.	143
5.19	Simulation with fit lines as given by the bin-method regression.	144
5.20	Confidence intervals with final fit. After manual optimization of the cumulative distribution function, and constraint to the region with a full bijective mapping between injected and recovered strains (below 1×10^{-24}), a $1\text{-}\sigma$ value of 0.37 in the slope was found to give accurate confidence intervals. The reason for the aforementioned restriction of the plot to $h_0\text{-effective} < 1 \times 10^{-24}$ can be seen in the distortion at levels above that. The chosen $1.74 \pm 0.37\sigma$, however, yielded the necessary correction factor.	145
6.1	Detection efficiency of 500 injections (each at H1, L1) into S6 data at 142 Hz, given threshold $\log_{10} p = -7.75$. The least-squares curve fit is to a symmetric sigmoid, which matches the data well at high detection efficiency. Since the 95% efficiency region is the most interesting, the relatively poor fit at low efficiency is not much concern, although the fit could be improved with additional parameters.	160
6.2	Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 142 Hz (injections also done at 162, 222 Hz)	161

6.3	S6 h_0 heatmap shows real data features, such as 46.7 Hz cal line	164
6.4	Joint 95% upper limits for Scorpius X-1 in S6 data, using a confidence interval given by reported $h_0 \times (1 + 1.0) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. This spectrum covers 40 to 2040 Hz using the lower upper limit from either interferometer (H1 or L1) when both yielded data. A total of 28.8 Hz were in bands that yielded no real upper limit (because the quarter root of the test statistic was imaginary) in either interferometer, generally due to excessive noise in that band. Bands were left-closed and right open, e.g., $[40.0, 40.1)$, $[40.1, 40.2)$. . . $[2039.9, 2040.0)$. These results are promising, given that the Advanced LIGO noise floor is expected to be up to 5 times lower at the most sensitive current frequencies of a few hundred Hz. Advanced LIGO will also make even larger relative improvements from 10 to 40 Hz, where the torque-balance equation allows higher GW strain than at high frequencies. Together with recent and anticipated improvements in TwoSpect analyses, order-of-magnitude better upper limits should be anticipated for full-sensitivity Advanced LIGO results.	165
6.5	Quick look at J1751-305, H1 $\log_{10} p$, 435 Hz ν_0 rotation frequency.	170
6.6	Quick look at J1751-305, L1 $\log_{10} p$, 435 Hz ν_0 rotation frequency.	171
6.7	Quick look at J1751-305, H1 $\log_{10} p$, 621 Hz r -mode.	172
6.8	Quick look at J1751-305, L1 $\log_{10} p$, 621 Hz r -mode.	173
6.9	Quick look at J1751-305, H1 $\log_{10} p$, 870 Hz $2\nu_0$, twice rotation frequency.	174
6.10	Quick look at J1751-305, L1 $\log_{10} p$, 870 Hz $2\nu_0$, twice rotation frequency. The dominant feature is believed to be an instrumental line.	175
7.1	Kiosk interferometer schematic	182
7.2	Interferometer fringes during construction in Ann Arbor	183
7.3	AutoCAD corner view of plexiglass and aluminum interferometer enclosure. Plexiglass in light gray, aluminum in dark gray. Plexiglass is attached to aluminum with small plexiglass blocks (a refinement to the design), placed three to the long Y -arm, two to the short X -arm, and three to the corner station. The triangular prism is not anchored to the aluminum block. It is open to the bottom so that a projection screen can be inserted, typically atop a flexible rubber sheet. In the World Science Festival exhibit, this projection screen included a photodiode and attached signal cable.	187
7.4	Front view of interferometer aluminum and plexiglass enclosure, with dimensions.	187
7.5	Dimensions and hole locations of the aluminum plate (dark gray) relative to plexiglass (light gray), with proposed hole locations given by Table 7.2. These hole locations were as tapped.	188
7.6	Interferometer assembly in Ann Arbor, Michigan. Projection screen not yet installed. Note the outer end of the Y -arm (bottom left of photo), with perforated-and-covered plexiglass to allow sound passage. Humans (<i>e.g.</i> , Evan Goetz at left, Keith Riles at right) can be protected during alignment processes by laser safety curtains, seen in back.	189
7.7	World Science Festival interferometer installed in the New York City exhibition hall, June 2010. Optics aligned and baffles installed to protect visitors from scattered light. Also visible: left center, an initial LIGO input mirror and suspension.	190
7.8	Helping to host the exhibit at the World Science Festival. Optics aligned, with fringe pattern, but an edge of the beam-splitter reflection is visible in the projected image. Additional features: the reflected blue lights of the sculpture above and posters of the exhibit.	193
8.1	LIGO Hanford Observatory sunset, initial detector era. Photo by author. Like the Hanford desert wiped clean by the Missoula floods, the gravitational wave sky may relate cosmic tales of cataclysm and rebirth in the distant past.	200
B.1	Detection efficiency of 500 injections (each at H1, L1) into S6 data at 162 Hz, given threshold $\log_{10} p = -7.75$	210

B.2	Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 162 Hz (injections also done at 142, 222 Hz)	211
B.3	Detection efficiency of 500 injections (each at H1, L1) into S6 data at 222 Hz, given threshold $\log_{10} p = -7.75$	212
B.4	Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 222 Hz (injections also done at 142, 162 Hz)	213
B.5	H1: loudest h_0 in 0.1 Hz bands, effectively circular polarization	214
B.6	L1: loudest h_0 in 0.1 Hz bands, effectively circular polarization	215
C.1	H1: loudest $h_0 \times (1 + 0.8) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. The ‘Tentative h_0 95 percent UL’ statement was later revised; the 95% confidence level is higher by 2.0/1.8. This plot shown corresponds to about 90% confidence in random polarization ULs in the final analysis.	218
C.2	L1: loudest $h_0 \times (1 + 0.8) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. The ‘Tentative h_0 95 percent UL’ statement was later revised; the 95% confidence level is higher by 2.0/1.8. This plot shown corresponds to about 90% confidence in random polarization ULs in the final analysis.	219

LIST OF TABLES

Table

6.1	List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 40 to 360 Hz.	167
7.1	Bill of materials (lasers, lenses, mirrors, optic mounts and screws) for interferometer assembly. Manufacturers: ¹ CVI Melles-Griot, ² Edmund Optics, ³ OptoSigma, ⁴ ThorLabs	181
7.2	Hole locations (in inches from origin) for the WSF interferometer aluminum base-board, plotted with suggested alterations on Figure 7.5.	186
C.1	List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 360 to 900 Hz.	220
C.2	List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 900 to 1700 Hz.	221
C.3	List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 1700 to 2040 Hz.	222

LIST OF APPENDICES

Appendix

A.	Fields and curvature	203
	A.1 Mathematical conventions	203
	A.2 Electrodynamics	205
B.	Scorpius X-1 real data detection efficiency	209
	B.1 Real S6 data: detection efficiency at 162 Hz	209
	B.1.1 Real S6 data: h_0 recovered vs injected	211
	B.2 Real data: detection efficiency at 222 Hz	212
	B.2.1 Real S6 data: h_0 recovered vs injected	213
	B.3 S6: Scorpius X-1 upper limits, raw circular output	216
C.	Scorpius X-1 further preliminary upper limits and outliers	217
	C.1 Upper limits at 90% confidence levels	217
	C.2 Preliminary high frequency outliers	217

CHAPTER I

Introduction

Space should reverberate with gravitational waves. Light shows part of cosmic history; now, primeval epochs and secret stellar reaches might be seen in patterns of light transformed by gravity. General Relativity and related theories of gravitation posit [76] that changing quadrupolar masses radiate gravitationally, just as accelerating dipolar charges do electromagnetically. In those waves we might see black holes and neutron stars colliding, supernovae, the dawn of the Big Bang and rotating neutron stars – and the potential for unanticipated insights, into other objects or laws of physics, is too tantalizing to ignore. As yet, we have made no direct detections. Hulse and Taylor [103] observed a neutron star in a binary system, PSR 1913+16, with an orbit shrinking just as gravitational radiation would predict.

Following on the pioneering work of Joseph Weber with bar detectors [171] and Robert Forward with tabletop interferometers [82], kilometer-scale interferometers were built at the end of the last millenium to look for gravitational radiation. Laser light in these instruments travels orthogonal paths and is reflected back; shifts in the combined pattern are scrutinized for indications that gravitational waves stretched space itself. LIGO (the Laser Interferometer Gravitational-wave Observatory) [21,84], Virgo [27], and GEO600 [180,100], soon to be joined by KAGRA [112],

are kilometer-scale interferometers, gravitational wave antennae standing on the threshold of discovery. This thesis includes analysis of LIGO data.

Vibrations in spacetime's metric require many steps to detect. The author's work has pursued a series of clearer perspectives on detection: filtering & regressing out correlated noise, helping cancel fluctuations in the electromagnetic field with quantum optics, and looking for continuous waves from promising neutron stars in binary systems. Chapter II describes General Relativity's prediction of gravitational waves and the design & operation of LIGO. Noise intrinsic to the optical configuration of these instruments is subtracted post-facto by feedforward filtering using recorded servo data in Chapter III. Quantum optical squeezing reduces relevant uncertainties per Heisenberg's principle in Chapter IV. Then the search begins. Astrophysicists expect to find signals from four categories of cosmic sources: inspiralling binary systems of stellar remnants, supernovae and similar bursts, stochastic background, and continuous waves from neutron stars. Einstein's theory predicts the intensity, speed, and polarization of gravitational waves that could be emitted from these sources. Low-mass X-ray binaries (LMXBs) should lead astronomically long lifetimes, radiating continuous waves from their constituent neutron stars. Chapters V and VI use these expectations to enhance and run Fourier-domain analyses of simulated and real data. These searches target the LMXB Scorpius X-1 and X-ray transient J1751-305. Astronomy has grown from humanity's first glimpses into the night sky with the unaided eye. With every new instrument, from Galileo's telescope through radio antennae and neutrino catchers, our understanding of the cosmos has grown. Communicating that understanding is the subject of Chapter VII. Gravity pervades the universe like no other force: we must hear its tale.



Figure 1.1: LIGO Hanford Observatory, at top in the plain in this north-looking photograph. Hanford Route 10 passes closest to the corner station joining the arms. Highway 240 goes northwest. The Yakima River flows southwards, joining the Columbia River downstream.



Figure 1.2: View closer to LIGO Hanford Observatory. Ripples in the plain catch wind. The corner station connects the X and Y arms, 4 kilometers long, which each have a mid- and end-station. The corner station also contains the laser, beam-splitter, and photo-detector of the gravitational wave interferometer.

CHAPTER II

Gravitational Waves and LIGO

2.1 Cosmic sources of gravitational waves

Gravity's power induces ripples in space. Pulsar 1913+16, discovered by Hulse and Taylor in radio waves, not only followed a pattern of orbital decay consistent with radiative energy loss to gravitational radiation – it continued to do so [173,172], as seen in Figure 2.1, after the 1993 physics Nobel Prize. This year, there has been much debate as to whether or not the BICEP2 [29] and Planck [30] probes of the cosmic microwave background have seen evidence of B -mode polarizations that would indicate primordial gravitational fluctuations due to inflation. Eventual identification of the polarization is expected by many, regardless [45]. We still may ask whether any gravitational waves will be directly detectable on Earth. We may ask whether they appear in detectors in a way consistent with General Relativity. The basic fact of their emission, however, appears settled. Gravity has the power to make waves.

Before delving into general relativistic emission, let us consider the astrophysical sources expected to emit gravitational waves. Physics prompts our search, but astronomy enriches it. When gravitational waves are heard by the interferometers, we will be hearing the songs of dead stars, rippling through the fabric of spacetime. Gravitational wave (henceforth also abbreviated GW) searches presently focus on




Figure 2.1: The Hulse-Taylor binary PSR 1913+16, orbital period change over time consistent with emission of gravitational radiation from its system [172]. Gravitational waves emission depletes orbital energy, causing the binary stars to slowly inspiral into closer orbits. Shrinking pulsar radius R leads to the shorter period T – Kepler’s third law, $T \propto R^{3/2}$ (first stated in *Harmonices Mundi* [99]). Timing error can be ruled out, because the orbital decay is not linear.

four distinct types of cosmic sources. This categorization of sources was first presented no later than the 1983 LIGO Blue Book proposal [63,151] and has since guided research focus. The quadripartite division:

- Burst (rarely, ‘supernova’)
- Compact binary coalescence (or ‘inspiral’)
- Continuous wave (or ‘pulsar’)
- Stochastic

This thesis concentrates on continuous waves (CWs) – sine waves. CWs are most likely to emanate from neutron stars. Neutron star CWs can be modulated by orbital motion, spun-up from accretion or spun-down from radiated energy. Given a sufficiently large ellipticity, $\epsilon \approx 10^{-7}$ [136] to within an order of magnitude for a neutron star rotating on the order of 1 kHz, a crust deformation (alternatively an *r*-mode [138,137]) would radiate sufficient gravitational radiation to be a plausibly detectable source. In an isolated neutron star, GW radiation would rapidly deplete rotational energy [138]. Binary systems, where the neutron star could be recycled and spun-up by a partner [139,168], would last longer in contrast. Scorpius X-1 offers a canonical case [24], although TwoSpect (a search for neutron stars in binary systems, detailed in this thesis) anticipates an abundance of other low-mass X-ray binary (LMXB) systems of interest. Given the paucity of insight on the interiors of collapsed stellar remnants, direct detection of GWs from neutron stars would prove informative [117]. Just as we might infer details from neutron star binary coalescences favoring one equation of state [113,145], we might also extract parameters from continuous waves suggesting the existence of quark stars or gravitars [136], and will have an unparalleled peek into the interior of the densest stable three-dimensional

objects in the universe. Their simple waveforms might even facilitate the calibration of other types of GW data, whereas binary mergers should sound like standard sirens to compare against electromagnetic and neutrino observations [142] (see next paragraph). CWs are conceptually elegant and astronomically enticing. Yet other sources of GWs have a comparable pull on our attention.

Inspirals or compact binary coalescences occur when two stellar remnants draw nearer in their orbits, radiating gravitational radiation and finally merging in a titanic release of energy. While sometimes invisible – except as short-hard gamma ray bursts (GRBs) – these events compete with supernovae as the most explosive in the modern universe. Were GW observatories to see their waveforms, they could be compared with those predicted through post-Newtonian approximation and numerical relativity. As GW amplitude should diminish inversely with distance, we would then have standard candles or *standard sirens* by which to calibrate and measure the universe. Advanced LIGO [33, 86] may prove sensitive to neutron star-neutron star and stellar mass black hole-neutron star mergers, and, if low-frequency sensitivity is sufficient and the sources exist, to intermediate-mass black holes. Proposed space-based observatories such as the Laser Interferometer Space Antenna (LISA) and ensuing DeciHertz Gravitational-wave Observatory (DECIGO) & Big Bang Observer (BBO) could detect supermassive black hole mergers. If launched, they would see a low-frequency noise floor due not to seismic vibration, as in LIGO, but to white dwarf binaries throughout the galaxy. Since the waveforms are well-predicted, we could even investigate deviations from General Relativity, perhaps seeing new physics in the ringdown of highly massive black holes.

Physical insight could also come from burst searches. Burst searches share with inspiral searches the property of looking for a single transient event, as opposed to

a source spread over long duration. Analytical programs for bursts can sometimes be applied to inspiral or detector characterization tasks as well; rare noise events are a limit on their astronomical sensitivity. Yet the immediate focus lies with supernovae [60, 135] and perhaps gamma-ray bursts. Because the waveform is unknown, burst searches rely significantly more on the coincidence between multiple detectors to distinguish signal from noise. Just as with neutrino observations of supernova 1987A, the burst program would hope for a fortuitously nearby cataclysm to be seen simultaneously – or nearly so, the time of flight indicating a direction – by multiple detectors in a global gravitational-wave detector network. Due to the versatility of this method, some researchers have proposed looking for longitudinal polarization in addition to plus and cross orthogonal polarization (possible in non-general relativistic terms; see Section 2.1.1). Any detection would be quite exciting for probing still mysterious systems with electromagnetic and neutrino measurements, and it would help, in conjunction with multi-messenger coordinated searches with those observatories, to ascertain at precisely what speed gravity travels through space-time and to what extent it is attenuated or altered.

The background of space-time itself may hide GW signatures. Searches for the stochastic GW background look not for single events but for persistent phenomena buried in many months of correlated signals between networks of detectors. In doing so, they hope in particular to see the earliest turbulence of the universe. Long before the cosmic electromagnetic background was emitted 380000 years after the Big Bang – now redshifted into microwaves – GWs were travelling unimpeded. While the opacity of the infant cosmos conflates electromagnetic signals from different times and places, the transparency of the universe to gravity means that we might see the inflationary epoch. Unfortunately, this signal is thought to be far below the

sensitivity of existing detectors in the LIGO band. While LIGO did set a new upper limit on the energy density of GWs, measured as a fraction, Ω_{gw} of the critical closure density of the universe [163], the inflationary background at LIGO frequencies is predicted to be about ten orders of magnitude lower. Alternative theories, such as ekpyrotic/cyclic universes or an axionic inflaton, make other predictions; an anomalously high stochastic background could thus prove cosmologically significant.

All GW searches strive to open up new directions in astronomy. While the most exciting possibility is that we will see the unexpected, we think that our present algorithms will permit serendipity. Continuous wave and inspiral methods both search against waveform templates; burst and stochastic have no template and rely on correlation and coincidence. Continuous wave and stochastic searches analyze weeks, months, even years of data in search of persistent features; inspiral and burst searches look for transient events. In the abstract dimensions of search groups, we are complete. Our attention is narrow now – in the focus on audio frequencies of tens to a few thousand Hertz at present – narrowness that will in time be broadened by CMB polarization, pulsar-timing and space-based interferometry for low frequencies and possibly by atom interferometry [69] around 1 to 10 Hz. To appreciate our choice of focus in these nascent days of the field, we must turn back a century to understand its origins in Einstein’s mathematics.

2.1.1 History from General Relativity

Einstein’s theory unified a sequence of historical insights. Since 1676, when Rømer used the moons of Jupiter to measure the finite speed of light, just before Newton’s 1687 *Principia Mathematica* [99], the question of gravity’s propagation beckoned. Bringing together the work of Minkowski and Poincaré, the 1905 special theory of relativity highlighted the universality of the speed of light, but only in 1915,

with the presentation of the Einstein field equations of General Relativity, based in Riemannian geometry, did a means to an answer emerge. In 1916, Einstein predicted GWs. At last, gravity had a theoretical speed: the same as for electromagnetism, that of light in the vacuum, c . In the linear approximation to the nonlinear theory of General Relativity (henceforth GR), the waves were mathematically similar to the waves of electromagnetism, as will be shown in Section 2.2. GR offers a consistent explanation for how changes in the distribution of matter change gravitational fields. Yet the detectability of the waves, even in principle, would remain an open question for another half century. Uncertainty in whether fluctuations within spacetime could be detectable with instruments themselves changed by the same fluctuations dominated the debate. Thought notions, *Gedankenexperiment*, such as beads-on-rods led to consensus that GWs carried and could deposit physical energy and thus be detected. Consult Misner, Thorne, and Wheeler’s *Gravitation* [131] and Sean Carroll’s lectures notes [55], as well as other history books of the field for an account of the controversy. Saulson [150] explicates both why gravitational waves should be detectable in theory and why travel time is superior to wavelength in understanding GW interferometry. See *Gravity’s Shadow* [63] for sociological perspective, and *Gravity’s Ghost* [64] for insight into the detection criteria that have since evolved¹.

Discussions of GWs frequently begin with derivations of the wave equations from Einstein’s field equations. General relativity, however, is not the only theory to predict GWs: waves are a natural consequence of a class of similar theories, which make a range of testable predictions (such as number of polarization modes, from two to six, and possibly speeds different from c) [178]. Waves are expected even if some small variation from Einstein’s theory is discovered, pointing a way perhaps toward a

¹See also a recent review by Riles [146].

quantum theory of gravity [148]. Eventual observation of some form of gravitational waves is highly probable, given our knowledge of the Hulse-Taylor binary.

Before deriving the answer to the detectability question from the field equations, a contrast with the situation in other fields of astronomy is in order.

2.1.2 Contrast with electromagnetic and particle astronomy

With astronomy, detection came first, then theory. Visible light astronomy began with the earliest humans. Records of the star Sirius are known from Egyptian astronomers, the planet Venus from Babylonians, sunspots from the Chinese, and eclipses from the Greeks. Thus the telescope, though revolutionary, was not unimaginable.

Infrared radiation, as first seen by William Herschel, was just beyond the visible red light of a prism. As the wave nature of light came to be understood, culminating in Maxwell's equations, the existence of invisible electromagnetic radiation was put on sound theoretical footing. The only remaining questions pertained to whether this radiation would prove astronomically interesting, especially in the extreme low- (radio) and high- (X- and γ -ray) frequencies found at the end of the 19th Century. Fortuitously, unlike with GWs today, these novel bands of electromagnetic spectrum could be easily generated and detected by scientists in small laboratories, as with Marconi and Marie & Pierre Curie. Within half a century, radio astronomy began with Karl Jansky [153], with Grote Reber soon observing the Milky Way [94]. By the early 21st Century, radio astronomy ranged from common rooftop designs [1] (nonetheless sensitive enough to infer the galactic rotation curve in the hydrogen line [126]) to planetary-scale very long baseline interferometry and plans for a Square Kilometer Array. Electromagnetic, or photon, astronomy has the advantage of calibration with familiar sources.

Other particles besides the photon began to play a role in 20th Century astronomy. Muons from space were detected long ago, as in C.D. Anderson's 1949 paper on what was then called the mesotron [115]. Solar neutrinos were, after much difficulty, seen by Ray Davis [41], and neutrino astronomy was a well established field by the turn of the millenium, following the detection of Supernova 1987A in the Large Magellanic Cloud. New types of neutrino detectors² could yield additional information and better sensitivity. As with electromagnetic waves, humans can generate a measurable neutrino flux. Since the Savannah River nuclear reactor experiments, the experimental detectability of neutrinos has been settled. Solar neutrino observations spurred a deeper understanding of neutrino oscillation to coincide with theory [40]; now, astrophysics with neutrino observatories is becoming mature. Nuclear fusion in the Sun has been well-established, and arrays such as Antares [32] and IceCube [2] will move the field outwards toward cosmic sources.

GWs from the PSR 1913+16 system itself [173] are too low in frequency to be seen with existing interferometers, but they confirm that the phenomenon is real. Although no terrestrial sources of GWs can be feasibly generated, calibration [10] of interferometers is nonetheless accurate to a few percent. Direct detection of GWs would let us infer the true strength of astrophysical gravitational radiation and compare that power with theory. Surprising new insights are probable in any new field, not unlike the discovery of neutrino oscillation in the solar spectrum or the anomalous galactic rotation curve due to dark matter.

Yet how can gravity be seen? GWs are conceptually simple, with mathematics illustrated by electrodynamic analogy in Appendix A. A short derivation follows.

²For example, time-projection chambers using electron bubbles in cryogenic noble gas such as helium or neon, the subject of some prototype construction and simulation research by the author [46, 124].

2.2 General relativity

Einstein's theory of General Relativity (hence GR) describes gravitation as curvature. Gravity is a pseudoforce, unlike Newton's force $\mathbf{F} = -G_C M_1 M_2 r^{-2} \hat{\mathbf{r}}$ for two bodies of masses M_1 and M_2 separated by a vector distance \mathbf{r} in a universe with a gravitational constant G_C (henceforth, units are set where $G_C \equiv 1$). Einstein posits that objects always follow an extremized path, $\delta s = 0$ for arclength s – a geodesic worldline – when the curvature of spacetime is considered.

The curvature of spacetime is described by the Riemann tensor, $R^\mu_{\nu\rho\sigma}$, which physically is required to be a solution to the Einstein field equations. Section 2.2.2 clarifies the connection between these field equations and the Ricci scalar curvature R , a trace of the Riemann tensor, whence they can be derived. This Riemann tensor is constructed from the Christoffel connection, $\Gamma^\mu_{\rho\sigma}$, which in turn can be expressed in terms of the metric, $g_{\mu\nu}$. That metric measures distance across spacetime: $ds^2 = g_{\mu\nu} dx^\mu dx^\nu$. Hence gravity is a pseudoforce in GR: curvature distortions (arising from matter) change the distance along worldlines and thus affect trajectories.

For more detail, please see Appendix A. Carroll and *Gravitation* are useful references [55, 131], consulting others for the Einstein-Hilbert/Palatini action [78].

2.2.1 Symmetry and action principles

Einstein's geodesic equation is a replacement for not only Newton's law of gravitation but also for the Laws of Motion. An object's acceleration is described in terms of its proper time and the Christoffel connection. The supposition that objects follow a geodesic path follows from Fermat's principle of least time, which also inspired the least action principle, $\delta\mathcal{S} = 0$. In turn, interferometry is best understood in terms of least-time phase or *time-of-flight*.

Gauge symmetries and the associated conservation laws of Noether’s theorem have been critical to the development of quantum field theory. Whereas electromagnetism arises from $U(1)$ unitary Lie group gauge symmetries, the electroweak force from $SU(2) \times U(1)$, and quantum chromodynamics from $SU(3)$, GR arises from $GL(4, \mathbb{R})$ general linear Lie group symmetries. These symmetries are *diffeomorphism invariances* with respect to Lorentz boosts and translations, meaning we can *pullback* vector fields, such as the other forces, through a coordinate change without changing their physics. Ergo the conservation laws that emerge from GR are the familiar ones from mechanics – stress-energy, momentum, and so on, with subtleties (particular for non-static universes).

2.2.2 Derivation of field equations

A common approach to GW derivations is linearized GR [81]. Linearization proceeds [31] from the metric $g_{\mu\nu}$ approximation as the sum of a Minkowski (flat space-time) component, $\eta_{\mu\nu}$, plus a perturbation h , where $|h_{\mu\nu}| \ll |\eta_{\mu\nu}|$, so $|h_{\mu\nu}| \ll 1$:

$$(2.1) \quad g_{\mu\nu} \approx \eta_{\mu\nu} + h_{\mu\nu},$$

the weak-field approximation, leading to a wave equation, notated by a d’Alembertian $\square = \partial_\mu \partial^\mu$,

$$(2.2) \quad -\frac{1}{2}\square \left(h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h^\lambda{}_\lambda \right) = \frac{8\pi G_C}{c^4}T_{\mu\nu}.$$

The right-hand side indicates that GWs are sourced by the stress-energy tensor, $T_{\mu\nu}$, which subsumes the electromagnetic stress-energy, pressure, and density terms. For gravitational sources of astrophysical interest, the majority of sourcing comes

from matter density – since mass-energy is conserved, monopolar radiation is forbidden (formally, by the Birkhoff theorem), and conservation of linear and angular momentum forbids the dipole terms. GWs are thus sourced by quadrupoles I_{ij} and higher only, although the higher-order terms are generally smaller and neglected. The wave does carry energy density ρ_{gw} and power P [42]:

$$(2.3) \quad \rho_{gw} = \frac{c^2}{16\pi G_C} \langle |h_{+,0}|^2 + |h_{\times,0}|^2 \rangle,$$

$$(2.4) \quad P = \frac{G_C}{5c^5} \langle (\partial_t^3 I_{ij})^2 \rangle \approx 5.5 \times 10^{-54} W^{-1} \langle (\partial_t^3 I_{ij})^2 \rangle.$$

Equations 2.1-2.4 represent the testable predictions of GR about GWs. GR itself is found by finding the field equations from the action [78], optimizing (extremizing) the integral of curvature: the Ricci scalar R . Here $|g|$ indicates the determinant of g . Add two terms to R : a hypothetical constant, Λ (termed the cosmological constant), and the matter Lagrangian density \mathcal{L}_M . The sum generates the Einstein-Hilbert action \mathcal{S} :

$$(2.5) \quad \mathcal{S} = \int \left(\frac{1}{8\pi} (R - 2\Lambda) + \mathcal{L}_M \right) \sqrt{-|g|} d^4x,$$

Extremizing $\delta\mathcal{S} = 0$ involves expanding $R = R^\mu{}_\mu$, where $R_{\mu\nu} = R^\lambda{}_{\mu\lambda\nu}$, the Ricci tensor. In turn, the Ricci tensor is contracted from the Riemann tensor, $R^\mu{}_{\nu\rho\sigma}$ (the curvature tensor, equivalently the 2-form $R^\mu{}_\nu$). Maurer-Cartan structure equations show how the Riemann tensor is an operator, the sum of a ‘differential’ plus connection ω , applied to the connection ω itself,

$$(2.6) \quad R_\nu^\mu = d\omega_\nu^\mu + \omega_\rho^\mu \wedge \omega_\nu^\rho,$$

$$(2.7) \quad \omega_{\mu\sigma'}^{\rho'} = e_\nu^{\alpha'} e_{\beta'}^\lambda \Gamma_{\mu\lambda}^\nu - e_{\beta'}^\lambda \partial_\mu e_\lambda^{\alpha'},$$

$$(2.8) \quad \Gamma_{\mu\nu}^\sigma = \frac{1}{2} g^{\sigma\rho} (g_{\nu\rho,\mu} + g_{\rho\mu,\nu} - g_{\mu\nu,\rho}).$$

where the last equation defines the *Christoffel* connection. Just as metric induces distance on the spacetime Riemannian manifold, and the connection corrects geodesic paths, so the curvature tensor accounts for geodesic deviation between paths.

Substituting into the action, we obtain [55] the *Einstein field equations* by varying over the metric $g_{\mu\nu}$, we make use of the fact that $\delta R = (g^{\mu\nu} \delta R_{\mu\nu} + R_{\mu\nu} \delta g^{\mu\nu})$ (the first term of which vanishes due to Stokes' theorem). The equations are often simplified using the Einstein tensor, $G_{\mu\nu} \equiv R_{\mu\nu} - (1/2)Rg_{\mu\nu}$:

$$(2.9) \quad \delta\mathcal{S} = \int d^4x \left(\frac{\delta R \sqrt{-|g|} + (R - 2\Lambda)\delta\sqrt{-|g|}}{8\pi} + \delta(\sqrt{-|g|}\mathcal{L}_M) \right),$$

$$(2.10) \quad = \int d^4x \sqrt{-|g|} \left(\frac{\delta g^{\mu\nu}}{8\pi} [G_{\mu\nu} + \Lambda g_{\mu\nu}] + \frac{\delta g^{\mu\nu}}{\sqrt{-|g|}} \frac{\delta(\sqrt{-|g|}\mathcal{L}_M)}{\delta g_{\mu\nu}} \right).$$

Since the integral must vanish everywhere,

$$(2.11) \quad \frac{1}{8\pi} [G_{\mu\nu} + \Lambda g_{\mu\nu}] = -\frac{1}{\sqrt{-|g|}} \frac{\delta(\sqrt{-|g|}\mathcal{L}_M)}{\delta g_{\mu\nu}}.$$

The field equations are sourced by the right-hand side, identified (e.g., by comparison with Newtonian gravity) with the stress-energy tensor, $T_{\mu\nu}$:

$$(2.12) \quad G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi T_{\mu\nu}.$$

Einstein's concern was the different nature of the source term. Much of the effort toward unifying GR with the other forces can be understood as trying to fuse the matter Lagrangian with the Ricci scalar in an intuitive way.

The linearized wave derivation in transverse-traceless gauge is standard in GWs. Radiation in the non-linear near-field is explored by numerical relativistic studies [78]. Equation 2.12, however, also allows derivations using the curvature. This derivation makes it easier to compare interferometer observations (of length scale L much smaller than GW wavelength λ) with less-linear spacetimes or where L is comparable to λ . In vacuum, the Riemann tensor satisfies a non-linear covariant wave equation in any gauge [111], $0 = R_{\mu\nu\rho\sigma;\lambda}^\lambda + S_{\mu\nu\rho\sigma}$, where the $S_{\mu\nu\rho\sigma}$ is quadratic in the Riemann tensor and therefore negligible. The first-order perturbative response of the Riemann tensor in Minkowski space of an observer at rest, in conventional transverse-traceless gauge (see below), to a metric perturbation $\bar{h}_{\mu\nu}$ is then,

$$(2.13) \quad R_{0\mu 0\nu} = -\frac{1}{2}\bar{h}_{\mu\nu,00}.$$

Obtaining this first-order response recovers the common formalism [42], where a metric perturbation is substituted into Equation 2.12 to yield,

$$(2.14) \quad G_{\mu\nu} = -\frac{1}{2} (h_{\mu\nu,\lambda}^\lambda - h_{\mu,\lambda\nu}^\lambda - h_{\nu,\lambda\mu}^\lambda + \eta_{\mu\nu}(h^{\lambda\sigma})_{,\lambda\sigma} - \eta_{\mu\nu}(h_\sigma^\sigma)_{,\lambda\lambda} - (h_\sigma^\sigma)_{,\mu\nu}),$$

simplified by switching coordinates to a traceless gauge, defined by $\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - (1/2)\eta_{\mu\nu}h_\lambda^\lambda$, that obeys the harmonic (transverse) condition, $\bar{h}_{\mu\nu} = 0$, into

$$(2.15) \quad G_{\mu\nu} = -\frac{1}{2}\bar{h}_{\mu\nu,00} = 8\pi T_{\mu\nu}.$$

Equation 2.15 is the origin of Equation 2.2 and also corresponds to Equation 2.13. Moreover, outside a source, we can impose the transverse-traceless gauge, which verifies that in vacuum $\bar{h}_{\mu\nu} = h_{\mu\nu}$. The curvature wave equation is what is fundamental.

However described, the wave, parametrized by path $x^\mu(\lambda)$ follows the same geodesic equation as light, where the covariant derivative by μ of V^ν is defined by $\nabla_\mu V^\nu = (\partial_\mu V^\nu + \Gamma_{\mu\lambda}^\nu V^\lambda)$ and along a path [55] by $D_\lambda = dx_\lambda^\mu \nabla_\mu$,

$$(2.16) \quad D_\lambda \partial_\lambda x^\mu = 0$$

In GR, the six possible independent components of the waveform reduce to just two h_+ & h_\times terms, although other theories of gravity predict additional polarizations. These plus h_+ and cross h_\times polarizations make the metric waveform along the z -axis ($k_\mu = \text{diag}(1, 0, 0, \omega)$, initial phase ϕ_0) in transverse-traceless gauge:

$$(2.17) \quad h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & -h_+ & h_\times & 0 \\ 0 & h_\times & h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \Re \left(e^{\sqrt{-1}(k_\mu x^\mu + \phi_0)} \right).$$

Tensor polarization corresponds to the prediction of hypothetical gravitons being spin-two particles, analogous to the spin-one photon with a vector-polarized electromagnetic field. GWs in GR travel at the same speed of light as electromagnetic waves, for which experimental data [132] has matched theory [92] so well that c is now embedded into our system of units. Any deviation measured in the speed of GWs from c , as might be seen in coincident electromagnetic observations of the same source, would be of great theoretical interest. Unlike light, which has complex interactions with interstellar media [54,122], GWs should pass almost unimpeded through the stars. While the transverse-traceless equation describes the (near-)vacuum, far-field GWs we seek to detect, the source term in Equation 2.15 must be understood to justify our expectation of astrophysical sources.

2.2.3 Radiation from quadrupoles

Section 2.2.2 notes that monopole and dipole sources are forbidden by conservation of mass-energy and momentum. This leaves us to derive radiation from quadrupolar sources in the stress-energy tensor. Starting from the field equations, a Green's function G (impulse response) lets us invert the linear differential operator \square by solving for a point, Dirac delta source δ [55]. As with electromagnetism, the retarded Green's function from a source at y^σ to an observer at x^σ is given (where the Heaviside step-function is θ) by,

$$(2.18) \quad G(x^\sigma - y^\sigma) = -\frac{1}{4\pi|x^i - y^i|} \delta(|x^i - y^i| - (x^0 - y^0)) \theta(x^0 - y^0).$$

Hence the source integral is (n.b., $\sqrt{-|g|} = 1$ in flat space),

$$(2.19) \quad \bar{h}_{\mu\nu} = -16\pi \int d^4x \sqrt{-|g|} G(x^\sigma - y^\sigma) T_{\mu\nu}(y^\sigma),$$

$$(2.20) \quad = 4 \int \frac{d^3x}{|x^i - y^i|} T_{\mu\nu}(t - |x^i - y^i|, y^i).$$

Note that the latter equation is integrated over time. Converting to far-field at distance $r = |x^i - y^i|$ and retarded time $t_r = t - r$, then Fourier transforming \bar{h} and T over time into $\tilde{\bar{h}}$ and \tilde{T} respectively, we can impose the Lorenz gauge ($\tilde{\bar{h}}^{0\nu} = (\sqrt{-1}/\omega)\tilde{\bar{h}}^{i\nu}$),

$$(2.21) \quad \tilde{\bar{h}}_{ij}(t, x^i) = -\frac{2}{r} e^{\sqrt{-1}\omega r} \omega^2 \tilde{I}_{ij}(t_r),$$

where I_{ij} is defined as the quadrupole moment with the Kronecker δ ,

$$(2.22) \quad I_{ij} = \frac{1}{3} \int d^3x T_{00} (3x_i x_k - r^2 \delta_{ij}).$$

Inverting the Fourier transform yields the time-domain equation for gravitational radiation in terms of a quadrupole moment,

$$(2.23) \quad \bar{h}_{ij}(t, x^i) = \frac{2}{r} I_{ij,00}(t_r).$$

Deducing radiated power requires a second-order perturbative expansion [55, 42], yielding Equation 2.4. While this second-order expansion makes many assumptions, the implied power in the gravitational wave strain is physically (but not literally) observable, as in PSR 1913+16. The goal of gravitational wave astronomers now is not just to infer radiated power but to measure $h_{\mu\nu}$ directly.

Both polarization states, h_{\times} & h_{+} , can be measured directly. The interferometer antenna patterns for each polarization are detailed in other references [91], but in particular the LIGO interferometers are most sensitive to a wave travelling perpendicular to, and aligned to maximally stretch, each arm. In coordinate system of Equation 2.17, this would be a plus-polarized h_{+} wave travelling along the z -axis, with arms aligned on x - and y -axes. Conversely, a h_{\times} wave along the z -axis would not stretch the arms and would then not be seen. The expected polarization for continuous wave signals is discussed in Section 2.3.2. GW sources in general can emit, as with light, a mixture of polarizations.

Should more than the expected two GW polarizations be seen, it would, as with deviations from the speed of light, be of great theoretical interest. A full six states are possible in GR-like theories [178, 179]. This full table of six states, with standard two polarizations at top, are represented in Figure 2.2.

Given our expectations for GWs are predicted by GR, LIGO and allies have built highly-sensitive interferometers to search for the signals anticipated in Section 2.3.

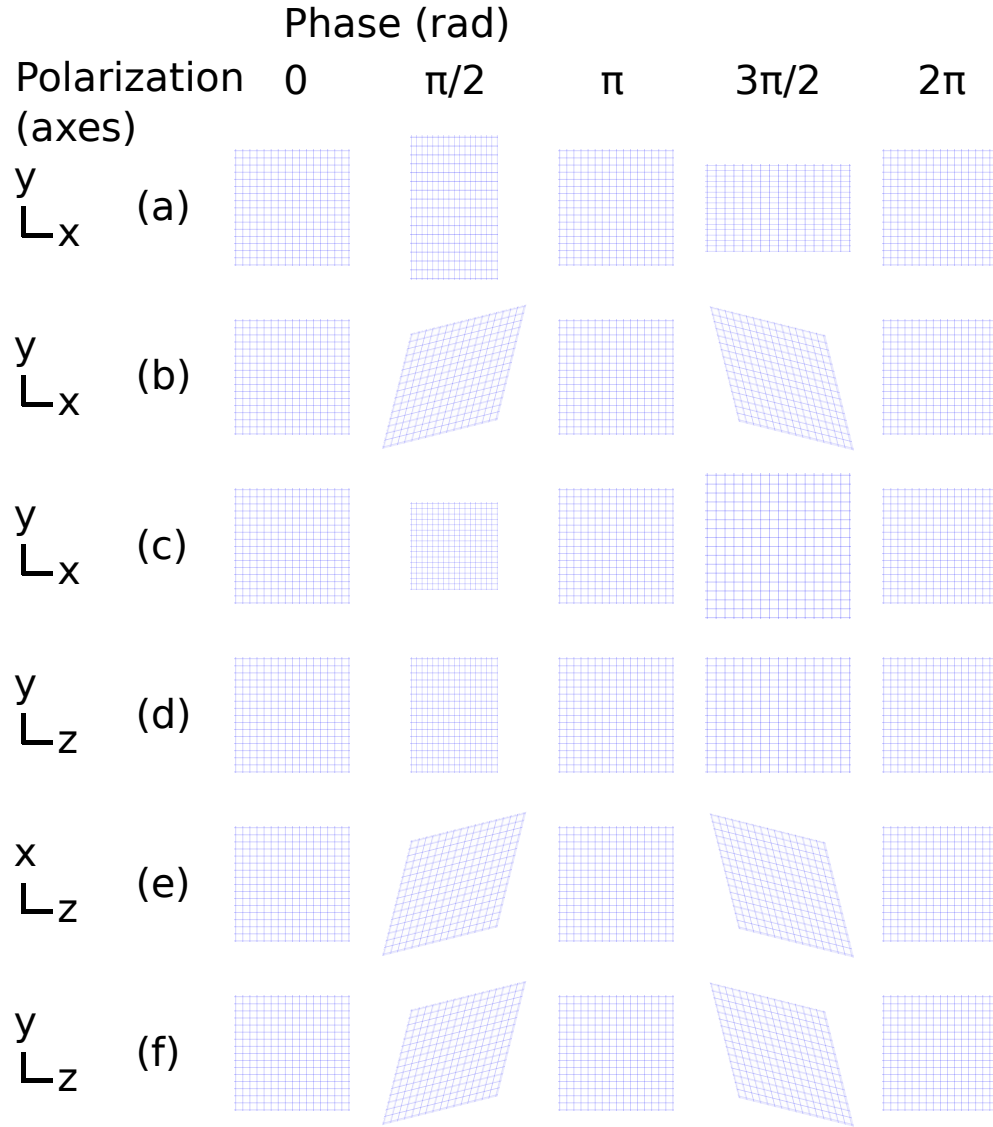


Figure 2.2: Six GW polarizations are permitted by metric theories of gravity [178]. Here, all are shown propagating in the z -direction (note different axes along left column: the wave travels into the plane of the page for the top three, along the plane of the page for the bottom three). GR permits the top two, (a) and (b); scalar-tensor-vector theory permits also (c); the last three are found in other theories [179]. Conventionally, (a) is labelled h_+ and (b) is h_\times . Shown is the metric perturbation; physically, this corresponds to measured relative motion of test particles. A complete wave cycle is depicted. Off-page (not shown) third- and time-axis are not perturbed. For illustration, the figure shows strain amplitude $h = 0.25$, large enough that the linearized GR model would become inaccurate. This strain is over 20 orders of magnitude larger than the strains LIGO typically expects; LIGO is designed with GR and h_+ and h_\times polarizations in mind.

2.3 Astrophysical estimates

Spacetime is stiff. Even large radiated power corresponds to low GW amplitude. Since anthropogenic gravitational waves are too quiet to plausibly detect [149], attention has turned to the cosmos. One Hubble time is roughly 8×10^{60} in geometric units, and one-and-a-half-generation interferometers were built with optimal ‘horizon’ distances on the order of $10^{59} P_L$ Planck lengths, $\mathcal{O}(50 \text{ Megaparsec})$, in hopes of seeing perhaps one inspiral per year. Such astronomical scales make GWs a distinctly classical phenomenon, yet their effects on Earth are on the quantum scale. These interferometers had displacement sensitivity of at best $8 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}} = 5 \times 10^{15} P_L/\sqrt{\text{Hz}}$, or strain sensitivity $2 \times 10^{-23}/\sqrt{\text{Hz}}$. Advanced detectors should push strain sensitivity down to as low as $4 \times 10^{-24}/\sqrt{\text{Hz}}$. Astrophysical estimates, elucidated next, will let us understand what these sensitivities might unveil for the sources of Section 2.1.

2.3.1 Sources: burst, continuous, inspiral and stochastic

Ground-based interferometric detectors concentrate on the aforementioned four searches: burst, continuous, inspiral, and stochastic [146]. As the focus of this thesis, continuous sources are expounded on in Section 2.3.2.

Compact binary coalescences, or ‘inspirals’, are sufficiently well-defined by numerical relativistic simulations to permit templated searches. Templating allows for high-sensitivity matched filtering, characterized by an *inspiral range* [80] for an average-orientation, average-sky location coalescence of two neutron stars to be detected at a signal-to-noise ratio (nominally 8). Equivalently, to a constant factor, one can calculate the optimal orientation & location *horizon distance*. These ranges yield a detectable volume, which, integrated over the time of a science observation run, yields an expected number of detected events. Six LIGO science runs have been

conducted to date, with the most recent, S5 (2005 November to 2007 October) and S6 (2009 July to 2010 October), attaining $\mathcal{O}(1 \text{ year})$ coincident data between multiple observatories at inspiral ranges up to 15 Megaparsec in S5 and 20 Megaparsec in S6; their successor instruments each have a planned inspiral range of 200 Mpc [97]. To date, no true GW events have been discovered. Neutron star binary merger rates between 2×10^{-4} to 0.2 (most likely 0.02) per year would be expected for an instrument at S5 sensitivity; advanced detectors anticipate rates of 0.4 to 400 (most likely 40) per year [11].

Conversely, the absence of gravitational wave detections, coupled with confidence in calibration, allows setting upper limits. LIGO and Virgo have estimated the rate of non-spinning binary black hole mergers (3.3×10^{-7} per cubic Megaparsec per year) [5] and binary neutron star ($1.3 \times 10^{-4} \text{Mpc}^{-3} \text{yr}^{-1}$) or neutron star-black hole ($3.1 \times 10^{-5} \text{Mpc}^{-3} \text{yr}^{-1}$) mergers [15], with 90% confidence. If Advanced LIGO and Virgo find no gravitational waves, resulting upper limit rates may begin to conflict with the astrophysical expectations, which could be interesting in its own right, though no conflict is seen yet. Nonetheless, detection is the goal.

Burst searches rely on coincident observations among detectors – and possible electromagnetic or neutrino counterparts – even more than inspiral searches. Short-hard gamma ray bursts (GRBs), thought to arise from the merger of a neutron star progenitor with another neutron star or black hole, can be seen with satellite observatories such as Swift and Fermi and correlated with LIGO data [14]. Long GRBs [6] might have long-lived gravitational wave counterparts. GRB detections thus trigger worldwide alerts for GW observatories to quiet local noise sources temporarily in case a corresponding GW signal proves detectable. GW observatories are also able, inversely, to provide trigger information so that electromagnetic observatories can

search for candidate events [8]. Bursts also have the capability for purely coincident analyses, without any prior assumptions about waveform [13]. As yet, no burst detections have been made, but direct coordination with the Swift gamma-ray [77] and Antares neutrino [32] observatories, and many more, has already begun and is foreseen to expand in the advanced detector era.

Stochastic background measurements could offer insight into early universe cosmology [93, 118]. Although reliant on correlated signals between detectors, as burst searches exploit, stochastic search algorithms integrate over long durations rather than seeking distinct events. While the expected background from the Big Bang and foreground objects (possibly ranging from white dwarves to superstrings) is well below the expected sensitivity of initial and advanced detectors, upper limits have already contributed to physics by excluding some models of universe evolution and string theory [163], bettering indirect limits extant from Big Bang nucleosynthesis by limiting the cosmic critical energy density due to gravitational waves to be $\Omega_{gw} < 6.9 \times 10^{-6}$ at 95% confidence in the LIGO band. Upper limits at high frequencies were also improved by a factor of seven [16]. Over time, methods [36, 83, 20, 19] have refined so that not only all-sky but also directional limits can be obtained. Anisotropies could be imparted by the Earth's motion or by the distribution of sources in this or nearby galaxies [35]. Radiometer combination of data streams [44] could search for these point sources, and can further be adapted to multipolar decomposition [125] of the sort that has proved crucial for understanding the cosmic microwave background [133]. While prospects for the next generation are tentative, GWs promise a way unavailable by any other instruments to see a background from the dawn of time.

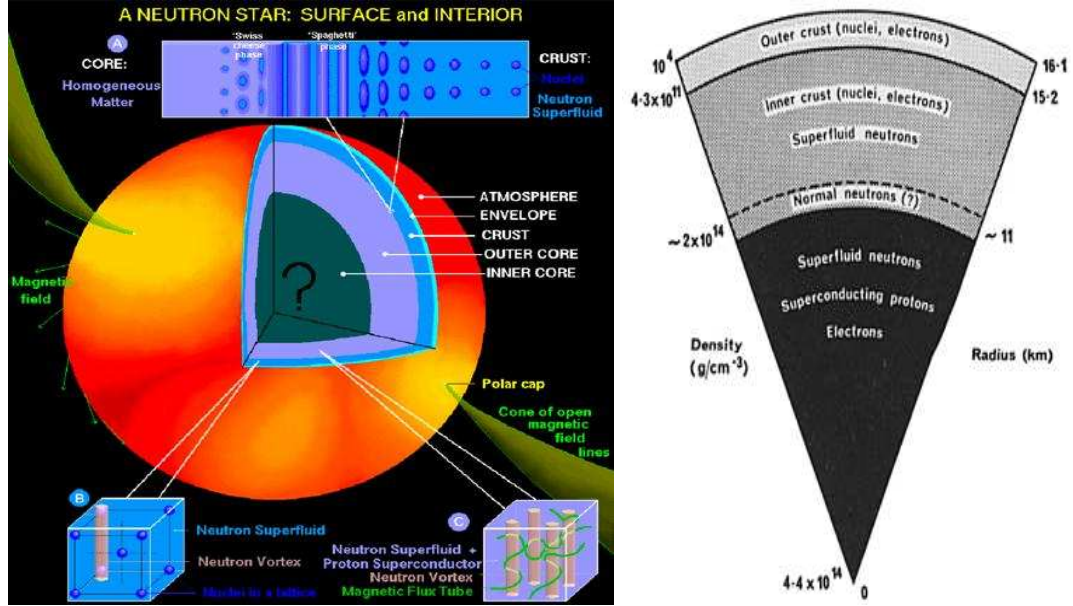


Figure 2.3: Hypothetical internal structures of a neutron star. Left: theoretical neutron superfluid vortices and magnetic field lines (citation: Dany Page, http://www.learner.org/courses/physics/visual/visual.html?shortname=a_neutron_star). Right: putative depths of internal layers (citation: http://heasarc.gsfc.nasa.gov/docs/objects/binaries/neutron_star_structure.html).

2.3.2 Continuous waves from neutron stars

Spinning neutron stars (NS) [141] motivate clear yet computationally challenging GW searches. Perhaps one NS is born per century in the Milky Way [134]. In the solar-system reference frame, an isolated NS would emit a pure sine wave, as predicted by (with leading constant G_C/c^4) Equation 2.21. One emission mechanism of gravitational radiation with amplitude h_0 from a NS is from a surface deformation defined as ellipticity ϵ (outer crust, Figure 2.3), leading to a quadrupole moment I [181, 22] and thus emission at frequency $f = 2\nu$, twice the NS spin frequency ν :

$$(2.24) \quad \epsilon = \frac{I_{xx} - I_{yy}}{I_{zz}},$$

$$(2.25) \quad h_0 = \frac{4\pi^2 G_C}{c^4} \frac{I_{zz} f^2}{r} \epsilon.$$

Emission might also arise from r -mode oscillations or free precession [152]. Expected sensitivities for terrestrial interferometers, best around 100 Hz, highlight NS with comparable periods (millisecond pulsars) as suggestive targets. For GWs from a rotating NS with a spin axis inclined at angle ι to an observer's line-of-sight, the relative polarization amplitudes [66] would be,

$$(2.26) \quad h_+ = \frac{1}{2}h_0(1 + \cos^2 \iota),$$

$$(2.27) \quad h_\times = h_0 \cos \iota.$$

For $\cos \iota = 1$, where the spin axis is along the line of sight, gravitational waves will thus be circular polarized; $\cos \iota = 0$, an orthogonally-aligned spin axis, implies linear polarization. Excepting the case of torque balance [139, 168] in low-mass X-ray binary systems (LMXB), NS rotation frequency will decay due to gravitational radiation, yielding a relationship between the 'spindown age' [51] τ and the gravitational radiation amplitude:

$$(2.28) \quad \tau = -f/\dot{f},$$

$$(2.29) \quad h_0 = \frac{1}{r} \sqrt{\frac{5G_C I_{zz}}{2c^2 \tau}}.$$

Although pure sine waves are considerably simpler than the numerical relativistic templates needed for inspirals, the low expected h_0 and consequent long-duration integration needed to detect NS mean that the search over parameter space is extremely sensitive to mismatched templates. This sensitivity implies that the algorithms must use fine-grained templates, necessitating long computational time. Algorithms have been developed to search for these waves, starting from maximum likelihood estimates of isolated NS [106]. These have become progressively more sophisticated,

incorporating barycentric solar-system correction [17], distributed computing [25], and semi-coherent methods [26]. Many isolated pulsars [18, 22], including the Crab and Vela pulsar remnants [9], as well as the LMXB Scorpius X-1 [23] and the direction of the Galactic Center [3] have been the target of these methods.

Many NS in known catalogs, such as ATNF (Australia Telescope National Facility [119]), that pulsate in the LIGO frequency band are in binary systems. This thesis describes, in Chapters V & VI where this discussion continues, a computationally-efficient means, TwoSpect [90] of searching for NS in binary systems, particularly LMXBs. First, we discuss how gravitational wave interferometers sensitivity can be improved to make astronomy possible.

2.4 Laser Interferometer Gravitational-wave Observatories

Design and construction of the Laser Interferometer Gravitational-wave Observatories required extensive prototyping and commissioning both during Initial LIGO [121, 31, 42] and Enhanced LIGO [85, 156, 70]. The core features of interferometer design were established even before construction, summarized by Saulson [149], but many technical noise sources were encountered. Motivations, fundamental noises, and the developments of commissioning all merit attention.

2.4.1 From Weber bars to interferometry

Between the earliest interferometry of Michelson [130] and present-day LIGO came the resonant bar detectors of Joseph Weber [171]. Bar detection of GWs developed rapidly through the 1960s, but contradictory results dampened enthusiasm [149, 63]. Nevertheless, great effort was devoted to technically excellent bars such as ALLEGRO, AURIGA, EXPLORER, NAUTILUS and NIOBE over the subsequent decades. Meanwhile, prototypes by Rainer Weiss, Robert Forward [82], Ronald Drever and

many others demonstrated that interferometry could also achieve the sensitivities desired for gravitational wave detection. Interferometry promised better fundamental noise sources and broadband operation, meaning that a single instrument could listen to a wider frequency band. Xylophones of bars were proposed, but never built, that might scale the few Hz bandwidths of bars to the hundreds to thousands of Hz of interferometers. Gravitational wave interferometry has carried the prospects for GW astronomy toward unprecedented levels of sensitivity, thanks to performance factors detailed in Section 2.4.2.

2.4.2 Gravitational wave interferometry in theory

Michelson invented modern interferometry [130] in his 1887 experiments with Morley to try to measure the velocity of the Earth with respect to the luminiferous ether. Although commonly claimed to have spurred the work of Minkowski, the implications of this work for special relativity may not have been fully understood for several decades [64]. Yet the value of the experimental technique was undeniable. Michelson & Morley had invented an accurate relative phase-meter for light.

Existing GW observatories, at their cores, are also Michelson interferometers. In contemporary GW interferometers, monochromatic laser light (typically 1064 nm) is used. This light is split into two equal beams by the beam-splitter, reflects off mirrors at the ends of two arms, X and Y, and returns to interfere at the beam-splitter after times T_x and T_y . The power P at the beam-splitter output ‘dark’ port, is by design near zero except when the measureable (e.g., ether, a GW, or other disturbance) causes a fringe shift $\Delta\phi$ related to the angular frequency ω of the light and the time-of-flight down the x and y arms, T_x and T_y . Energy is conserved; the change of sign on the beam-splitter reflective surface means that destructive interference on the ‘dark’ port implies constructive interference on the opposite ‘bright’ port. Power

is classically proportional [105] to the squared electric field [149]:

$$(2.30) \quad P = \frac{\epsilon_0 c E_0^2}{2} |1 - e^{i\Delta\phi}|^2 = 2\epsilon_0 c E_0^2 (1 - \cos \Delta\phi),$$

$$(2.31) \quad \Delta\phi = 2\pi(L_s + L_0)/\lambda + \omega(T_y - T_x).$$

In Enhanced and Advanced LIGO, L_s is known as the Schnupp asymmetry (a constant difference in arm length related to RF-modulated sidebands) [31], and L_0 is a smaller constant offset from a fringe to enable non-radio-frequency-modulated DC readout [85]. A putative, even smaller ‘DARM’ phase shift can read out $\omega(T_y - T_x)$, due to differential arm motion, encoding any time-varying GW signal in the relative time-of-flight between the arms. Operating near the maximum of $dP/d(\Delta\phi)$ would increase susceptibility to unwanted power fluctuations, so the interferometer was initially operated with radio-frequency modulation at minimum $dP/d(\Delta\phi)$ and P (a dark fringe), later changed to a slight offset and DC readout once the noise sources were better understood.

The GW sensitivity of the power P can be directly improved by increasing effective laser power (e.g., a brighter laser or use of ‘power recycling’), and that of $\Delta\phi$ can be increased by making $T_{x,y}$ storage times longer (e.g., with Fabry-Perot cavities). Doing so allows more accurate measurements of P in the presence of shot noise (fluctuations in photon counting, entering via the E_0 term), radiation pressure (photon momentum-induced fluctuations in the mirrors causing changes in $T_y - T_x$), and other motions that change $\Delta\phi$ (such as thermal and seismic). Yet these noise sources implicate quantum effects, and the photodiodes used to measure power do so with a certain quantum efficiency. More realistic descriptions by Caves [57, 58] clarify the quantum mechanical origin of shot and radiation pressure noise. Quan-

tum squeezing, the subject of Chapter IV, can mimic the effect of higher power in Equation 2.30. GWs, eminently classical phenomena, are small enough in amplitude that quantum behavior is a key limit on detection.

Gravitational wave signals enter the interferometer by changing the spacetime curvature in the arms. In the simplest, linearized picture, the spacetime curvature arises from a GW with amplitude h at frequency f . When the GW is linearly polarized in h_+ and aligned with x-axis and y-axis arms of equal length L (see [106] for the antenna pattern correction when it is misaligned), it changes the time-of-flight in the arms according to Equation 2.17:

$$(2.32) \quad T_x = \int_0^{\frac{2L_x}{c}} \sqrt{|g_{xx}|} dt \approx \int_0^{\frac{2L_x}{c}} \left(1 - \frac{h_+(t, x(t))}{2} \right) dt,$$

$$(2.33) \quad T_y = \int_0^{\frac{2L_y}{c}} \sqrt{|g_{yy}|} dt \approx \int_0^{\frac{2L_y}{c}} \left(1 + \frac{h_+(t, y(t))}{2} \right) dt,$$

We have been able to neglect issues of the proper-time of the mirrors (as opposed to the photodiode) because the laser frequency ω is unchanged ($g_{tt} = 0$ in transverse-traceless gauge). The mismatch in the integrals leads to a detectable time-varying signal,

$$(2.34) \quad \phi_{\text{DARM}} \equiv \omega(T_y - T_x) = \omega \int_0^{\frac{2L}{c}} \frac{h_+(t, x(t)) + h_+(t, y(t))}{2} dt.$$

Measured by Equation 2.30, ϕ_{DARM} encodes the GW signal. Comparing ϕ_{DARM} from different observatories, or waiting while a single observatory is Doppler-modulated by terrestrial motion, can pinpoint the origin of such a signal.

We often assume that the time-of-flight is much less than the GW period, $T_{x,y} \ll 1/f$. This assumption is easily relaxed by accounting for the destructive self-interference

of the signal light from a high f signal [149]. A straightforward interpretation emerges for a quasi-static plane-wave (when phase offset $2\pi L_0/\lambda \gg \phi_{\text{DARM}}$) – the interferometer linearly transduces GW strain (via $h_+ = \phi_{\text{DARM}}c/(2L\omega)$) into power:

$$(2.35) \quad \frac{dP}{dh_+} = \frac{2L\omega}{c} \frac{dP}{d(\Delta\phi)} = 4L\omega\epsilon_0 E_0^2 \sin\left(\frac{2\pi}{\lambda}(L_s + L_0)\right),$$

$$(2.36) \quad P = \left[4L\omega\epsilon_0 E_0^2 \sin\left(\frac{2\pi}{\lambda}(L_s + L_0)\right)\right] h_+.$$

2.4.3 Interferometer theory in practice

Interferometers transduce GW strain into power in a manner like Equation 2.36 only because they operate in a linearly-controlled regime [85] with noise sources suppressed as well as possible [149, 174]. A complete discussion of noise is beyond our scope, but the key limits on a ground-based GW interferometer arise, for a facility in vacuum, from microseisms, gravitational gradients, thermal coating and thermal suspension noise, and quantum effects (radiation pressure and shot noise).

Seismic noise dominates the lowest frequencies, and it is suppressed by use of servo controls, as well as by passive damping from pendula. Each stage of a multi-stage pendulum reduces, for frequencies above its natural resonant frequency f_0 , the coupling of seismic noise into the interferometer by a factor $(f/f_0)^2$.

Gravitational gradients due to clouds, human activity, and density fluctuations in the ground are an area of concern, but may be addressed through active feedforward servos [72], as may magnetic couplings.

Thermal coating noise is reduced by using larger laser spot sizes on the surfaces of mirrors and optimizing the material properties of the coatings; the mirror itself must be made of a material with a high Q , such as fused silica, where Q is the quality factor for internal oscillations. Suspension thermal noise is also reduced by using high- Q

materials, *per* the fluctuation-dissipation theorem [149]. These effects predominate in the logarithmic middle of LIGO’s frequency range, termed the ‘bucket’.

Shot noise is naively reduced by increasing laser power, which indeed yields improvements up to a point. At high laser power, however, thermal distortions due to absorption in the mirrors are inevitable, leading to the need for thermal compensation systems [42]. High laser power also buffets the mirrors with radiation pressure noise, requiring heavier mirrors. Quantum squeezing can circumvent the issue of thermal distortions, but it too induces radiation pressure noise (typically worse, because the Heisenberg uncertainty area increases); a solution may be squeezing at high (shot noise-limited) frequencies and anti-squeezing at low (radiation pressure-limited) frequencies, as discussed in Chapter IV.

GW interferometers on Earth are limited in size (thus storage time and strain sensitivity) by the planet’s curvature (which couples vertical displacement noise into horizontal) as well as by finances. GEO600 uses delay lines, and the other interferometers use Fabry-Perot arms to extend the storage time of light. Fabry-Perot arms with high-finesse [144, 143] (roughly 137 in Initial LIGO) are locked [48] using the Pound-Drever-Hall (PDH) technique. Electro-optic modulators (EOMs) provide a radio-frequency error signal for PDH locking the Fabry-Perot arms.

Figure 2.4 presents a schematic overview of the key systems in Enhanced LIGO.

2.4.4 Observatory operation

Satisfying the noise requirements above, the first-generation LIGO detector was designed with the following features:

- Power recycled Michelson interferometers with 4-km Fabry-Perot arms³
- Hanford, Washington and Livingston, Louisiana observatories

³A second 2-km interferometer was sited at Hanford as well during Initial LIGO.

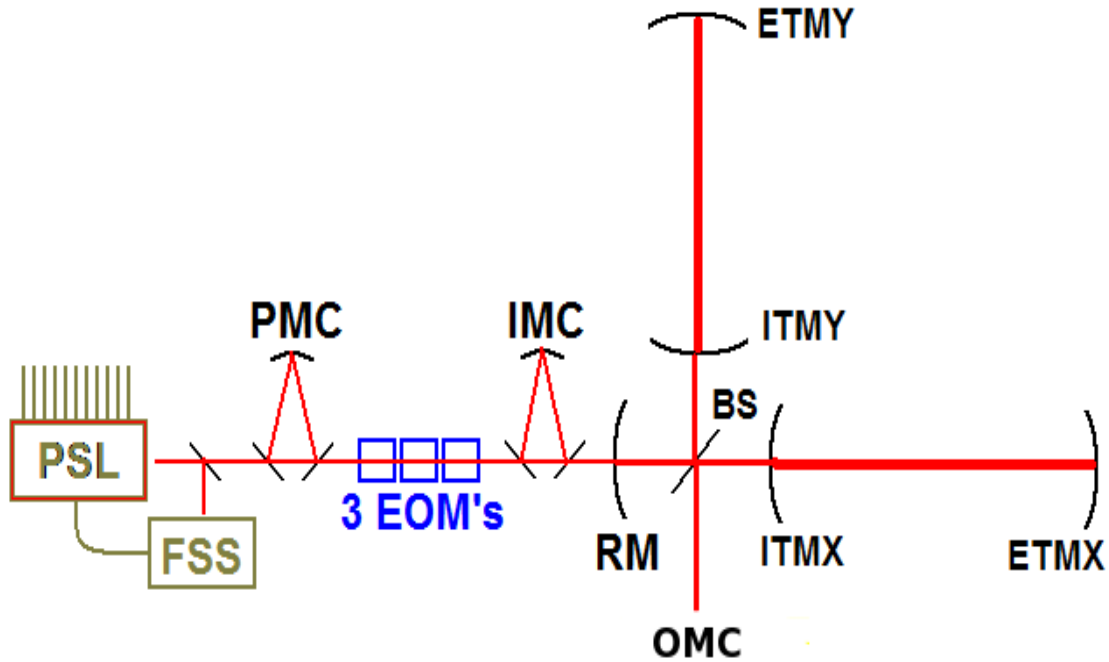


Figure 2.4: Enhanced LIGO primary systems and optics. From left, the phase-stabilized laser (PSL) produces coherent light, kept at constant frequency by the frequency-stabilization servo (FSS). The pre-mode cleaner excludes non-Gaussian beamshapes, allowing only the TEM_{00} mode to pass. The beam is phase-modulated using three electro-optical modulators (EOMs) before being further shaped in the input mode cleaner (IMC). The beam passes through the power recycling mirror (RM) and is split at the beam-splitter (BS). Along the X and Y arms, light resonates in the Fabry-Perot cavity formed between the input test mass (ITM) mirror and the end test (mass) ETM mirror, before recombining at the beamsplitting. Any light with the same phase as before the arms is returned to the interferometer by the RM, but if phase-shifted, perhaps by a gravitational wave, it exits through the output mode cleaner (OMC): a photodiode just downstream of the OMC records the signal.

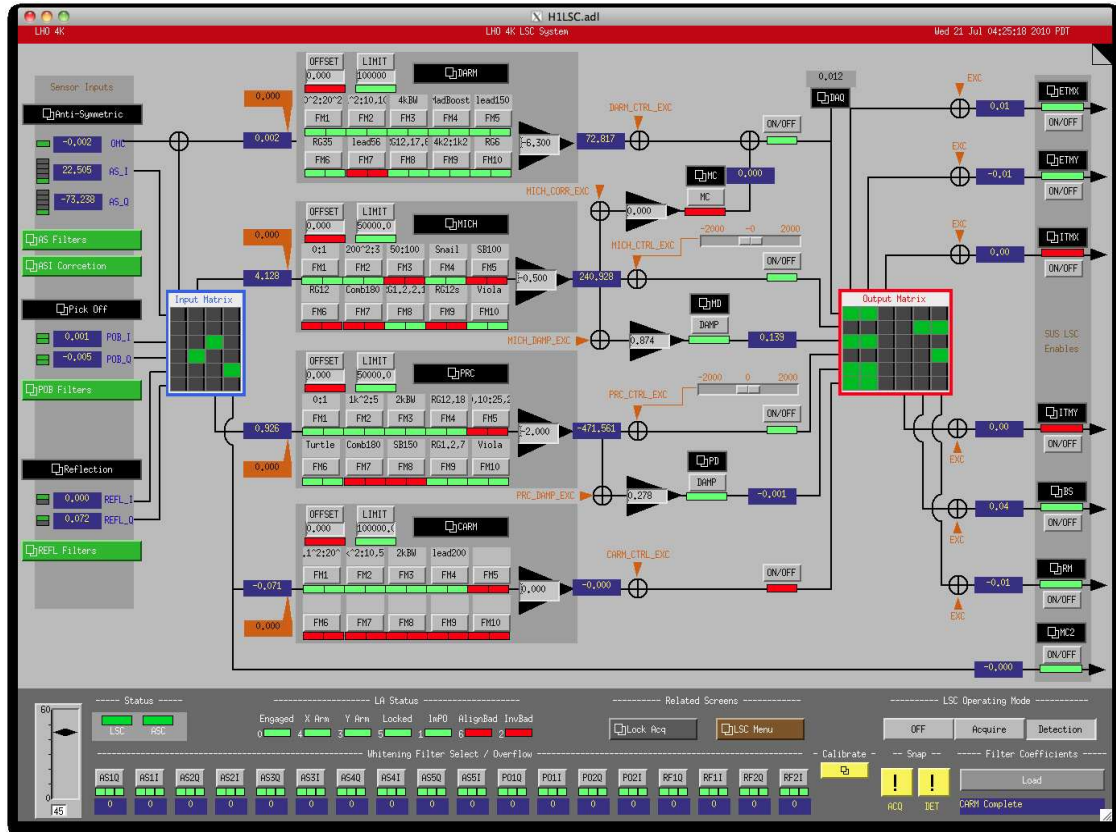


Figure 2.5: Screenshot of MEDM control panel. MEDM is a Motif Epic and Display Manager for EPICS, the Experimental Physics and Industrial Control System, which lets operators control LIGO. MEDM allows operators to run locking scripts as well as alignments and tests, and to activate and de-activate filters.

- 4-km perpendicular beam tubes
- 10^9 torr vacuum
- 10 W Nd:YAG 1064 nm laser (Enhanced LIGO upgraded to 20 W)
- 10-kg fused silica primary optics
- 4-stage seismic isolation (including active hydraulics at Livingston)
- Laser frequency stabilization
- Angular sensing and control (wavefront sensors, optical levers)
- Length sensing and control (magnet coils, common mode)
- Pre- & input-mode cleaning (Enhanced LIGO added output mode cleaning)
- Power recycling
- Digitally filtered servos and readout
- RF heterodyne (Initial LIGO) or DC homodyne (Enhanced LIGO) readout

LIGO's digital systems were managed with MEDM and EPICS, seen in Figure 2.5. Extensive effort on installation was required to attain the necessary angular, length, and auxiliary servo control needed for Initial LIGO operation [87]. Science Run 5 succeeded in obtaining the desired sensitivity and duty cycle. Although the RF heterodyne technique⁴ had good low-frequency sensitivity, it limited the shot noise performance compared to a DC homodyne instrument by a factor of $\sqrt{3/2}$, and squeezing was also easier in DC. The switch to DC readout occurred prior to S6, and in the leadup to that science run, substantial detector characterization was needed to understand the upgraded interferometer.

⁴Curiously, the RF power sidebands as measured by the author were not always equal in the recycling cavity [123].




Figure 2.6: Omega scan of an audible broadband glitch. Shortly before the start of Science Run 6, detector characterization took place to identify categories of glitches, such as the ‘gremlin’, and to eliminate them. The burst group analysis pipeline, Omega, generated time-frequency spectrograms that made this identification easier. Glitches are a limiting factor in the identification of rare events and therefore of potential gravitational wave burst signals.

2.5 Detector characterization and development

Commissioning Enhanced LIGO in the three months encompassing the start of S6 around July 2009 required eliminating the residual noise sources lingering from installation. In particular, the Output Mode Cleaner [156] was new, and its servo-locking scheme (involving both piezoelectric and thermal actuation) required some artistry. DC readout generally increased the low-frequency noise floor. New glitches sprang forth as a result, and the author’s work at LIGO Hanford Observatory centered on identifying them.

Common tools for investigators on-site are both time-domain (DataViewer) and frequency-domain (Diagnostic Test Tools) data-plotters. While those perspectives allowed diagnosis of at least one flaw in the interferometer’s filters, a combined, spectrogrammatic view using the burst group’s Omega analysis pipeline proved a needed third, unified view, seen in Figure 2.6. This pipeline was run on many pre- and early-S6 glitches using a manual infrastructure⁵.

Effective engines for tracking and, preferably, eliminating glitches and persistent lines in the detector spectrum – the process of detector characterization – are critical to distinguishing rare events.

2.5.1 Feedforward filtering

Feedback servos keep LIGO operational, holding the relative positions of the mirrors and auxiliary systems in stationary ‘lock’ points. Closing the loop in feedback requires a way to actuate a physical control signal to cancel out whatever influence is causing an error signal that pushes the system away from its lock point. When this direct cancellation is not possible – for instance, when ambient magnetic fields due to 60 Hz mains lines cannot be escaped – feedforward remains. Feedforward, an open-loop technique, can be done purely in software as discussed in Chapter III, but it requires an accurate estimator for the influence of a noise source into the system.

In the case of Enhanced LIGO commissioning, an example was implemented with the aforementioned 60 Hz line using a magnetometer to supply the feedforward correction [156], which greatly enhanced sensitivity near the Crab pulsar frequency by suppressing magnetic-coupled noise. Subsequently, into the following years, this author began to work on the long-known [31, 42] potential improvement from cancelling the influence of auxiliary length control noises. These noises exist in the mirror servos

⁵Omega scan diagnostics were later automated by Tomoki Isogai.

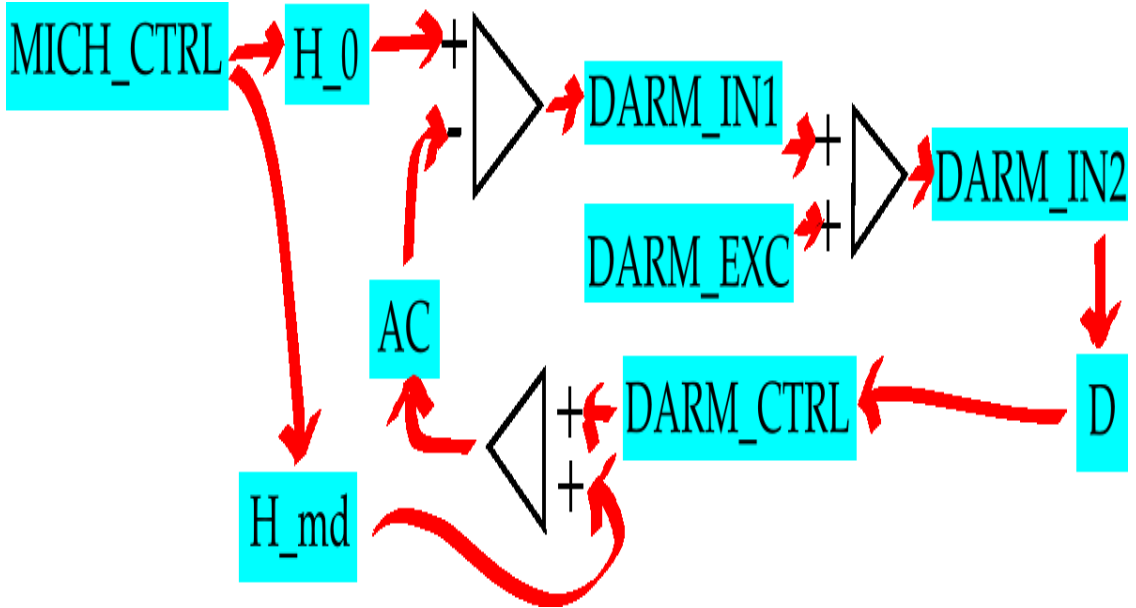


Figure 2.7: Real-time servo loop diagram. The MICH_CTRL signal can be thought to leak into the true DARM signal via a transfer function, H_0 . (The letter ‘h’ is typical for transfer functions as well as gravitational wave strain; coincidentally, this is a noise transfer into the channel for displacement, DARM, which is proportional to strain $h(t)$). H_0 leaks into error signal, which is otherwise kept null thanks to the servo cancellation provided by DARM_CTRL times a physical actuation function AC. The measured error signal is DARM_IN1. If desired, an excitation can be supplied via DARM_EXC for a sum of DARM_IN2. This error signal passes through digital filters D to yield the aforementioned control signal DARM_CTRL. The auxiliary length cancellation loop is simply adding H_{md} to DARM_CTRL in order to subtract out the corruption of H_0 .

and contaminate measurement channels for gravitational waves. Loops for nulling this effect exist, diagrammed in Figure 2.7, but they require periodic retuning [108]. The direction of this work was to find an automated, and more precise, way to tune these servos. A finely tuned servo filter was implemented in late September 2010⁶. This filter reduced the noisiest residual coupling, from the differential inner Fabry-Perot mirror motion to the differential arm motion that encodes strain, by almost an order of magnitude.

Figure 2.8 illustrates the closed loop gain. Control theory shows how this gain is a function of frequency. The closed loop gain is the geometric sum of time-delayed (by

⁶After, and unrelated to, the famed Big Dog blind injection [146] in the same month, simulating an inspiral event.

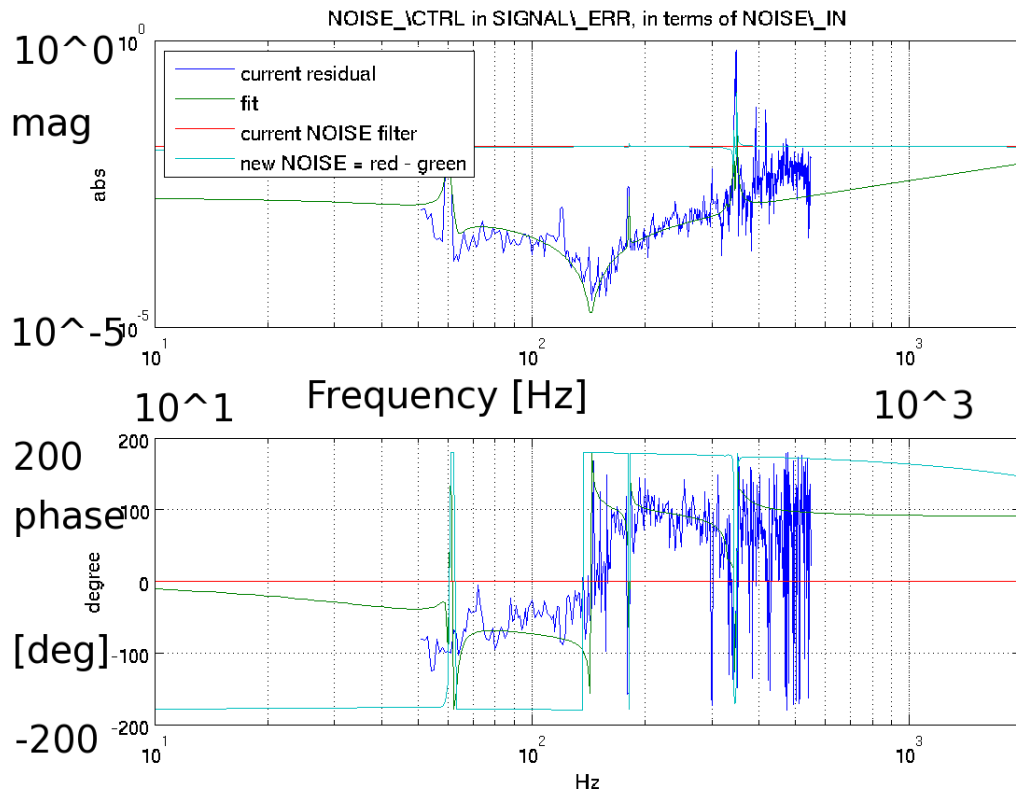


Figure 2.8: Real-time work on a LIGO noise filter. This Bode plot, made in Matlab, shows the correction to the existing MICH damping (cancellation) loop needed mid-2010, toward the end of Science Run 6. The correction is small, because the majority of the coupling fits the flat model expected from theory. This transfer function estimate suffices for post-factor correction. However, in order to be incorporated into the control scheme shown in Figure 2.7, measurements of the open loop gain G and actuation function AC are necessary to incorporate the filter correctly into the closed loop response.




Figure 2.9: Early work on post-facto noise filtering. After testing out MICH damping filters offline, post-facto, the correction factors for AC and G_{closed} were incorporated and the entire filter imported using Foton into the EPICS control system, where it was used real-time from the September equinox of 2010 until the end of Science Run 6.

τ) open loop gain responses, by the DARM_CTRL control channel, to error signals from the error signal DARM_ERR: $G_{\text{closed}} = \lim_{p \rightarrow \infty} \sum_n^p G e^{in\omega\tau} = (1 + G e^{i\omega\tau})^{-1}$. Because MICH damping sums with DARM_CTRL, it acquires an additional factor of the actuation function AC for a net transfer function of $H_{md} G_{\text{closed}} \text{AC}$ in the closed loop. Compared to an open-loop subtraction (effect estimated in Figure 2.9), a real-time servo will differ by a factor of $G_{\text{closed}} \text{AC}$. See Chapter III for full details of open-loop subtraction and the extension to a post-facto, veto-safeguarded improvement in the sensitivity of all S6 data.

2.5.2 Phase camera

Improving length control is intuitive, since length directly correlates with LIGO's inference of GW strain. Angular control, however, is just as necessary: LIGO uses



Figure 2.10: Optical table layout with Fabry-Perot cavity and Pound-Drever-Hall locking for phase camera. Faraday isolator and polarizing beam-splitter at upper right; piezo-electrically-actuated mirror at center left. Difficulties with stability, despite a plexi-glass enclosure and floated table, meant that locks were fractions of a second at most, hampering efforts.

quadrature-photodiode wave-front sensors [121] to minimize misalignment. Misalignment causes direct angle-to-length coupling, adding noise to the system by corrupting the phase measurement at readout; it also permits power fluctuations in the Fabry-Perot arms, reducing shot noise-limited performance and, when the relative power returning from the arms (the contrast defect) changes, directly corrupting the power measurement at readout [70]. Maintaining lock also becomes hard if misalignment grows too severe. Therefore more sophisticated systems have been researched in hopes of reducing angular fluctuations still further, including exploration of a phase camera at Michigan [66].

Locking uses the Pound-Drever-Hall (PDH) method [48, 121]. Mathematically, the method assumes two mirrors: an input, with amplitude reflectivity r_1 , and an end, with amplitude reflectivity r_2 . Upon the input mirror is incident a coherent electric field, due to light, E_{inc} at frequency ω_0 . A field E_{refl} is reflected. The key to the technique is that the incident electric field is phase modulated by an amplitude Γ with modulation frequency Ω , by a device such as an electro-optic modulator (EOM, e.g., a Pockels cell). Typically this modulation is done at radio-frequency. This EOM modulates an electric field of amplitude E_0 into higher and lower frequency sidebands, amplitude E_+ and E_- , as determined by expansion in Bessel functions $J_n(\Gamma)$ – higher order sidebands exist at smaller amplitudes. From the reflected signal, intensity I encodes an error signal of the arm length l , for given light wavenumber $k = \omega/c$, at the modulation frequency (besides additional DC and $2f$ signals):

$$(2.37) \quad E_{\text{refl}} = \frac{r_1 - r_2 e^{-2ikl}}{1 - r_1 r_2 e^{-2ikl}} E_{\text{inc}},$$

$$(2.38) \quad E_{\text{inc}} = E_0 e^{i\Gamma \cos(\Omega t)} \approx E_0 e^{i\omega_0 t} [J_0(\Gamma) + J_1(\Gamma) e^{i\Omega t} + J_{-1}(\Gamma) e^{-i\Omega t}],$$

$$(2.39) \quad E_{\text{refl}} \approx e^{i\omega_0 t} [E_0^{\text{refl}} + E_+^{\text{refl}} e^{i\Omega t} + E_-^{\text{refl}} e^{-i\Omega t}],$$

$$(2.40) \quad \begin{aligned} I = & [|E_0|^2 + |E_+|^2 + |E_-|^2] + \\ & [(E_0^* E_+ + E_0 E_+^*) e^{i\Omega t} + \text{C.C.}] + \\ & [E_+ E_-^* e^{2i\Omega t} + \text{C.C.}]. \end{aligned}$$

Radio-frequency photodiodes are generally needed to record this intensity with its modulation. Typically, the photodiode current is demodulated with a mixer that uses as local oscillator the same sine wave that drives the Pockels cell EOM phase modulation. Alternative photoresistor readouts were explored in hopes that they might have faster response times since the standard electronics references were

compiled [102, 155], because of their potential desirability for a low-cost many-pixel readout. Part of this author's work included a Spice simulation and physical construction of a photoresistor circuit. Sadly, the inexpensive photoresistors used had rise times of milliseconds, far too slow for a direct RF readout scheme. This result refocused attention on photodiodes.

Using the usual RF photodiode, a Fabry-Perot cavity was built prior to this author's work, seen in Figure 2.10. My efforts began with characterizing the RF electronics of the analog locking system. A signal analyzer helped re-select a cable of appropriate length and phase delay for the system. Numerous attempts to stabilize the cavity and derive a useful PDH locking signal were made. The laser was upgraded from Helium-Neon to solid-state. A faulty EOM replaced, and its frequency response was studied by this author (helped by Gustafson) to choose a better modulation frequency where the EOM could supply adequate modulation depth. Unfortunately, locks proved too fleeting to pursue a program of phase camera angular measurements. Much progress nonetheless was made by Dergachev in realizing effective methods of high-speed digitization of the RF-modulated signal in free space [66]. The full program, for investigating a multi-pixel RF sensor with digital (in-software) demodulation, has yet to be realized. If pursued, more sophisticated angular controls might be available for future interferometers.

2.5.3 Advanced observatories and beyond

Advanced LIGO (aLIGO) [33, 86] is intended to improve upon the first-generation LIGO inspiral range tenfold, to 200 Mpc. To accomplish this, it simultaneously must lower the noise floor at the most sensitive frequencies, to a strain noise $4 \times 10^{-24} / \sqrt{\text{Hz}}$, and push the low-frequency wall of seismic noise down, so that frequencies as low as 10 Hz (rather than about 50 Hz) can be probed.

The aLIGO optical systems are far superior to those of initial LIGO. The 10-kg fused silica test mass mirrors have been upgraded to 40 kg for better radiation pressure performance. Fused silica fibers with high mechanical Q suspend the mirrors, instead of steel piano wire; fibers are welded directly to the test mass. Using fused silica in this manner is intended to reduce the suspension thermal noise that dominates in the middle of the LIGO frequency band. Rather than a single pendulum, each test mass sits at the bottom of a four-stage ‘quad’, actively-servoed pendulum chain, reducing seismic noise, in principle, by a factor of $(f/f_0)^8$. The quad pendulum itself is attached to a new in-vacuum seismic isolation table (similar tables hold the auxiliary optics). This multi-stage pendula-and-seismic-isolation scheme is needed to reduce the effect of ground motion, which was the limiting factor for frequencies below about 40 Hz in initial LIGO. Images of the quadruple pendula and the mirror weld can be found on the observatory webpages (*e.g.*, http://www.ligo-wa.caltech.edu/galleries/gallery_092311.html).

These test masses indeed need to be heavier. Laser power has increased from 20 W in Enhanced LIGO to 200 W in aLIGO. Fabry-Perot arm cavity finesse has also increased, raising the stored arm power. The arm cavities can be locked independently of each other and brought into the Michelson interferometer gradually, which should improve ease of use, allow quick recovery from lock-losses, and thereby improve reliable duty cycle, coupled with other system-wide improvements in control architecture. Resonant sideband extraction [149, 86] is one of the most inventive changes, adding an additional signal recycling mirror to the optical configuration, which allows tuning the cavity pole of the Fabry-Perot arms and trading between different levels of high- and- medium-frequency shot noise-limited performance.

Although not part of the reference design, the success of the LIGO Hanford squeez-

ing experiment [4, 62, 74] noted in Chapter IV make quantum optics a likely feature of any enhanced follow-up.

Once commissioned, aLIGO's second-generation design should make possible detections at the rate of 40 per year [11]. Hanford and Livingston have both upgraded their 4 km interferometers and have begun commissioning. The 2-km Hanford interferometer had been intended to be upgraded, but its optics have found a new and likely more useful prospect – part of the growing network of observatories described in Section 2.5.4.

2.5.4 Worldwide network

Gravitational wave observations are new to astronomy. It has long been recognized that multiple observatories with independent confirmation of a signal would be more persuasive evidence of detection than an isolated site. Yet even once the field is established, additional observatories will permit better science. Sky localization of inspiral and burst sources relies on relative arrival times of a GW signal [149]. Multiple, widely separated observatories give the best baseline. Moreover, the data for continuous wave searches from separate sites can often be coherently combined to yield a quieter noise spectrum. LIGO is fortunate thus not to be alone in the pursuit of gravitational wave astronomy.

Optics for a third interferometer will form the core of the nascent LIGO India project [104], which will dramatically improve the accuracy of sky position measurements for gravitational wave transients. In Japan, KAGRA [112] is being built; when completed with sapphire mirrors and cryogenic systems, it should reach comparable sensitivity to Advanced LIGO interferometers. Advanced VIRGO [28] benefits from the first-generation superattenuator's superb seismic isolation performance; from it will be suspended new optics that reflect a more powerful beam. Together, these

second-generation ground-based interferometers should suffice to make direct detection of gravitational waves a reality.

Around the same time as these observatories make first detection, a different technique of GW astronomy, pulsar timing [101], may open up vastly lower frequencies ($\mathcal{O}(10^{-9})$ Hz). Pulsar timing uses radio telescopes to measure the arrival time of pulses, with characteristic wavelengths of order of one light-year. The two techniques should provide complementary information about the gravitational sky, as distinctive as gamma ray telescopes are from radio.

Yet GW astronomy would become a true precision science if a third generation of instruments followed. Third-generation interferometry could take place underground, either in proposed American DUSEL or the European Einstein Telescope [142], potentially with xylophone interferometers tuned to optimize sensitivity in different frequency bands. Prospects in space appear even grander in the long-term, despite short-term setbacks. Following NASA withdrawal from the Laser Interferometer Space Antenna, LISA, in 2011, the European Space Agency has made a concerted push to launch LISA Pathfinder in 2015, with plans for a somewhat-reduced yet very capable eLISA in coming decades [167]. If eLISA launches, it may open the way for measurements of the most elusive gravitational signature: the background of the universe itself, analogous to the cosmic microwave background. Both the Deci-Hertz Gravitational Observatory (DECIGO) [39] and Big Bang Observer (BBO) [98] proposals could, by midcentury, open up a vision of the cosmos in its earliest days.

2.6 Summary

Initial LIGO, during the science run S6, would have been able to see the coalescence of two neutrons stars at about twenty Megaparsecs, out in the Virgo cluster

of galaxies, from sixty-five million years ago, when dinosaurs still walked the Earth. In the first week of aLIGO lock at Livingston, following Memorial Day 2014, aLIGO had a temporal range extending only as far back as when early humans began their diaspora from Africa – a terrestrial parallel to the expansion of the cosmos. Just after Labor Day, aLIGO had exceeded the range of S6. Progress on aLIGO has been swift. When completed, the Hanford and Livingston second-generation interferometers should see back ten times beyond what S6 could, six hundred and fifty million years, to before the Cambrian explosion of life. Perhaps in the third or fourth generation of interferometers, our view of the gravitational sky may stretch back to the age of the observable universe. Even then, we will not have seen all that can be seen. With the two long-range forces of the universe, electromagnetism and gravitation, giving two complementary views of spacetime, we still must build great machines to explore the sights they show, we must understand what we are seeing, and we must propagate that understanding.

This thesis is a prelude to those efforts, from the building of the quantum optical squeezer, and the feedforward regression and continuous waves binary search, to our public interferometer exhibitions. Feedforward regression provides a microcosm of the complexities of gravitational wave interferometry, so there we will begin.

CHAPTER III

Feedforward: Auxiliary MICH-PRC Subtraction

Most of this chapter is drawn from Meadors, Kawabe, and Riles, ‘Improving LIGO sensitivity with feedforward subtraction of auxiliary channel noise’ Classical and Quantum Gravity 31 (2014) 105014 [127], a paper that carries LIGO Document Number LIGO-P1300193.

LIGO, the Laser Interferometer Gravitational-wave Observatory, has been designed and constructed to measure gravitational wave strain via differential arm length measurements. The LIGO 4-km Michelson arms with Fabry-Perot cavities have auxiliary length control servos for suppressing Michelson motion of the beam-splitter and arm cavity input mirrors, which degrades interferometer sensitivity. We demonstrate how a post-facto pipeline (AMPS) improves a data sample from LIGO Science Run 6 with feedforward subtraction. Dividing data into 1024-second windows, we numerically fit filter functions representing the frequency-domain transfer functions from Michelson length channels into the gravitational-wave strain data channel for each window, then subtract the filtered Michelson channel noise (witness) from the strain channel (target). In this chapter we describe the algorithm, assess achievable improvements in sensitivity to astrophysical sources, and consider relevance to future interferometry.

3.1 Introduction

Antennae for gravitational wave observations [165] require precise understanding of noise sources to attain peak sensitivity. Some of these noises arise from auxiliary degrees of freedom in interferometric antennae. Feedforward control can correct these auxiliary control noises. Cluster computing on archived data makes previous methods of feedforward correction scalable to year-long science runs. Computing can also adjust for the non-stationarity inherent in these noise couplings. This chapter describes such a computational method and the improvements it might provide for searches with LIGO (Laser Interferometer Gravitational-wave Observatory).

As a network with GEO600 [180, 100] and Virgo [27], Enhanced LIGO [21, 84] produced data during LIGO Science Run 6 (S6) that was the most sensitive yet taken in the search for gravitational waves of astrophysical origin reaching the Earth: in this chapter, we further enhance LIGO sensitivity via post-run software corrections.

LIGO infers gravitational-wave strain $h(t)$ at each of its two observatories [Hanford, Washington and Livingston, Louisiana] from the length difference between 4-km Michelson interferometer arms [149] using a calibration response function [10]. Each arm contains a Fabry-Perot resonant cavity locked using the Pound-Drever-Hall technique [71, 49], comprised of an input test mass, near the Michelson beam-splitter, and an end test mass. A power-recycling mirror sits between the laser and the beam-splitter. These six core optics form coupled optical cavities with four length degrees of freedom, each of which is servoed to maintain optical resonance by minimizing motion (see Section 3.2). The effective change in the differential arm length L_- (colloquially DARM) caused by gravitational waves is encoded in the intensity of the light reaching the anti-symmetric port of the Michelson interferometer and is read out by

DC homodyne [84]. Auxiliary length control for the beam-splitter and input mirrors is becoming more complex, as in Advanced LIGO, which will add a signal recycling cavity. This chapter describes post-facto software improvements of detector noise using adaptive feedforward subtraction in a pipeline called Auxiliary MICH-PRC Subtraction (AMPS) [116]: these improvements refine LIGO’s gravitational-wave sensitivity to astrophysical sources.

AMPS improves LIGO S6 data (2009 July 07 to 2010 October 20), as this chapter will show. S6 gravitational wave strain (target) is corrected based on auxiliary length noise measurements (witness). Enhanced LIGO generated the S6 data with high laser power and DC readout to prepare for Advanced LIGO. The motion of the beam-splitter and input mirrors of the Fabry-Perot cavities is known [31, 42] to cause cross-talk in the gravitational wave strain channel, which compounds a noise floor fundamentally limited by seismic, thermal suspension, and laser shot noise. Observed S6 cross-talk included differential Michelson (MICH) as well as power-recycling cavity length (PRC). The DARM readout, as explained in Section 3.2, is intrinsically sensitive to MICH divided by a factor of arm cavity gain, Equation 3.10. (Theoretically, physical $h(t)$ is imprinted in MICH, but the cavity gain and relative smallness of the Michelson cavity make the effect about five orders of magnitude smaller than in DARM, so it is ignored). Methods [109] to tune real-time feedforward filters for LIGO servo cross-talk are our starting point, but we seek to automate and improve retuning.

Post-facto, adaptive feedforward simplifies cross-talk subtraction. AMPS uses Matlab 2012a [164]. The witness-to-target transfer function is estimated in discrete time windows of 1024 seconds and fit to a zero-pole-gain filter with Vectfit [68, 96, 95]. Safeguards ensure a statistically significant fit that does not further degrade the

target signal. Noise from the witnesses passes through respective filters, then is subtracted from the strain target channel. The correction lowers the noise floor, benefitting any gravitational-wave searches using this data.

3.2 Description of the feedforward method

Gravitational-wave antennae around the world share features and form a collaborative network. Amongst kilometer-scale Michelson interferometers, GEO600 in Hannover, Germany uses folded arms with both power- and signal-recycling, LIGO and Virgo use Fabry-Perot cavities coupled with power- (and potentially signal-) recycling cavities. The Japanese interferometer KAGRA, under construction, will have a similar optical layout to LIGO and Virgo but with cryogenically-cooled mirrors in an underground laboratory. Although nomenclature here pertains to LIGO, the core problem of this chapter applies directly to all power-recycled Michelson interferometers with Fabry-Perot arms. It could extend to other instruments with multiple degrees of freedom that obtain a signal from a target contaminated by control noise from auxiliary degrees of freedom, especially when those auxiliaries are controlled using a lower signal-to-noise ratio (SNR) error signal than for the target and when the witnesses are highly independent.

LIGO core optics include the beam-splitter (BS) and power-recycling mirror (PRM), which is situated between the laser and the beam-splitter. The four LIGO mirror test masses (TM) are named by arm (X or Y) and input (I) vs end (E) of the Fabry-Perot cavities. LIGO controls four optical pathlength degrees of freedom [87]. DARM is a signal of *differential arm* length, which is calibrated into the primary part of the gravitational strain measurement, $h(t)$. CARM yields *common arm* length, and is controlled with a common mode servo using laser frequency. MICH *Michelson* and

PRC *power-recycling cavity* length refer only to input test masses.

$$(3.1) \quad \text{Strain: } h(t) = \frac{\delta(L_-(t))}{\langle L_+ \rangle},$$

$$(3.2) \quad \text{Common arm length: CARM} \propto \delta(L_+) = \frac{\delta(L_y + L_x)}{2},$$

$$(3.3) \quad \text{Differential arm length: DARM} \propto \delta(L_-) = \delta(L_y - L_x),$$

$$(3.4) \quad \text{Power-recycling cavity length: PRC} \propto \delta(l_+) = \frac{\delta(l_y + l_x)}{2},$$

$$(3.5) \quad \text{(Inner) Michelson length: MICH} \propto \delta(l_-) = \delta(l_y - l_x),$$

$$(3.6) \quad L_y \equiv z(\text{ETMY}) - z(\text{ITMY}),$$

$$(3.7) \quad L_x \equiv z(\text{ETMX}) - z(\text{ITMX}),$$

$$(3.8) \quad l_y \equiv z(\text{ITMY}) - z(\text{RM}),$$

$$(3.9) \quad l_x \equiv z(\text{ITMX}) - z(\text{RM}).$$

Average arm length is $\langle L_+ \rangle$, about 4 km in LIGO. The distance function $z(\mathcal{X})$ indicates the distance (note that $z(\text{ITMY})$ is a function of both the ITMY and BS position), along the optical path, from the laser to an optic \mathcal{X} . The variation δ denotes a change with respect to nominal value. DARM length is thus defined as $\delta(L_y - L_x)$ and MICH length as $\delta(l_y - l_x)$. In practice, DARM and MICH are the names given to the channels that predominantly measure those quantities. Unless stated otherwise, the terms DARM, MICH, PRC and CARM will refer to the measured channels, which are related to the lengths through calibration and are cross-contaminated (e.g., $\text{DARM} = \delta(L_-) + \pi/(2\mathcal{F})\delta(l_-)$, where \mathcal{F} is cavity finesse). The terms will not refer to the ideal physical lengths in Equations 3.2 through 3.5.

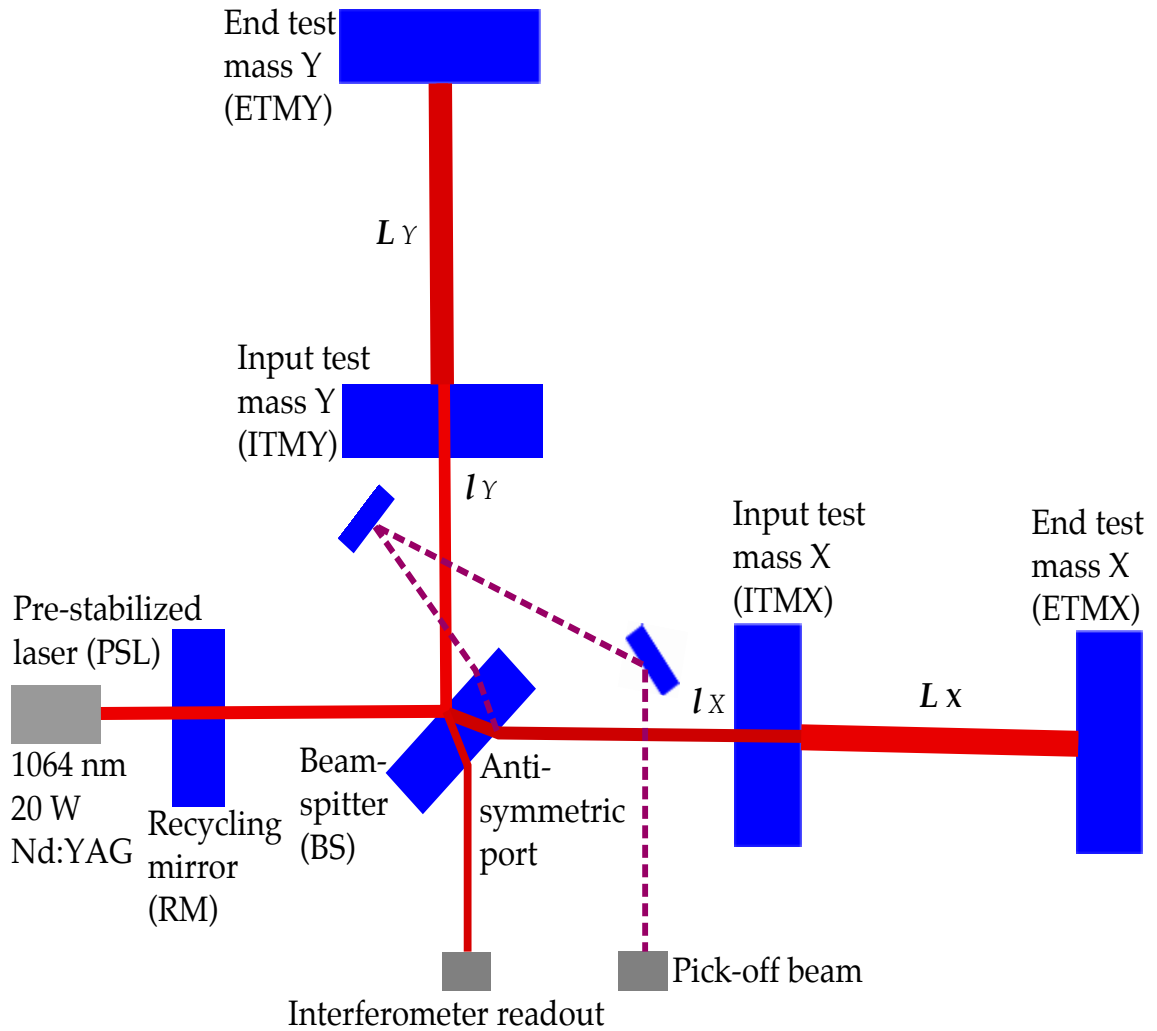


Figure 3.1: Gravitational wave strain $h(t)$ is derived from differential arm motion (DARM), read-out from a photodiode downstream of the antisymmetric port. An internal reflection off an anti-reflective coating, on either the beam-splitter (BS) or an input test mass (ITM), provides the Michelson (MICH) channel. The DARM readout channel predominantly measures the small change in differential arm length, $\delta(L_-) \equiv \delta(L_y - L_x)$, while MICH measures that in the Michelson length $\delta(l_-) \equiv \delta(l_y - l_x)$. There is also a small coupling from $\delta(l_-)$ to the DARM channel. To a lesser extent, changes in the length of PRC, which is defined as $\delta(l_+) \equiv \delta(l_y + l_x)/2$ and is measured in quadrature demodulation with respect to the MICH pick-off, also add noise to DARM.

As Equation 3.3 and 3.5 imply and Figure 3.1 illustrates, MICH noise confuses the physical interpretation of DARM. An arm cavity gain $r'_c/r_c \simeq 139/0.990$, where r_c is the arm cavity reflectivity for the LIGO laser carrier frequency and r'_c is the derivative of r_c with respect to round trip phase [87,42], amplifies DARM motion for Initial and Enhanced LIGO. Where $\mathcal{F} \simeq 219$, the gain is given by Equation 3.10:

$$(3.10) \quad \frac{r'_c}{r_c} = \frac{2\mathcal{F}}{\pi} \simeq (139/0.990).$$

A priori MICH noise will leak into measurements of DARM with a transfer function equal to the inverse of Equation 3.10 [154]. Coherence measurements confirm this coupling dominates the transfer function, but residuals suggest other effects exist. PRC is also indirectly correlated with DARM. These correlations are physical consequences of the interferometer design.

In Enhanced LIGO, DARM is measured with a photodiode at the interferometer ‘dark’ antisymmetric port of the beamsplitter. Independent photodiodes for MICH and PRC, used for feedback on their respective auxiliary length control servos, provide the witness channels for canceling cross-talk into DARM. The MICH and PRC photodiodes receive a beam from an internal reflection in the beam-splitter. This beam carries a radio-frequency modulation; one demodulation quadrature provides MICH, the other PRC.

3.2.1 Auxiliary noise coherence at sensitive frequencies

Coherence, the Fourier frequency-dependent analog of statistical covariance, quantifies cross-talk. On a scale of 0 (none) to 1 (full), magnitude-squared-coherence (MS-coherence) represents the normalized fraction of power of a frequency bin in the spectrum of one channel that can be found in the same frequency bin in the spectrum

of another channel. First, we must define the cross-power of two time-series. Where P_{xy} is cross-power spectral density, we can describe how the coherence at a given frequency f and time t [176] is given by Equation 3.11.

$$(3.11) \quad C_{xy}(f, t) = \sqrt{\frac{|P_{xy}(f, t)|^2}{P_{xx}(f, t)P_{yy}(f, t)}}.$$

The calibrated strain channel for $h(t)$ (internally, the discrete-time calibrated strain channel for the physical strain $h(t)$ is called ‘Hoft’), is, with high confidence, coherent with MICH and PRC, as seen in Figure 3.2. MICH- $h(t)$ coherence is sometimes as large as 0.1 in the 100 to 300 Hz band; PRC- $h(t)$ is an order of magnitude lower. Unfortunately, this is the most sensitive band for Initial and Enhanced LIGO. Because this band is so sensitive, it generates much of LIGO’s scientific value. Removing a fraction of 0.01 to 0.1 of the noise is useful, particularly when doing so demonstrates a technique applicable to future cases with potentially-worse noise.

3.2.2 Estimating filters

Allen, Hua, and Ottewill (AHO) [37] proposed the filtering scheme that this chapter employs. Where there is a strong correlation between a signal (target) channel and a noise (witness) channel, the noise can be partially cancelled if a witness-to-target transfer function, convolved with the witness, is applied to the measured target. Equations 3.12 through 3.13 capture this method. It is analogous to frequency-domain principal component analysis (PCA) using Gram-Schmidt orthogonalization. In the original theory, superscript ^(b) indicates a frequency band that we denote as domain (f). Equation 8 in AHO corresponds to Equation 3.13 here.

The transfer function T_{xy} , from the cross-power ratio of arbitrary channels x and y , guides the estimated feedforward filter g . Figure 3.3 shows the fit to the transfer function. Viewed as an inverse Fourier transform F^{-1} , decoupling signal (target,

subscript s) from noise (witness, subscript n):

$$(3.12) \quad g(t) = F^{-1}(\text{fit}[T_{sn}(f)]).$$

Finally, the post-filtering signal (target) \hat{s} is given by the convolution (\times) with γ , the transfer function coupling noise (witness) into signal (target), s pre-filter signal, n noise, and with channels indexed by j and curly brackets indicating an observable quantity:

$$(3.13) \quad \hat{s}(t) = \{s + \Sigma_j (\gamma_j \times n_j)\}(t) - \Sigma_j (g_j(t) \times \{n_j\}(t)).$$

Blind application of this method could produce incorrect noise reduction. Application of this chapter's method to uncorrelated channels would lead to arbitrary noise reduction by an average analytic factor of $(1 - 1/F)$, where F is the number of bins in a fitted frequency span (equal to the the number of time-domain averages). Given 1-s windowing with 50%-overlap on 1024 s, $F = 2047$, for a false noise reduction of about 0.05%. The ideal of 1024-s windows is not always achievable with LIGO duty cycles. In these cases, AMPS incorporates some filters estimated on as little as 32 s of data, for which the reduction would be 3%, but only when these filters are averaged together with longer-duration (512 s or greater) filters. No isolated filter uses less than 60 s of data, which could yield a false reduction of 1.5%. AHO clarify that subtraction is tenable so long as covariance is present at a statistically significant level. They set a benchmark of an order-of-magnitude above the magnitude-square covariance expectation value of $1/F$. Since the AMPS pipeline emphasize fits in regions where the magnitude-squared coherence is greater than 3%, and often 10% or more, it usually satisfies their criterion.

As detailed in Section 3.5.1, filters for each 1024-s window are blended to estimate the target. Figure 3.4 illustrates the short-term consistency of transfer functions dur-




Figure 3.2: Sample coherence measurements between $h(t)$ and auxiliary control channels for LIGO Hanford Observatory, H1: 2010 March 21. MICH- $h(t)$ coherence on left, PRC- $h(t)$ coherence on right. Statistically significant coherence justifies fitting; in frequency bands, about 80 to 400 Hz, where coherence rose above background levels, the transfer function fit was weighted more heavily. Units of coherence spectral density ($\text{Hz}^{-1/2}$) vs frequency (Hz).

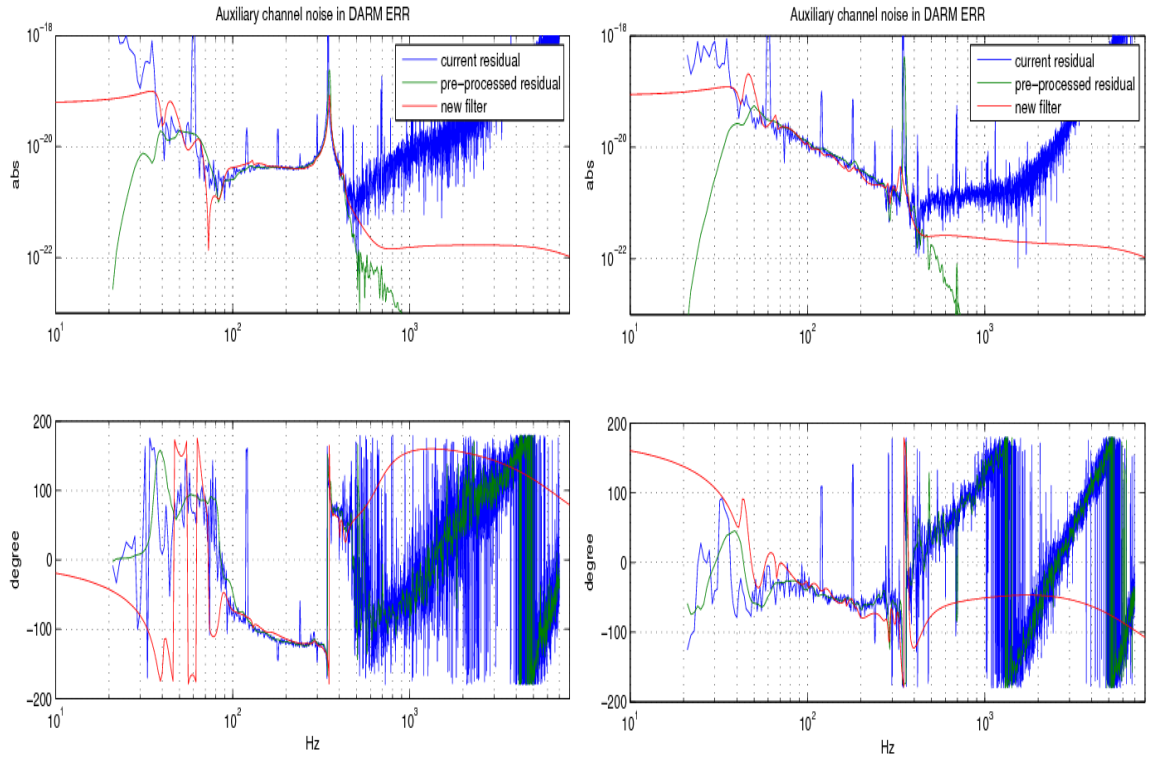


Figure 3.3: Sample transfer function measurements (amplitude and phase) from LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ on left, PRC- $h(t)$ on right. Transfer function fit in coherent band – note the difference between raw data residual and the ‘pre-processed residual’, which has been smoothed and weighted to emphasize known-coherent bands. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) and phase (degrees) vs frequency (Hz).




Figure 3.4: Sample Bode plots of fitted ZPK filter functions (amplitude and phase) for multiple 1024 s windows in a science segment, at LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ on left, PRC- $h(t)$ on right. Colors only represent different time windows. The similarity in the high-coherence, 80 to 400 Hz band leads us to conclude that the filter design is fairly stable throughout a science segment. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) and phase (degrees) vs frequency (Hz).

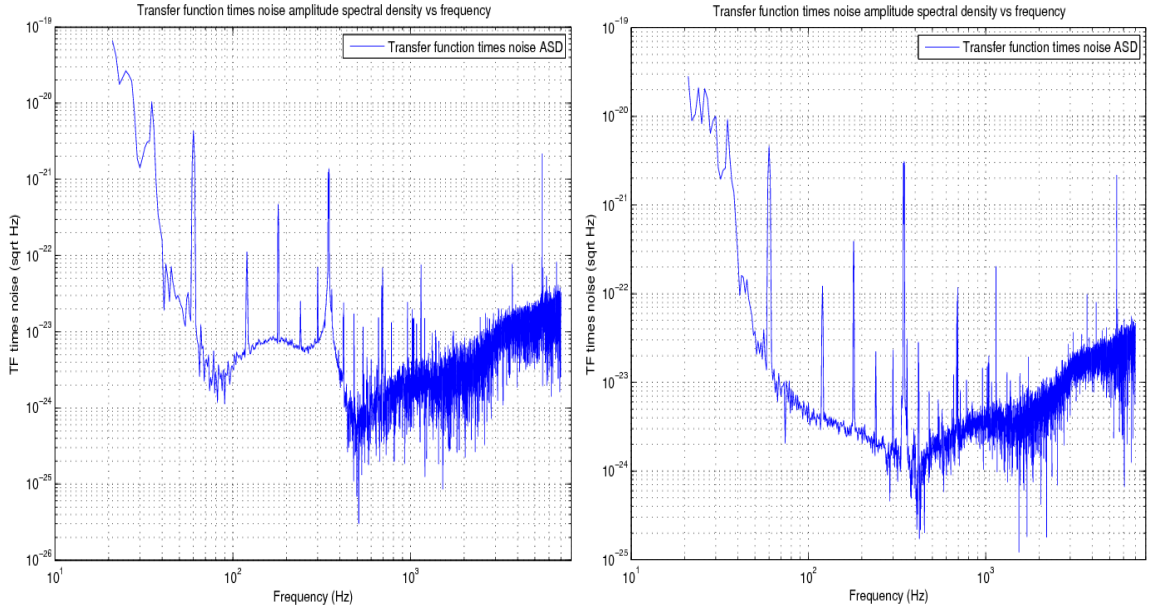


Figure 3.5: Sample subtracted spectra for one window, representing the applied feedforward corrections for each channel during that window, at LIGO Hanford Observatory, H1: 2010 March 21; MICH- $h(t)$ correction on left, PRC- $h(t)$ correction (after MICH- $h(t)$ correction is applied) on right. Units of amplitude spectral density ($\text{Hz}^{-1/2}$) vs frequency (Hz).

ing a few-hour science segment, presenting Bode plots of MICH- $h(t)$ and PRC- $h(t)$ transfer function fits for consecutive windows from 2010 March 21 at LIGO Hanford Observatory (magnitude, top and phase, bottom vs frequency [Hz]). Figure 3.5 presents sample subtractions. These figures show that AMPS efficiently builds on AHO for operational data for a kilometer-scale gravitational wave interferometer.

Section 3.3 compares transfer function estimators, with Section 3.3.2 being the chosen method. Section 3.4 discusses safeguards and vetoes, and Sections 3.5 and 3.5.1 discuss the details of implementation and verification.

3.3 Feedforward in- and out-of-loop methods

Feedforward subtraction must meet operational constraints. Existing manually designed filters have long worked, but are more labor-intensive than automated de-

sign; new methods, such as Wiener filtering now being considered for seismic and gravity-gradient cancellation [72] could lead to future improved performance.

3.3.1 Manually designed rational filtering in-loop

Manual designs of feedforward functions prove time-consuming. Transfer functions must be manually measured, fit, copied and incorporated into the control system. Additional transfer functions are needed for servo in-loop gain and actuation functions. Manual design is an inefficient choice: it is labor-intensive, and S6 suggests that filter redesign should be performed often.

While involved, manual-designed rational filtering of MICH and PRC in-loop provides a key part of servo controls to date. Auxiliary controls introduce noise into the DARM channel, so without real-time correction, the performance would be much worse than design. Most MICH & PRC subtraction so far comes from real-time corrections; our pipeline makes one to two orders of magnitude smaller corrections.

3.3.2 Frequency-domain automated filter design

AMPS uses Vectfit for periodic re-design. Since the dynamic range in magnitude for transfer functions varies over tens of orders of magnitude, data is pre-processed and weighted to emphasize the most sensitive band. The method fits an infinite impulse response (IIR) filter onto the witness-to-target transfer function. Since coherence and AHO are linear and transitive, it targets $h(t)$ rather than DARM (noise coupling enters the signal there, but it is wasteful to duplicate the response function).

Each transfer function for a typical 1024 s of data is the average of 1024 independent ratios-of-Fourier-transforms (2047 Hann-windowed, 50%-overlapping FFTs of 1 s samples). Since FFT error scales with the inverse square root of the number of averages, the relative accuracy is $\mathcal{O}(1024^{-1/2})$. The minimum data length, 32 s,

yields $\mathcal{O}(32^{-1/2})$ relative accuracy. Outside the sensitive band, the fit is deweighted and the transfer function pre-processed, suppressing it by factors of $(f/f_{\text{knee}})^\alpha$, where $\alpha = 8$ at low frequencies and -8 at high. The f_{knee} values are, respectively, 50 and 400 Hz. AMPS smooths and deweights (Figure 3.3) known spectral peaks, including 60 Hz harmonics, the LIGO suspension violin modes, and calibration lines. Violin modes are internal resonances of mirror suspensions caused by thermal noise; calibration lines are injections used to track the response function. De-weighting and pre-processing prevent biasing the filter design with transfer function bands where coherence is low, which would introduce noise. This process leads to convergence with fewer parameters.

Vectfit converges iteratively, starting with a posited set of poles (32nd order here). About five iterations can converge to a good least-squares-fit for the state-space model, but we require fifteen iterations and complex left-half-plane stability for safety. Root-mean-square (RMS) error, is the threshold for rejecting the filter regression. From empirical studies, RMS error above 10^{-18} indicates poor fit. This test isolates a bad MICH- $h(t)$ correction from a good PRC- $h(t)$ one, or vice versa. We fit one channel at a time, as Section 3.5 discusses. After this test, the transfer function model is extracted.

Zero-pole-gain (ZPK) format is used to trim out-of-band zeroes and poles and multiply by a 2nd order Butterworth low-pass filter just below the Nyquist frequency of 8192 Hz, placing poles at 7 kHz to keep causality. A scale factor keeps the filter gain at 150 Hz the same value as without the low-pass filter. Then the ZPK model is refactored into second-order-sections (SOS) for numerical stability. Instead of the inverse Fourier transform of Equation 3.12, the model is converted from continuous time (or s -domain) to discrete time (or z -domain).

Each filter is applied to its respective witness: the estimated true $h(t)$ target equals the original $h(t)$ measurement minus the filtered witnesses. This procedure assumes that coupling from each witness into $h(t)$ is linear. Further, it assumes that 2nd-order coupling, from one witness into the other, is negligible (we estimate the relative contributions to be $\mathcal{O}(10^{-5})$). Spot-checks confirm that these simplifications are justified.

3.3.3 Wiener filters

Wiener filtering [177] would give an optimal filter that minimizes the squared error of the residual, for all the spectrum. Low-frequency MS-coherence of MICH & PRC with $h(t)$ is small, but Wiener fits them due to high RMS error in that band. Filtering at high RMS power, such as the seismic and Newtonian gravity gradient bands, can allow Wiener filtering directly, but MICH & PRC would need other methods. Noise whitening [73, 67] uses cost functions to limit out-of-band noise. Wiener filtering sub-spectra could also circumvent contamination, as with wavelet transforms [110].

3.3.4 Prospects for near-real-time filtering

AMPS runs a few times faster than real-time on a single 2013 CPU core whilst conducting tests and safeguards, documented in Section 3.4. The minimal time lag for a modal sample is one window (1024 s), undesirable for electromagnetic follow-up and multi-messenger astronomy [77, 32] sought for Advanced LIGO. Such speed is acceptable for secondary $h(t)$ -reconstruction when another $h(t)$ exists but inadequate for in-loop, real-time production.

Near-real-time filtering might be useful for countering upconversion and non-linear cross-coupling, using recent data for training sets, but this is not yet implemented. Until then, the existing method can generate a filter as-needed for real-time use, as

prototyped on H1 in the last month of S6.

3.4 Safeguard and veto methods

3.4.1 Calibration integrity

It is of vital importance for our noise subtraction scheme to keep the integrity of the calibration of LIGO instruments intact. If the witness channel contains a cross coupled term proportional to the differential arm length motion $\delta(L_-)$, this term is subtracted from the target and thus could change the calibration of the target in theory. A simple calculation shows that, for a known coupling mechanism, this effect is on the order of 10^{-5} for LIGO and other similarly configured instruments.

Even without considering any feedback control mechanism, theoretically MICH and DARM signal both have cross contamination terms proportional to $\pi/2\mathcal{F}$ [154]:

$$(3.14) \quad \text{DARM} \propto \delta(L_-) + \frac{\pi}{2\mathcal{F}}\delta(l_-),$$

$$(3.15) \quad \text{MICH} \propto \delta(l_{-,0}) + \frac{\pi}{2\mathcal{F}}\delta(L_-) + n,$$

where \mathcal{F} is the finesse of the arm cavities, n the sensing noise of the MICH, and $\delta(l_{-,0})$ the natural fluctuation of the Michelson path difference caused by seismic motion etc. Since $\delta(l_{-,0})$ is not coherent with $\delta(L_-)$, and since MICH is dominated by n in our frequency band of interest, we can ignore $\delta(l_{-,0})$ in this discussion.

The servo system with an open loop transfer function of G_M to keep MICH from going out of linear range would inject $\frac{\pi}{2\mathcal{F}}\delta(L_-) + n$ term to the physical Michelson length difference $\delta(l_-)$, which in turn affects the DARM signal:

$$(3.16) \quad \delta(l_-) = -\frac{G_M}{1+G_M} \left(\frac{\pi}{2\mathcal{F}}\delta(L_-) + n \right)$$

$$(3.17) \quad \text{DARM} \propto \delta(L_-) - \frac{\pi}{2\mathcal{F}} \frac{G_M}{1+G_M} \left(\frac{\pi}{2\mathcal{F}}\delta(L_-) + n \right).$$

The DARM signal, uncorrected by feedforward, is now contaminated by MICH noise

term n as well as a small correction term for $\delta(L_-)$, both due to the MICH feedback. Note that $G_M/(1 + G_M)$ is on the order of 1 or smaller, and $(\pi/2\mathcal{F})^2$ is on the order of 10^{-5} , so the correction term is on the order of 10^{-5} or smaller.

Feedforward subtraction looks at the MICH control signal,

$$(3.18) \quad \text{MICH}_{ctrl} \propto \frac{\pi}{2\mathcal{F}}\delta(L_-) + n,$$

and subtracts n from DARM, and in the process also subtracts $\frac{\pi}{2\mathcal{F}}\delta(L_-)$. Looking at Eqs. 3.17 and 3.18, as far as the noise reduction is observed in DARM, the small $\delta(L_-)$ term is also reduced, and the impact on the calibration is on the order of $(\pi/2\mathcal{F})^2 = \mathcal{O}(10^{-5})$ at most.

A similar argument can be made for PRC, but this time the cross coupling is not only dependent on $(\pi/2\mathcal{F})$ but also on the asymmetry of the arms, further reducing the coupling.

Note that the above mentioned discussion is equally applicable to real-time as well as post-facto feedforward. LIGO uses real time feedforward, and the offline feedforward described in this chapter subtracts only a small amount of noise left uncaught by the real-time system. Virgo implements similar real-time feedforward to remove the 50Hz line successfully [53]. In both of these cases, no measurable effect caused by feedforward has been reported.

Nevertheless, two checks were performed to see if there is any unknown mechanism to compromise the calibration of DARM, which are explained in Section 3.4.3.

3.4.2 Runtime safeguards

Safeguards and vetoes then verify data integrity against possible issues. These issues include degrading data, offsetting and incorrectly time-stamping the data, incorrectly subtracting $h(t)$ from itself, and introducing windowing artifacts.

Amplitude spectral density (using Welch’s method [175]) leads to an estimate of inspiral range \mathcal{R} . Inspiral range [80] in LIGO detector characterization refers to the orientation-and-direction-averaged distance at which a 1.4-1.4 solar mass neutron star binary coalescence could be detected with an SNR of 8 (for further details, see Section 4.2.2; the range is given precisely by Equation 4.6). A window’s filtered data is used only if it passes two cuts. The post-filter \mathcal{R} must be at least 99.9% of unfiltered \mathcal{R} . None of the 40 points in a ‘comb’ (each point being 5/16 Hz wide) of quiet bands can be noisier than 1.2 times uncorrected $h(t)$. The factors are chosen empirically to permit expected noise fluctuations; most surviving data is superior.

If cuts are triggered, the filter is rejected. To avoid discontinuities and add robustness against non-stationarity, the windowing procedure (Section 3.5.1) is re-run to merge successfully-filtered data smoothly with unfiltered h_0 . Cut tests are also re-run; if passed, the data is used, else the unfiltered data is progressively weighted further for eight more attempts. If all fail, the final attempt is written and the program proceeds. Empirically, almost all written data is an improvement.

3.4.3 Post-processing safeguards

Diagnostics check whether calibration lines (Figure 3.6) are preserved and injections (Figures 3.7 and 3.8) are recovered.

Calibration line studies seek to answer two questions: whether feedforward distorts the signal or adds noise.

In post-processing, ‘Short Fourier Transforms’ (SFTs) were made with a frequency resolution (1/1800 Hz) from corrected $h(t)$. These SFTs are much shorter than the science run; 1800 s is standard for continuous wave searches. Signal distortion is evaluated using the mean of three bins in $[393.1 - 1/1800, 393.1 + 1/1800]$ Hz. Evaluating multiple bins accounts for some spectral leakage; the bin-centered central line is much

larger. For the 10^6 s of H1 science time analyzed, the before-feedforward mean was $8.7261 \times 10^{-22}(\text{Hz})^{-1/2}$, whereas after it was $8.7569 \times 10^{-22}(\text{Hz})^{-1/2}$. Feedforward made the calibration line region noisier by $3.1 \times 10^{-24}(\text{Hz})^{-1/2}$ or 0.35%, perhaps due to a deweighted fit around the 393.1 Hz line, to avoid biasing more sensitive parts of the spectrum. It does not affect the calibration, since it is consistent with MICH and PRC leakage merely raising noise floor near the line. We also check the calibration line at 1144.3 Hz (before: $3.1190 \times 10^{-20}(\text{Hz})^{-1/2}$, after: $3.1188 \times 10^{-20}(\text{Hz})^{-1/2}$), which is actually less noisy. The 46.7 Hz line is too low-frequency to measure with these SFTs.

To test for noise addition, we searched for windowing artifacts, e.g., spectral combs with spacing of 1/1024 or 1/512 Hz, around a prominent line. No new combs or other artifacts were obvious in our 1800-s, 50%-overlapping Hann-windowed SFTs before/after comparison of approximately 10^6 s of H1 science time between GPS times 931.0×10^6 (2009 July 07) and 932.8×10^6 (2009 July 28), focused on the 393.1 Hz calibration line. Strictly speaking, the line visible in $h(t)$ is a residual from imperfect cancellation of control and error signals used in $h(t)$ construction from DARM error and control signals. The nature of the line does not affect our analysis, because neither MICH nor PRC contain or affect it.

Injection studies first examined compact binary coalescence and sine-Gaussian injections [13] for GPS seconds 931.0×10^6 (2009 July 07) to 932.8×10^6 (2009 July 28). We then calculated the matched-filter SNR of each injection in S6. Each SNR is directly proportional to the distance at which such a signal *can* be observed, and therefore also to the instrumental sensitivity to signals of that type [34, 140]. Higher SNR (mean 3.99% H1, 2.77% L1) was found. Effective distance as *recovered* is inversely proportional to signal, appeared nearly unchanged in these injections




Figure 3.6: Calibration line test: before-feedforward mean of the 393.1 Hz line and two neighboring FFT bins was 8.7261×10^{-22} , after was 8.7569×10^{-22} . Feedforward made the calibration line region noisier by 3.1×10^{-24} or 0.35%, suggesting that we correctly apply Hann-windowed feedforward without subtracting true $h(t)$. Moreover, no spectral line combs are observed to either side of the calibration line peak at 393.1 Hz, indicating that the method does not introduce windowing artifacts.

(mean -0.00347% H1, +0.307% L1), establishing that the SNR increase came from N decreasing rather than S increasing. The constant injection effective distance reinforces that the calibration is unchanged. These tests affirmed that feedforward data contains recoverable (slightly higher SNR ratio) injections at the correct time and phase.

We conclude that we are not subtracting $h(t)$ from itself.

3.5 Feedforward with MICH and PRC channels

A post-facto, linear filter is fitted and applied to either MICH or PRC (serially). Fits occur in the frequency-domain, and application occurs in time-domain. Reading in $h(t)$, correcting it with MICH and PRC, and writing the result (including data quality and state vectors), the pipeline is shown in Figure 3.9.

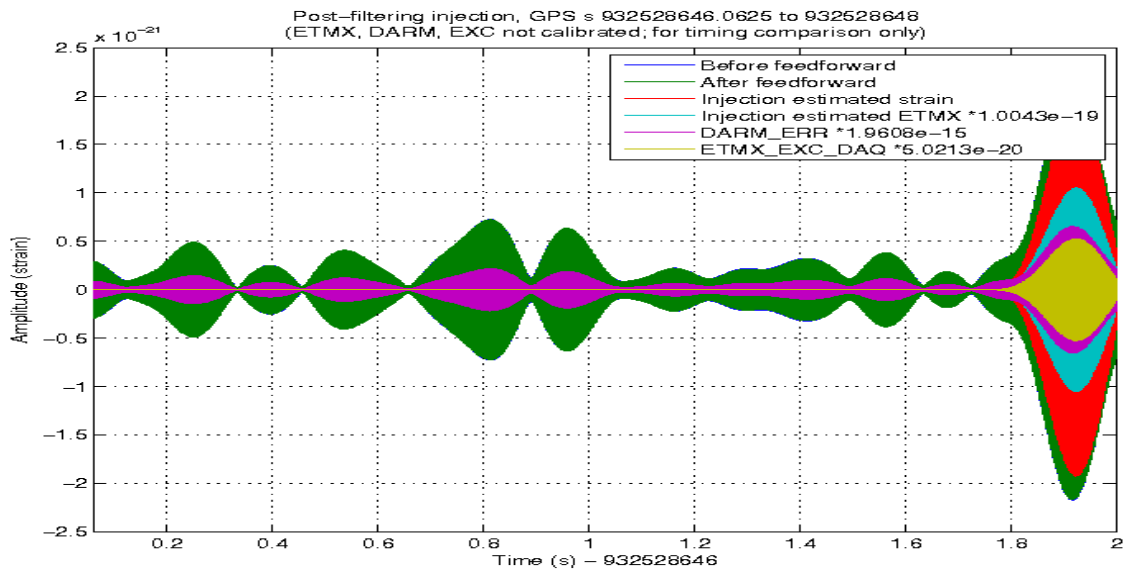


Figure 3.7: Time-domain plot of diagnostic channels from a burst injection. Colors are illustrative only to the fact that the envelopes of the traces increase after 1.8 s, indicating that the burst injection time is correct in the new data. ‘Before feedforward’ and ‘after feedforward’ traces occult each other in the graph, because they are almost identical. ‘Before feedforward’ is $h(t)$ data; ‘after feedforward’ is $h(t)$ with feedforward subtraction. ‘Injection estimated strain’ is the digital injection as intended to be introduced into strain, but the actual injection is made on the end test mass X (ETMX), so the calibrated ‘Injection estimated ETMX’ is also displayed. Raw ‘DARM_ERR’ and ‘ETMX_EXC_DAQ’ are redundant but reinforce the trend.

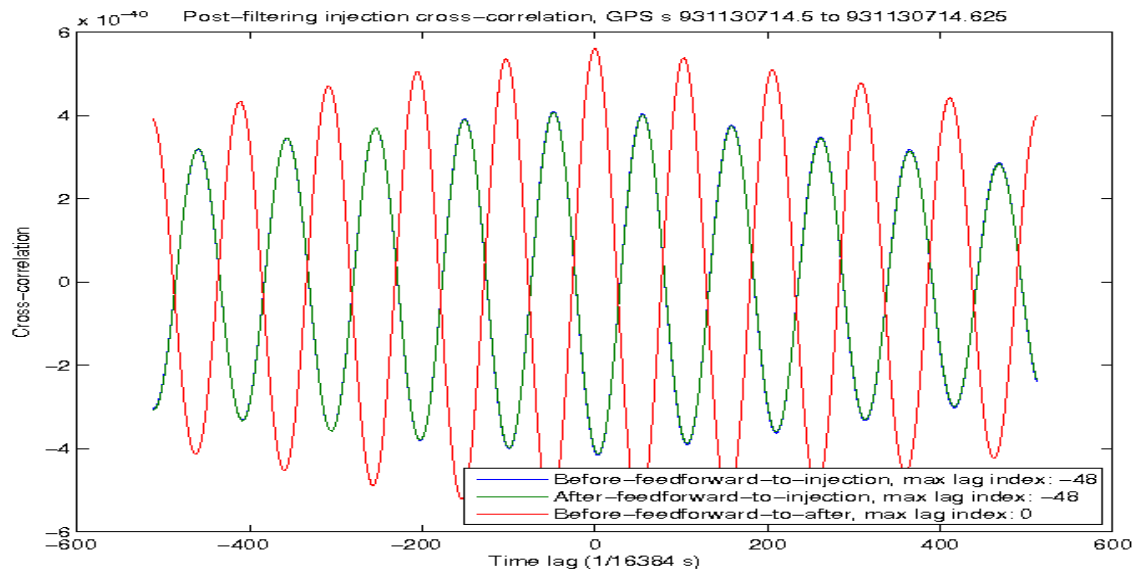


Figure 3.8: Cross-correlation pairwise between $h(t)$ pre-, post-feedforward, and ETMX injection data: the extrema and zero-crossings match. Note both before-feedforward (blue) and after-feedforward (green) strain traces are almost identical and therefore overlap. The strains appear inverted, but in the same way, due to a sign error in the hardware injections at this time. The absence of a time lag shift between before and after indicates that feedforward has not altered the phase of the data, at least for this injection. The equivalence in cross-correlation magnitude indicates that amplitude also is unaffected.

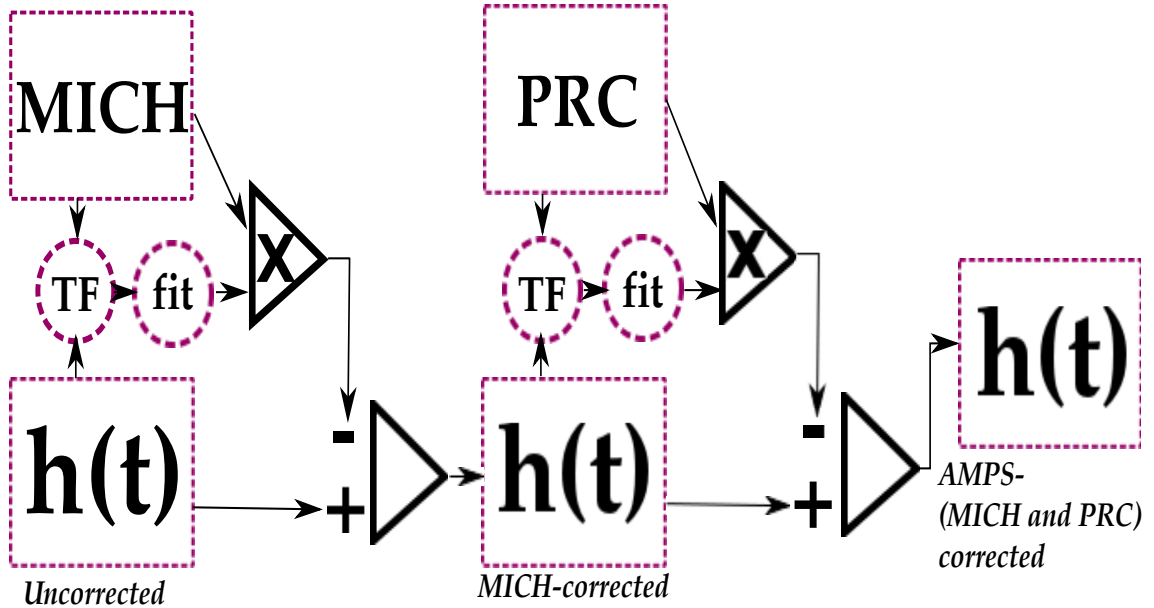


Figure 3.9: Feedforward subtraction pipeline to read in $h(t)$, MICH, PRC, and write out AMPS-corrected $h(t)$. Data flows schematically from left to right; the MICH- $h(t)$ stage output is used as input for the PRC- $h(t)$ stage, then data is written. Code online: <http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/detchar/AMPS/trunk/aletheia.m>

3.5.1 Filter fitting across science segments

We run one job process per science segment. Segments range in duration from seconds to days, with a median of hours. The interferometer is locked during each segment, meaning it is held fixed on a fringe, by servo control. Lock loss or noise degradations can define a segment end. Segments are divided into 50%-overlapping Hann windows, filtered, and smoothly re-merged. Windowing is idempotent for $h(t)$ itself; the difference from window to window is the correction added. The first 512 s derive only from the first window; every 512 s afterward, a new window commences, as in Figure 3.10.

Using Equation 3.13 with filters g , target $S = \{s + \sum_j (\gamma_j \times n_j)\}$ and witness $N_j = \{n_j\}$, we can evaluate $\hat{s}(t)$. Since the filters for different channels are calculated in series, with transfer functions T , Equation 3.19 has $g_1 \sim T_{S,N_1}$ but $g_2 \sim T_{(S-g_1 \times N_1),N_2}$.

Here, S , N_1 and N_2 are respectively $h(t)$, MICH and PRC.

$$(3.19) \quad \hat{s}(t) = S(t) - \Sigma_j (g_j(t) \times N(t)),$$

$$(3.20) \quad = S(t) - g_1(t) \times N_1(t) - g_2(t) \times N_2(t),$$

$$(3.21) \quad \sim S(t) - F^{-1}\text{fit} [T_{S,N_1}] \times N_1(t) - F^{-1}\text{fit} [T_{(S-g_1 \times N_1),N_2}] \times N_2(t).$$

Since $N_1(t)$ and $N_2(t)$ are added linearly to $S(t)$, we can analyze them independently. Analyze the first two terms of Equation 3.20 and take $N(t) = N_1(t)$. Let g_A and g_B be the earlier and later filters for $N(t)$ being time-domain merged in a Hann-window; they are respectively calculated from overlapping data sets $[S_A, N_A]$ and $[S_B, N_B]$. The sets are identical at time t , so $S(t) = S_A(t) = S_B(t)$, $N(t) = N_A(t) = N_B(t)$. Windowing merges data streams \hat{s}_A and \hat{s}_B over $\tau = 1024$ s, per Equation 3.22:

$$(3.22) \quad \hat{s}(t) = \frac{\hat{s}_A(t)}{2} \left[1 - \cos \frac{2\pi(t + \frac{\tau}{2})}{\tau} \right] + \frac{\hat{s}_B(t)}{2} \left[1 - \cos \frac{2\pi(t + \tau)}{\tau} \right],$$

$$(3.23) \quad = \frac{1}{2} \left(\hat{s}_A(t) + \hat{s}_B(t) + \cos \frac{2\pi t}{\tau} [\hat{s}_A(t) - \hat{s}_B(t)] \right),$$

$$(3.24) \quad = \frac{2S(t) - (g_A + g_B) \times N(t)}{2} - \frac{g_A - g_B}{2} \times N(t) \cos \frac{2\pi t}{\tau},$$

$$(3.25) \quad = S(t) - \frac{1}{2} \left(g_A \left[1 + \cos \frac{2\pi t}{\tau} \right] + g_B \left[1 - \cos \frac{2\pi t}{\tau} \right] \right) \times N(t).$$

Equation 3.25 shows that the windowing process equates to evolving filter coefficients with a 512 s cadence. Substitute $\hat{s}(t)$ into $S(t)$ with $N(t) = N_2(t)$ for the next noise channel to extend the result.

Science segments are subdivided into at most 16384 s. These subdivisions overlap for 512 s, so each side calculates identical filters for the overlap, but only the latter half writes the overlap, to avoid race conditions. Label g_W, g_X, g_Y, g_Z the final filters

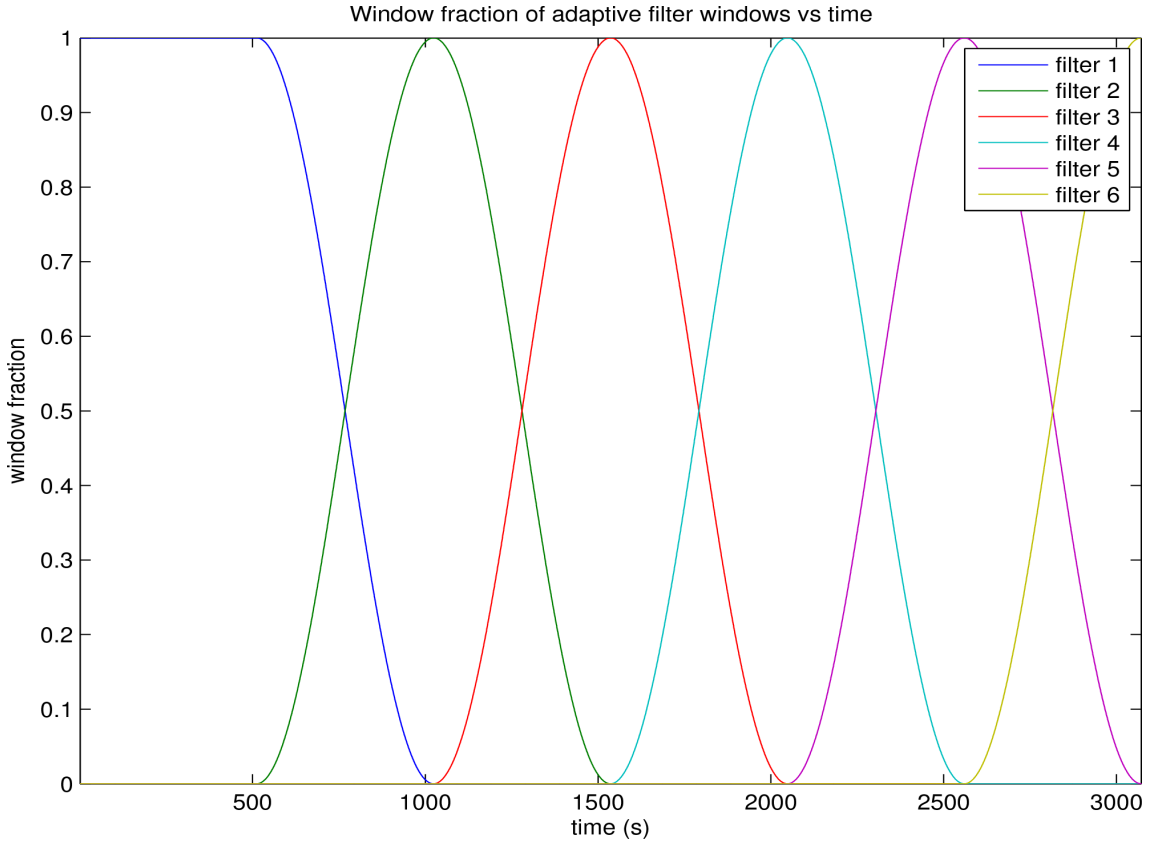


Figure 3.10: Schematic windowing for one LIGO science segment, illustrating windowing after an initial half-window offset. Filters are calculated for windows up to 1024-s, then 50%-overlapping Hann windows merged, giving a corrected measurement of $h(t)$. Code online: <http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/detchar/AMPS/trunk/eleutheria.m>

calculated in job 1; g_A, g_B, g_C, g_D are the first in job 2. Where each parenthesis contains 512 s and the addition sign denotes Hann-windowing of the filters, the end of job 1 is $\dots(g_W + g_X)(g_X + g_Y)(g_Y + g_Z)$ and the start of job 2 is $(g_A + g_B)(g_B + g_C)(g_C + g_D)$. Overlap denotes that filter g_A is derived from the same data as filter g_Y , and likewise $g_B \simeq g_Z$. Thus $(g_A + g_B) = (g_Y + g_Z)$. Segments shorter than 60 s are not filtered, and dangling windows shorter than 32 s are rolled into their predecessors. These provisions prevent filters based on insufficient data.

3.6 Results of feedforward

3.6.1 Post-processing diagnostics

Lower spectral noise floors reveal improvement in Figures 3.11 and 3.12. Feedforward most improves the spectra with elevated noise levels. It enhances sensitivity less when the interferometer is already optimized. This tendency is consistent with underlying thermal and shot noise. Glitches contaminate $h(t)$ less when the servo-to-strain coupling is minimized. Insofar as coupling non-stationarity evolves slower than the 512-s windowing timescale, adaptive filtering appears to reduce the impact of glitches.

Post-processing tests also compute average SFT spectra. For GPS second 931.0×10^6 (2009 July 07; GPS seconds count from 1980 January 01) to 932.8×10^6 (2009 July 28, about 10% of S6), the harmonic mean spectra are shown in Figure 3.13.

SFTs are high-pass filtered at 38 Hz. The harmonic mean spectrum shows several percent improvement from about 80 Hz up to the 330 Hz violin mode frequencies. Above 400 Hz, there is proportionally minor degradation, due to filter rolloff.

3.6.2 Feedforward benefits and potential

Inspirational range \mathcal{R} increases for both S6 LIGO observatories, which should generalize to any observatory with broadband noise due to contamination from auxiliary servos.

Figures 3.14 and 3.16 show the variation in achieved subtraction over about 10% of S6; figures 3.15 and 3.17 show the entire science run.

Figure 3.18 shows a screenshot of a webpage where LIGO data analysts can access summary graphs of the feedforward performance.

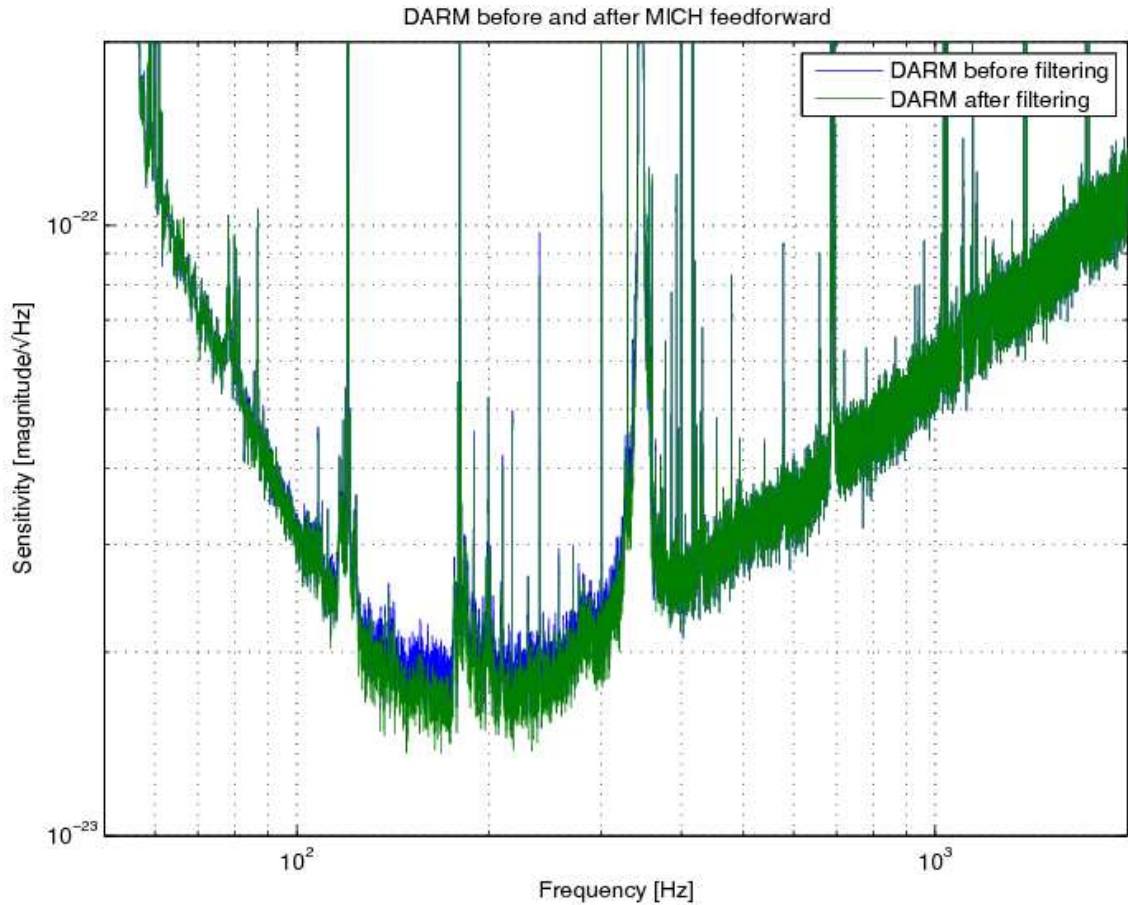


Figure 3.11: Exemplar of a typical case, +1.1 Mpc (5.9% inspiral range) (*GPS time 953164819 to 953165839, 2010 March 21*). Read ‘DARM’ as $h(t)$, ‘MICH’ as ‘MICH-PRC’. The most benefit is seen in the 80 to 400 Hz band, especially around 150 Hz, where LIGO is most sensitive. The main fundamental limit in this band is thermal suspension noise, but historically auxiliary channel noise has been a major contaminant. Note that the 60 Hz and harmonic lines are not subtracted, although a separate magnetometer servo does reduce their impact. The 340 to 360 Hz violin mode is not likely amenable to feedforward.




Figure 3.12: Best improvement seen in S6 for H1 $h(t)$, +4.4 Mpc (29% inspiral range) (*GPS 955187679 to 955188191, 2010 April 13*). Read ‘DARM’ as $h(t)$, ‘MICH’ as ‘MICH-PRC’. Such a loud cross-coupling would be noticed in real-time by the on-site staff. The elevated noise floor is unusual in science mode, but the fact that feedforward corrects it suggests the importance of controlling auxiliary channels to prevent such glitches. The post-feedforward spectrum is comparatively normal for science mode

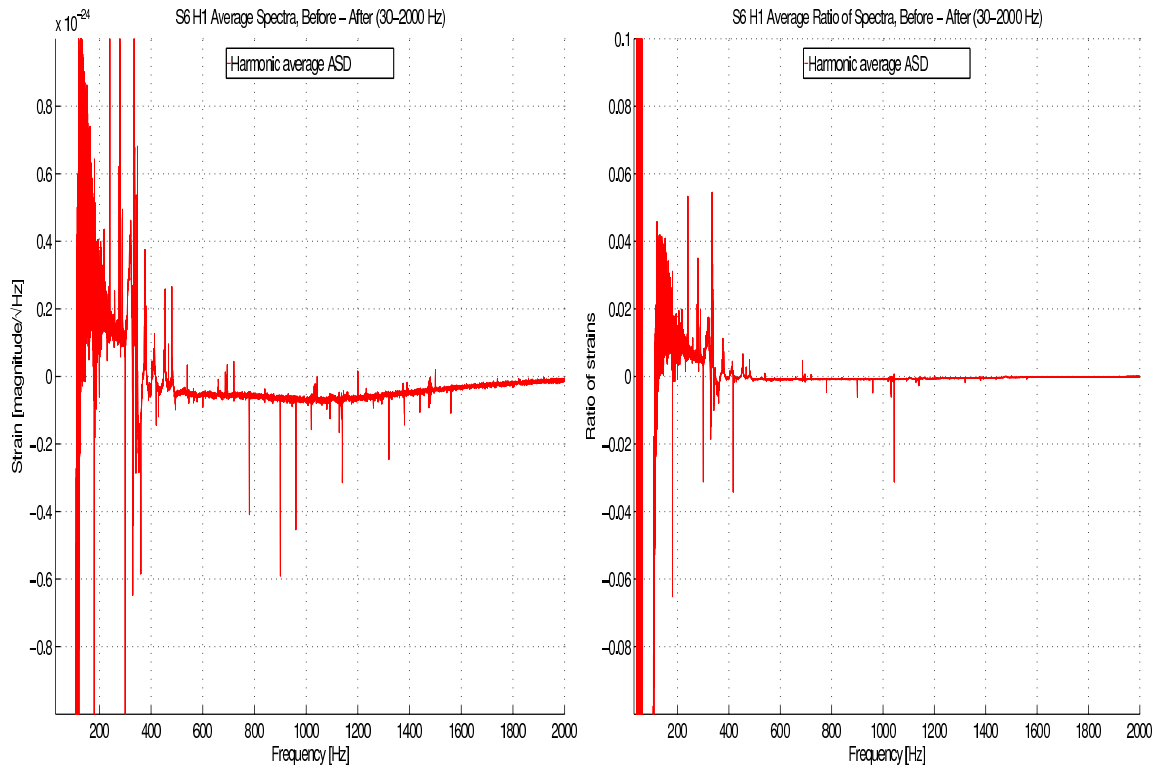


Figure 3.13: Harmonic mean, GPS seconds 931.0×10^6 (2009 July 07) to 932.8×10^6 (2009 July 28): (*before-after*) (L), (*before-after*)/*before* (R); greater than zero is improvement. The mean shows the absolute and relative difference of before and after, between the average of many spectra. Improvement from 80 to 400 Hz is noticeable; at higher frequencies there is degradation, negligible in relative terms, due to high-frequency filter rolloff. Frequencies below 50 Hz should be disregarded; they are usually not searched by LIGO, so spectra for this plot were generated with a high-pass filter at 38 Hz.

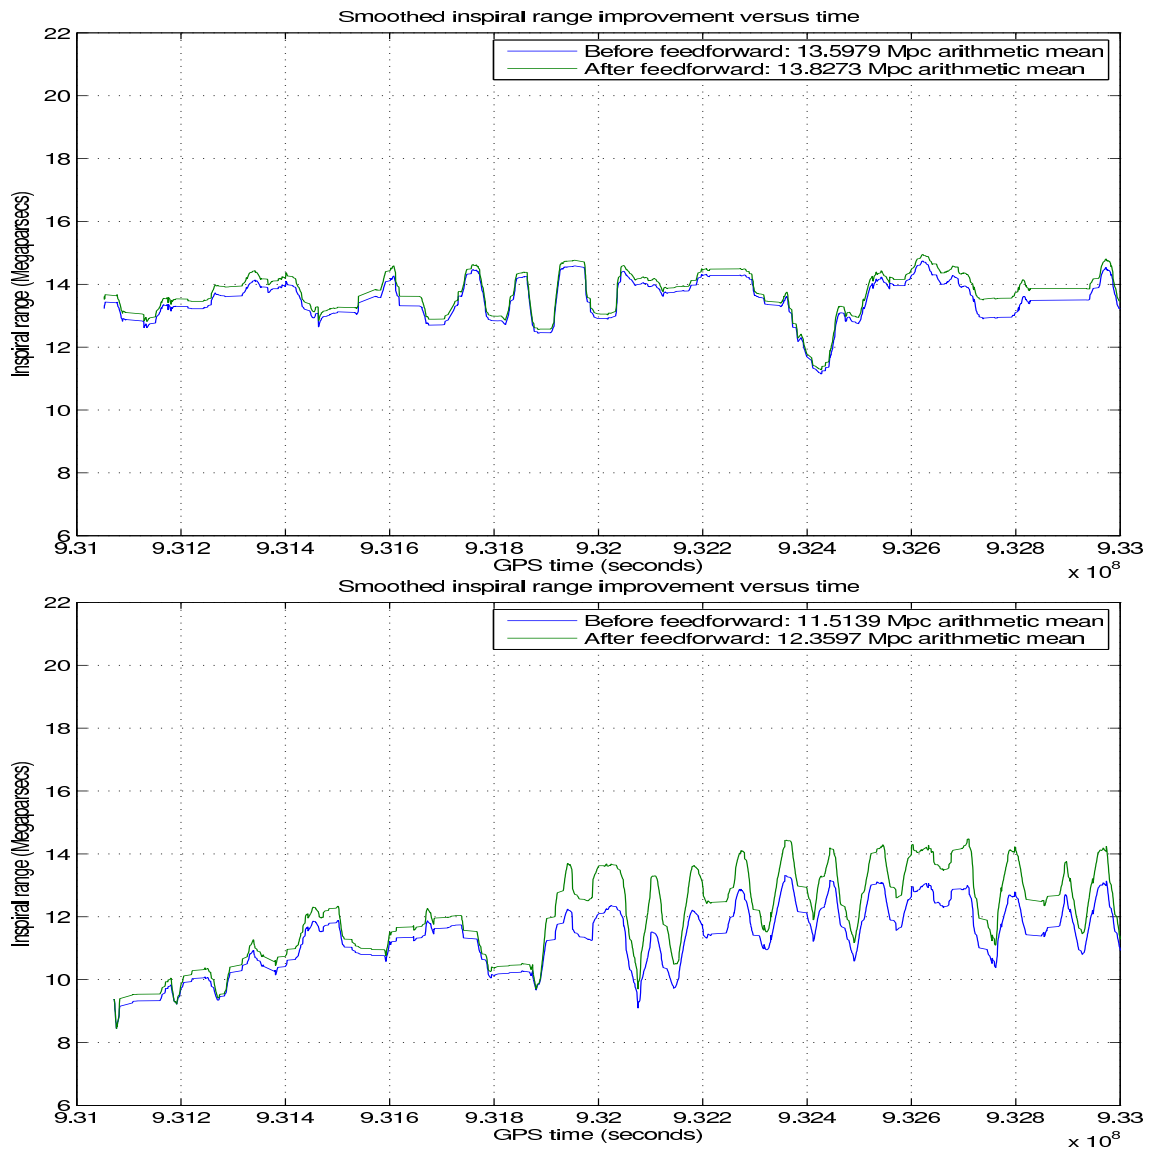


Figure 3.14: Inspiral range vs time for S6 (starting 2009 July 07) before GPS time 9.33e8 (2009 July 30): Each frequency band with reduced noise and consequent enhanced sensitivity contributes to the inspiral range integral, detailed further in Section 4.2.2. LIGO Hanford Observatory, H1 (top) gains 0.23 Mpc; LIGO Livingston Observatory, L1 (bottom) gains 0.84 Mpc. In this first month of S6, L1 saw greater benefit from post-facto feedforward correction; later data from H1 and L1 would improve by fluctuating amounts thanks to better real-time feedforward servos. Although H1 is less improved than L1 here, real-time tunings were made soon after.

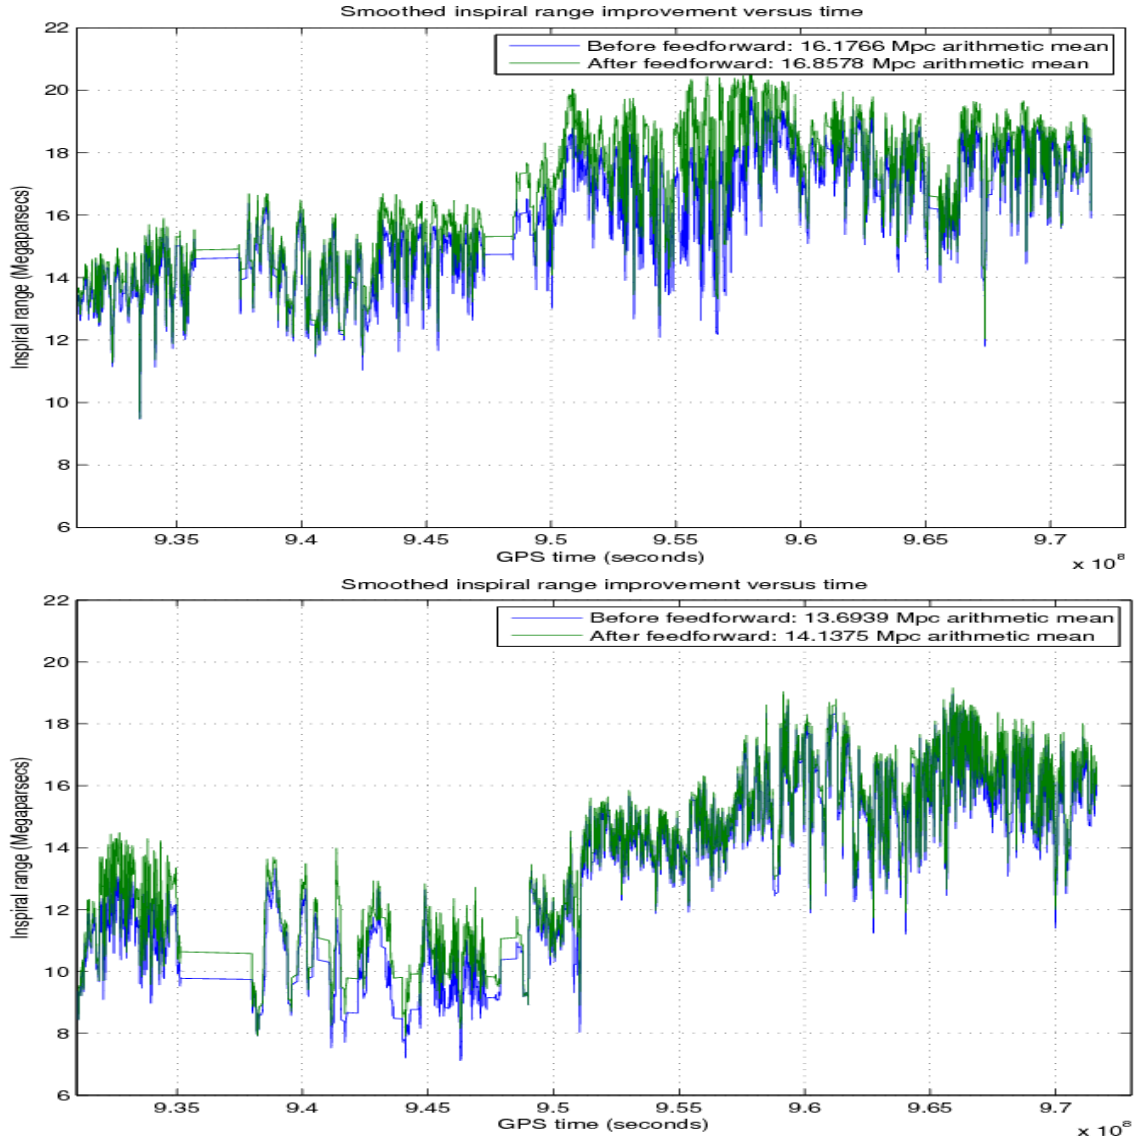


Figure 3.15: Inspiral range vs time for all S6 (2009 July 07 to 2010 October 20): LIGO Hanford Observatory, H1 (top) gains 0.68 Mpc; LIGO Livingston Observatory, L1 (bottom) gains 0.44 Mpc.




Figure 3.16: Inspiral range *fractional gain* vs time for S6 (starting 2009 July 07) before GPS time 9.33e8 (2009 July 30): LIGO Hanford Observatory, H1 (top) 1.68% better; LIGO Livingston Observatory, L1 (bottom) 7.00% better. This plot shows relative gain for the same data for we also show absolute gain.




Figure 3.17: Inspiral range *fractional gain* vs time for all S6 (2009 July 07 to 2010 October 20): LIGO Hanford Observatory, H1 (top) 4.14% better; LIGO Livingston Observatory, L1 (bottom) 3.60% better.

Diagnostics for GPS times 931000000 to 931100000

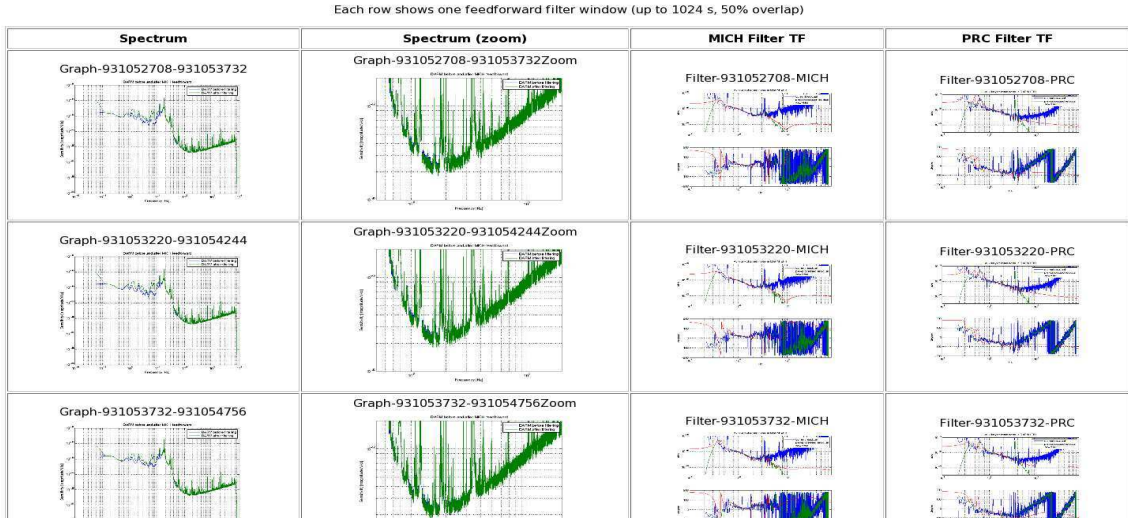


Figure 3.18: Screenshot of diagnostic web pages, indexed by window.

3.7 Conclusion

Auxiliary MICH-PRC Subtraction has cleaned LIGO S6 data, yielding better strain sensitivity and inspiral range. Frequency-domain-derived, time-domain-applied feedforward correction removes noise by fitting a rational transfer function between witness & target. Second order sections filter the witness channels, which then are subtracted from the measured target to produce an improved strain estimate. Diagnostics confirm that the corrected $h(t)$ benefits from dynamic, adaptive, algorithmic *post facto* feedforward subtraction, gaining several percent in detectable inspiral range. Such an improvement potentially enhances the performance of any LIGO search.

The subtraction leads to the lowest noise floor, around 150 Hz, of any time or interferometer so far (the highest performance to date at shot-noise limited frequencies has been obtained differently, with quantum optical squeezing [4, 74]). This record has only just been surpassed by Advanced LIGO. Thereafter, adaptive feedforward

filters, real-time or *post facto*, can be applied to mitigate noisy-but-inescapable couplings of the servo system. Signal recycling and filter cavities will further challenge commissioning. Angular and length sensing will need finer control servos. Advanced LIGO will also contain more physical and environmental monitors, from seismic and accelerometric to magnetic, that could provide witnesses for non-control-related noise. Altogether, more auxiliary channels and loops will exist, and while they may require sophisticated, non-linear methods, the subtraction technique presented here is a basis. Sensitive interferometry will benefit from simple, effective methods of suppressing instrumental influences.

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. *As noted at the beginning, most of this chapter is drawn from Meadors, Kawabe, and Riles, ‘Improving LIGO sensitivity with feedforward subtraction of auxiliary channel noise’ Classical and Quantum Gravity 31 (2014) 105014 [127], a paper that carries LIGO Document Number LIGO-P1300193.* This research was also made possible by the generous support of the National Science Foundation, awards 0855422 and 1205173, LIGO Hanford Observatory, the LIGO Scientific Collaboration, and the University of Michigan. The authors wish to thank Gregory Mendell as well as Stuart Anderson, Juan Barayoga and Dan Kozak for grid computing expertise, Ian Harry for investigating signal recovery before and after injections and providing conclusions about signal-to-noise ratio for matched filtering, and Jeff Kissel for refining MICH and PRC subtraction by hand. Tobin Fricke wrote the converter function from second-order-system to ZPK filtering and also reviewed this manuscript, as did Rana Adhikari and Jenne Driggers, who developed many of these methods at the Caltech 40 m interferometer.

Michael Coughlin, Jan Harms, and Nicolás Smith-Lefebvre all generously provided comments. Finally, we thank the referees for Classical and Quantum Gravity for helpful suggestions.

3.8 Postscript

This analysis is intended to make Advanced LIGO simpler. To that end, some of the tools involved were directed toward the first lock stretch of the advanced detector era, at LIGO Livingston’s L1 interferometer on 2014 May 26.

MEWD¹ attained the lock. MICH noise was dominating the spectrum. Based on the subtraction² of MICH & PRC noise from the S6 data set (and H1 in real-time), an investigation ensued. The Matlab function `f_domainsubtract` (from the 40 m, used in AMPS code to make a frequency domain transfer function) was run on data at CIT in

```
/archive/frames/A6/L0/L1/L-L1_R_1085/
```

using the `readFrames`³ function to pull 2100 s of frame data starting 1085202300.

Absent a `DARM_ERR` or `_CTRL` channel, the `_IN1_DQ` channels were used, following an example⁴ in the LLO aLog.

Uncalibrated PSDs from Matlab `pwelch` follow.

The function `f_domainsubtract` can estimate two quantities:

- coherence between different channels
- theoretically obtainable spectrum if noise is subtracted

The first set of graphs shows (in blue x’s) $\sqrt{1-\text{magnitude-squared coherence}}$

¹Denis Martynov, Anamaria Effler, Robert Ward, and Ryan deRosa.

²Meadors, Kawabe, Riles.

³by Gregory Mendell.

⁴due to Rana Adhikari; Kawabe subsequently confirmed that `IN1_DQ` has replaced `ERR` and `OUT_DQ` has replaced `CTRL`.




Figure 3.19: DARM power spectral density for first radio-frequency lock of L1 in Advanced LIGO. Note that the spectrum is uncalibrated.

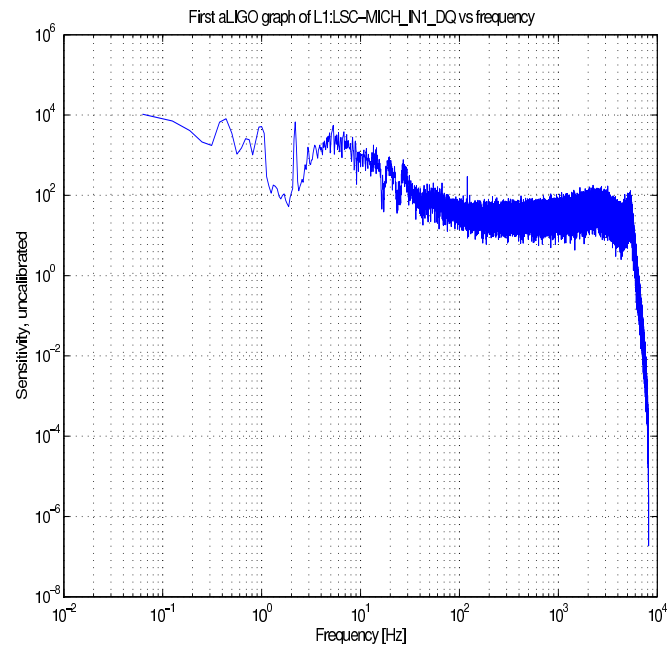


Figure 3.20: MICH power spectral density, first lock.




Figure 3.21: PRCL power spectral density, first lock. This corresponds to the PRC channel in Initial and Enhanced LIGO.




Figure 3.22: SRCL power spectral density, indicating the signal for the signal recycling cavity length.




Figure 3.23: MICH coherence with DARM.

and (in red) the ratio, (predicted amplitude spectral density (ASD) after subtraction)/(measured ASD before subtraction):

The second shows (in red) the ASD of the noise `_IN1_DQ` channel, (in green) the `DARM_IN1_DQ` channel, and (in black) the ASD of the `DARM_IN1_DQ` channel if the noise can be subtracted from it:

Summary

MICH and PRCL indeed cohere with DARM, in similar ways, each with a coherence (not MS-coherence) as great as 0.5 at 10 Hz, dropping off in either direction (about 0.05 at 3 Hz and 60 Hz).

In contrast, SRCL has a coherence as much as 0.3 at 10 Hz. It is even broader band, dropping off in either direction (about 0.05 at 3 Hz and 400 Hz, with a drop in coherence at 50 Hz).

These results corroborate the assessment that MICH noise is currently dominant and suggest that PRCL and SRCL noise are also significant. To be thorough this




Figure 3.24: PRCL coherence with DARM.




Figure 3.25: SRCL coherence with DARM.




Figure 3.26: MICH subtraction from DARM, projection, based on frequency domain transfer function. The full AMPS machinery has not yet been applied.




Figure 3.27: PRCL subtraction from DARM, projection.




Figure 3.28: SRCL subtraction from DARM, projection.

result should be verified using the OUT_DQ channels; these results were calculated using IN1_DQ.

Future

This is not the full feedforward code, which could be updated (not entirely trivially) to subtract out the noise and calculate a spectrum post-facto. This would require no work on the part of the commissioners. If a calibration exists, we could calculate a morale-boosting inspiral range.

If commissioners want a real-time filter, that would presuppose the existence of a MICH/PRCL/SRCL subtraction path in the model and closed loop response transfer functions. Once those exist, one could update (somewhat more trivial) the real-time calculation code to get a high-quality Foton filter.

A great deal of progress has already been made, and a great deal more can probably be made with manual filter design, but if sophistication is needed, the tools exist.

CHAPTER IV

Squeezing: Quantum Vacuum Phase Noise

Gravitational wave interferometry measures changes in light travel time far less than the light-crossing time of an atom; so the quantum nature of light has always influenced detector design. Advanced LIGO aims to reduce most of the noise sources described in Sections 2.4.2 and 2.4.3: even so, quantum shot and radiation pressure noise from light will remain. Carlton Caves first derived the quantum behavior of interferometer shot and radiation pressure noise [57]. Prior work had inferred the shot noise level from classical principles, accurately, but been ambiguous about radiation pressure. With the quantum noise clarified, it was realized that this noise could be reduced through so-called squeezing of the vacuum state [58]. This chapter describes part of how squeezing was successfully realized three decades later at the LIGO Hanford Observatory and describes the author's contributions to this milestone. For additional details, consult Chua [61] and Dwyer [75].

4.1 Squeezing theory

Quantum noise was initially treated as Poissonian photon counting (shot noise) and momentum coupling (radiation pressure), but it is better understood as the phase and amplitude uncertainty of the electromagnetic field at the interferometer output. Roughly speaking, phase uncertainty is equivalent to semi-classical shot

noise and amplitude uncertainty to radiation pressure. The standard quantum limit (SQL) on displacement sensitivity can be derived independently of either argument, as noted by Caves [58]:

$$(4.1) \quad (\Delta z)_{\text{SQL}} = (2\hbar\tau/m)^{1/2},$$

where in Equation 4.1 z is the mirror (test mass) displacement, τ is measurement duration, \hbar is the reduced Planck constant, and m is mirror mass. When one tries to attain this standard quantum limit with an interferometer, then shot and radiation pressure are found to balance at a particular optimum input power P_0 [58]:

$$(4.2) \quad P_0 = \frac{1}{2} \frac{mc^2}{\tau} \frac{1}{\omega\tau} \frac{1}{b^2} \equiv \frac{\hbar\omega\alpha_0^2}{\tau},$$

with ω the angular frequency of light and b the number of bounces (assuming delay lines, whereas a Fabry-Perot value is proportional to finesse) in Equation 4.2 and $\alpha_0^2 = P_0\tau/\hbar\omega$ is the number of photons. Equation 4.2 shows that the optimal power P_0 rises as $1/\tau^2$, that is, as the square of the gravitational wave frequency. High frequency sensitivity thus benefits from higher laser power. High mirror mass preserves low frequency performance when P_0 increases. Combined, these formulae imply that an interferometer should be built with the largest, most massive mirrors and most powerful lasers practicable.

4.1.1 Problems with lasers: thermal compensation

LIGO and its allied interferometers do try to operate with massive optics and powerful lasers, but there are limits. Vacuum enclosures place a constraint, only broken with significant expenditures, on the size of optical tables and mirror suspensions and therefore on mirror mass. Large optics, both substrates and coatings, are

also difficult and expensive to manufacture with high quality. This indirectly limits laser power through the SQL. Laser power, however, was not limited for this reason in initial and enhanced LIGO. In practice, thermal distortion of the test masses was a worse problem [42]. Absorption, both in the bulk of the fused silica substrate and in the HR coatings, of a few parts per million was sufficient to distort the radius of curvature of the Fabry-Perot cavities. While managed with a CO₂ laser-driven thermal compensation system (supplemented with ring heaters in aLIGO), this thermal distortion remains a serious concern.

Having an alternative to increased mirror size and laser power would grant flexibility. Hence squeezing: the standard quantum limit can be achieved through another means.

4.1.2 Quantum shot noise and radiation pressure

Squeezing as used in the context of gravitational wave interferometry has a specific meaning. While no alterations are made to the input laser light, the output ‘dark’ port of the interferometer, whence comes the light received by the photodetector, is changed. The output port in a squeezed interferometer receives a squeezed vacuum state from the opposite direction to the laser light. This squeezed vacuum state is prepared on a ‘squeezer’ table. This beam is injected with the correct polarization, using a Faraday isolator and wave plates to guard against laser light entering the squeezer, and propagates through the interferometer arms (requiring proper alignment) until it and the laser light are both incident on the same output photodetector. The desired effect is altering $\Delta E \Delta \phi$ uncertainty for the vacuum state of the electromagnetic field. Reducing $\Delta \phi$ at the cost of still-acceptable ΔE permits¹ a more precise measurement, equivalent to the beneficial effect of increasing laser power but

¹ ΔE can also be squeezed; see Section 4.1.3

without heating the mirrors. Loosely, more laser power boosts the shot-noise limited ‘signal’ of a high-frequency measurement whereas squeezing reduces the ‘noise’. Formally, the expectation value of quantum operators is improved.

Caves’s prescription [58] is as follows:

- Vacuum fluctuations couple through the anti-symmetric port, so inject a
- *Squeezed* \vec{E} field (uncertainty ellipse with smaller ΔE) defined by
- Squeeze operator S in Equation 4.3 (with squeeze angle θ , factor r , creation operator a):

$$(4.3) \quad S(\zeta) = \exp\left[\frac{1}{2}\zeta^* a^2 - \frac{1}{2}\zeta(a^\dagger)^2\right], \quad \zeta = r e^{i\theta}$$

thereby reducing shot noise by a factor of e^{-r} . The optimum power changes accordingly: Caves shows that $P_{\text{opt}} = P_0 e^{-2r}$, so long as $|\alpha| \gg \sinh^2 r$ and $|\alpha| \gg e^{2r} \sinh^2 r$.

For a 20-W interferometer using 1064 nm light to measure at 100 Hz, the latter squeezing condition would be an issue around $r \geq 10$, but to reduce optimal power by 1/2 only requires $r \approx 0.34$. Trying to reduce optimal power by orders of magnitude is unlikely to occur any time soon: rather, a modest laser as would ordinarily be used in an interferometer can be supplemented with a squeezed beam to increase its effective power. Squeezing is an addition, not a replacement. Yet the squeeze beam still must be physically generated, and this along with optical losses in the interferometer prove a much more significant challenge than theoretical constraints.

Squeezing can be physically generated through the use of an optical parametric oscillator (OPO). Optical parametric amplifiers, including an OPO, involve a nonlinear optical medium that is pumped with an electromagnetic wave at ω_p . An

OPO [58] is an optical parametric amplifier for which the output ‘signal’ (at ω_s) and ‘idler’ (at ω_i) beams² are degenerate (for any amplifier they satisfy $\omega_s + \omega_i = \omega_p$, the input, pump frequency). As Takahashi noted [159] (and others, see Wall for a review [170]), a degenerate amplifier will transform a coherent state into a state with unequal uncertainties in its two quadratures, to wit, phase and amplitude.

By placing a degenerate amplifier in the cavity resonant at ω_s , the amplifier becomes an optical parameter *oscillator*, with high effective interaction length and thus squeeze factor. Nonlinear media come in many varieties, but the material used in LIGO squeezing experiments (as in this chapter) has typically been PPKTP (periodically-poled potassium titanyl phosphate). For an effective second-order nonlinear susceptibility d , pump electric field amplitude E_p , effective interaction length L , and index of refraction at the signal frequency n_s , the amount of squeezing is [58],

$$(4.4) \quad r = \left(\frac{4\pi\omega_s L}{cn_s} \right) d|E_p|.$$

Another, heuristic³, understanding of the OPO (when $\omega_p = 2\omega$) is that a vacuum fluctuation ϵ at $Ae^{i\omega\epsilon t + \phi_\epsilon}$ couples to the pump and so induces a second beam, $Ae^{i(2\omega - (\omega_\epsilon t + \phi_\epsilon))}$. The superposition of these two beams has amplitude $\approx 2A$ but phase $\approx (\phi_\epsilon - \phi_\epsilon)$, *i.e.*, the amplitude fluctuations (radiation pressure) have increased but the phase fluctuations (shot noise) have been reduced.

One caveat: the pump must be phase matched to the main laser (at ω) of the overall interferometer. This is usually accomplished (as at the Hanford experiment) with a small beamsplitter diverting a portion of the laser (then, at Hanford, through a polarization-preserving fiber) and directing it to a second harmonic generator

²The nomenclature is historical, sometimes distinguished by $\omega_s \geq \omega_i$. After alignment is complete, the output beams contain no power by design.

³Popularized by Daniel Sigg, formally described as sidebands in Dwyer [75].

(SHG), which yielded $\omega_p = 2\omega$. SHGs themselves operate using non-linear crystals to frequency-double incident laser light. Imperfect phase matching is one of two chief reasons why squeezing can underperform. Phase matching the squeezer requires stable optical fiber over distances of at least 10 meters. The other reason for degraded squeezing is optical losses.

Optical losses can include absorption, scattering, and imperfect mode-matching. While beyond the scope of this chapter, these optical losses are the principal reason that interferometer squeezing at the H1 photodetector in our experiment was about half what was generated in the OPO and must be addressed as squeezed light sources become more widespread.

Demonstrated first at GEO600 [162], then on H1 at the LIGO Hanford Observatory [4], squeezing is now a proven technique. At Hanford, the subject of this chapter, 2.15 dB of squeezing (28%, equivalent to a 64% increase in laser power [4]) was achieved with the combined efforts of Sheila Dwyer, Sheon Chua, Lisa Barsotti, Matt Evans, Keita Kawabe, Daniel Sigg, Conor Mow-Lowry, Alexander Khalaidovski, Maxim Factourovich, Nicolás Smith-Lefebvre, Robert Schofield, Mike Landry, Cheryl Vorvick, and Richard Gustafson, along with the author. Figure 4.1 shows a diagram of the squeezer in relation to the interferometer. A picture of the squeezer and some of its scientists is shown in Figure 4.2. The squeezing experiment was well supported by the entire staff⁴ of Hanford Observatory, to whom the squeezers were grateful. This success encourages development of a production-quality squeezed light source for a future upgrade to Advanced LIGO, along the lines of Section 4.1.3. Squeezed light should increase aLIGO's performance and could prove critical to its successors, especially should they operate at cryogenic temperatures for which excess heat would

⁴Quite busy themselves with Advanced LIGO installation.

be problematic.

4.1.3 Squeezing filter cavities

Future squeezed light sources may prove more sophisticated still. Equation 4.3 includes an angle ω . When ζ is positive, squeezing is generated and phase fluctuations diminish, as discussed above. The converse, ζ negative, yields anti-squeezing, where amplitude fluctuations are smaller. This angle ω is controlled by adjusting the phase of the main laser pickoff as it enters the OPO. It would be beneficial if amplitude fluctuations could be reduced at the frequencies where they are the dominant noise sources and likewise simultaneously with phase fluctuations at other frequencies. Such simultaneous reduction can, in principle, be accomplished with a filter cavity, which can rotate the squeezing uncertainty ellipse in a frequency-dependent way, as Chua summarizes [61]. Prototypes at gravitational wave interferometer frequencies are under construction. If promising, they might be tested at the observatories in a few years time.

4.2 LIGO Hanford Observatory quantum vacuum squeezing

4.2.1 Collaboration and contributions

Squeezing tests at LIGO Hanford Observatory had been planned for several years in advance of the 2010-2011 experiment. Nergis Malvalvala at the Massachusetts Institute of Technology, David McClelland at Australian National University, and Daniel Sigg at Hanford made arrangements for a prototype squeezer to be developed at MIT, with an OPO developed by ANU. Experience with squeezing at GEO600 (Roman Schnabel and Henning Vahlbruch) informed planning. When the prototype squeezer had been made functional at MIT, it was disassembled and shipped cross-country to LHO. There it was reassembled and integrated into the 4-kilometer

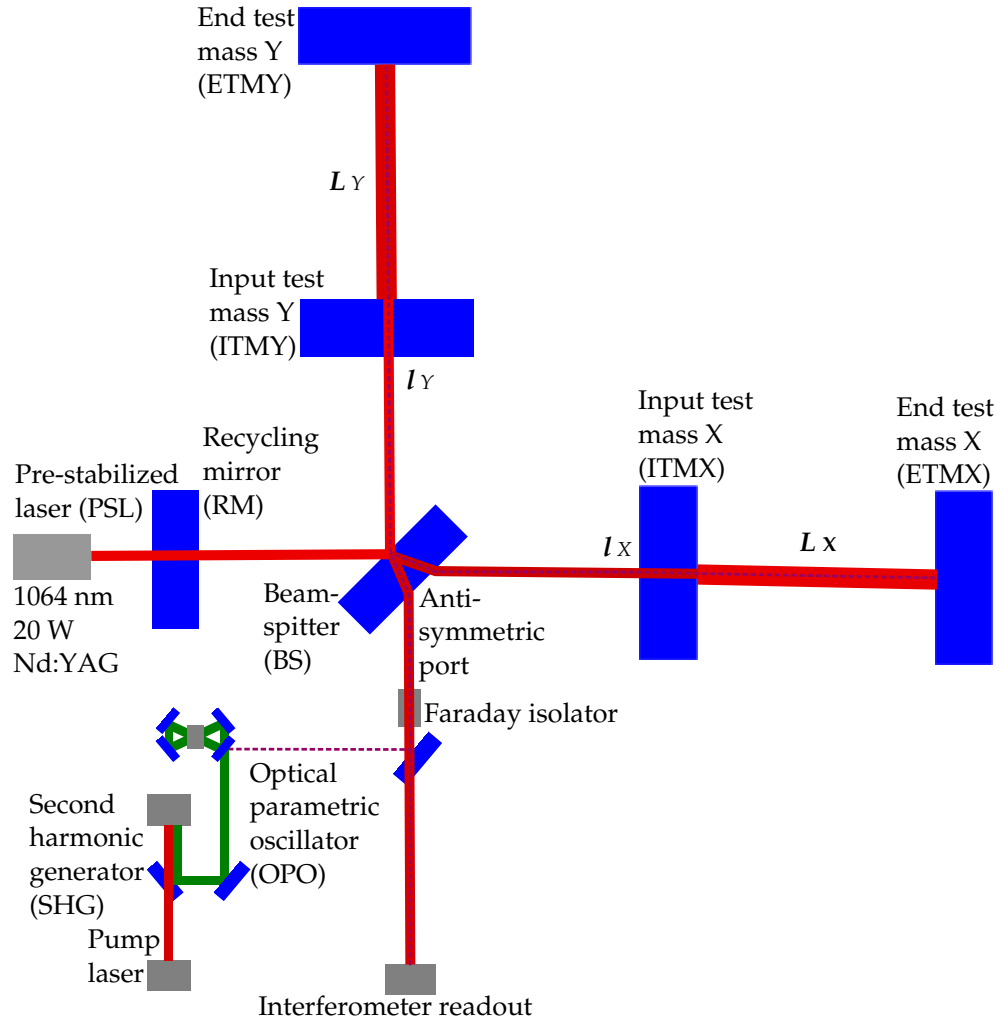


Figure 4.1: Diagram of squeezer integration into the LIGO interferometer. Red lines show 1064 nm laser light, green 532 nm; dashed are squeezed (1064 nm) beams. A pump laser, phase-locked to the PSL, generated laser light at 1064 nm, boosted to 532 nm by the second harmonic generator (SHG). This 532 nm beam resonates in the optical parametric oscillator, containing a periodically-poled KTP crystal for non-linear optical effects. The squeezed beam is generated there and exits the OPO, passes through the Faraday isolator and into the interferometer. Simplified and adapted figure [4].




Figure 4.2: Image by Lisa Barsotti in Hanford eLog.
Counterclockwise from lower left: Sheila Dwyer, Lisa Barsotti, Conor Mow-Lowry, Grant Meadors. This photograph shows the uncovered squeezer table with components from the MIT squeezer experiment unpacked at LIGO Hanford Observatory in November 2010. The table sat in a temporary location by HAM6, where it was recommissioned by Dwyer, Mow-Lowry, Sheon Chua, and Alexander Khalaidovksi until H1 could be brought back online for the squeezing experiment in late 2011.

interferometer by a team led by Barsotti, Sigg, and Kawabe, with Dwyer and Chua in charge of optics and operation. This author's contributions were in support: the design and installation of supports for the optical table, assistance with in-vacuum installation and measurements of both optics and electronics, as well as interpretation of the squeezing results in astrophysical significance.

Optical table support assembly

No space for an entire squeezing table or the constituent optics as then laid out existed inside the H1 vacuum enclosure. Thus a new location had to be found, situated close to the beam splitter's dark port. This site was by the vacuum chamber known as HAM4. HAM4 had a free viewport, initially covered by steel then provided a vacuum window by Vorvick and Gerardo Moreno, wherein the squeezed light beam could be injected toward the dark port. While the optics were installed and aligned on a spare, standard optical table, additional support for the table needed to be manufactured. Existing prototype legs for the table were incorporated because of their good theoretical isolation performance, but the table height would have been too short for the optics to reach the viewport without a periscope, and such a periscope would likely have caused unacceptable beam movement. Some interest existed in using the prototype legs together with added spacers rather than entirely new legs both because of cost and the potential for using the prototype legs for other aLIGO applications.

After several design iterations, the author designed and ordered four table leg spacers. These closely followed the design of Figure 4.4 with several subsequent adjustments (see caption). On arrival in March 2011 (Figure 4.3), the full leg assemblage was installed (with a clear, also spare, optical table on top) at the final site with the help of Landry, Dani Atkinson, the Apollo Sheetmetal crew, and the



Figure 4.3: Table legs testing. The squeezer table is raised to its final height by the leg extensions. From top to bottom: table (WISCT10), existing leg extensions (with flanges), new leg extensions (flangeless), triangular high table legs. This assemblage provided the squeezer a serendipitously-stable (as measured by Sheila Dwyer, Sheon Chua, and Robert Schofield) platform at low cost. Photo in temporary location; the actual squeezer table was anchored to these table legs, grouted, by HAM4.

author. The vibration isolation on the table was measured by Schofield and Chua, using mechanical shakers to obtain the transfer function ⁵. The legs having decent performance, Schofield recommended grouting the legs to the floor to achieve better stability. Once this was done in June, the clean optical table was deinstalled and the now-reassembled squeezer table emplaced on top⁶. Cables, including the optical fiber transporting light from the main laser for the phase-locked pump, were laid again. Finally, surrounded by a physically separate platform for accessibility onto the table, the squeezer was in position to inject a beam into H1.

⁵Later, final numbers were measured in October 2011: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=10/03/2011&anchor_to_scroll_to=2011:10:03:23:49:25-lisabar

⁶See http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=06/24/2011&anchor_to_scroll_to=2011:06:24:11:49:32-gmeadors (public URL) for details




Figure 4.4: Table legs SolidWorks schematic. This initial design for the table leg extensions on the squeezer table incorporated a flange, which was removed immediately prior to fabrication, replaced with a larger diameter flangeless tube with incorporated tap-holes. Flanges were thought necessary for a flexible alignment initially, they would be prohibitively expensive to machine, and welding would induce unacceptable distortions into the metal. Manufactured by Brockman Machine Works, of Kennewick, Washington. Design in consultation with Daniel Sigg, Lisa Barsotti, Keita Kawabe, Gerardo Moreno, Richard Savage.

Faraday isolator measurement

Squeezing is helpful only if it helps detector sensitivity more than harms it. Backscattered light from the interferometer, reflected by the squeezer back into the main detector [61, 4], is a serious concern, particularly phase noise in the scattered light fields. In order to mitigate this, an additional Faraday isolator is put into the optical path of the squeeze beam, which acts to supplement the intrinsic 40 dB isolation thanks to the OPO. Measuring the isolation ratio of this Faraday isolator was a task for Kawabe, Barsotti, Evans, and myself.

Our results ⁷ were that the Faraday isolator, as measured in air, transmitted 480 of 500 mW injected in the correct direction (96% transmission), attenuated 460 mW to 35 μ W when injected in the wrong direction (-41.3 dB isolation ratio), and would back-scatter 2.5 mW of 500 mW (-23 dB) input power toward the squeezer. These numbers sufficed to proceed with in-vacuum installation.

In-vacuum installation

Installing the additional Faraday isolation in vacuum was necessary before the squeezed beam from the table could be injected directly into the dark port. The new Faraday, but not the existing initial LIGO Faraday, had an opening to accommodate the beam. Because Advanced LIGO installation, concurrent with the squeezing experiment, involved opening the vacuum, our own installation team (Kawabe, Barsotti, Dwyer, Chua, Gustafson, and the author) were able to emplace and align the new Faraday isolation on the optical table inside HAM4, move a baffle in HAM5 (optically downstream), and direct the beam onto HAM6 (the optical terminus), the location of the output photodiode⁸.

⁷See http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=06/09/2011&anchor_to_scroll_to=2011:06:09:16:58:36-lisabar for more details.

⁸Detailed in http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=06/29/2011&anchor_to_scroll_to=2011:06:29:18:31:51-lisabar on the electronic log.

Once the squeezing beam was on the output photodiode, the vacuum enclosure pumped down and H1 restored to operation⁹, it was determined that the interferometer was not functioning as expected. The output mode cleaner transmission was about 25% lower than previously measured [169,156]. This loss was severe enough to warrant an additional vacuum incursion. Our task was to install the now-unused L1 OMC as a tentative replacement. First, Evans, Barsotti, Chua, Kawabe and myself verified¹⁰ that the L1 OMC had superior finesse to that inferred for the H1 OMC. A rapid replacement of the H1 OMC by the L1 followed, over the course of two days¹¹, involving Bram Slagmolen, Smith-Lefebvre, Kawabe, Evans, Dwyer, Chua, Factourovich and myself. Once installation was complete and the vacuum pumped down, H1 returned to operation.

In the optics lab, the H1 OMC was found [169] to have one of its control elements, a heater, out of place. It was occulting the beam inside its resonant bowtie cavity, reducing the finesse. The transplanted L1 OMC worked as intended inside H1.

Data digitization

Full control of the squeezer required more of its digital computer input and readout to be accessible proximal to the rest of the interferometer controls. Before H1 was resuscitated, the author had made and helped lay cables and, with Dave Barker, selected computer channels for the analog-to-digital conversion of squeezer data¹².

Commissioning began.

⁹And local wildlife rescued: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=09/23/2011&anchor_to_scroll_to=2011:09:23:21:37:25-sheilad

¹⁰See http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=10/21/2011&anchor_to_scroll_to=2011:10:21:11:12:21-sheon

¹¹Day 1: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=10/24/2011&anchor_to_scroll_to=2011:10:25:09:11:32-kawabe

Day 2: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=10/25/2011&anchor_to_scroll_to=2011:10:25:19:49:56-kawabe

¹²Fuller information in the electronic log: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=08/16/2011&anchor_to_scroll_to=2011:08:16:21:27:49-gmeadors and http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=08/26/2011&anchor_to_scroll_to=2011:08:26:15:08:16-gmeadors




Figure 4.5: Integrand of inspiral range as a function of frequency, with and without squeezing.

4.2.2 Squeezing's scientific benefit

Squeezing worked in spite of optical losses, did not make the interferometer worse, and in fact increased shot noise limited performance by over 2 dB [4]. The benefits translate directly into measurable scientific benefits.

Figures of merit: inspiral range

As noted in Chapter III, gravitational wave interferometers are regularly characterized by their inspiral range [80]. This figure of merit integrates the gravitational wave strain spectrum into a single number: the distance to which a pair of 1.4 solar mass neutron stars (canonically, with average orientation and sky location and a signal-to-noise ratio of 8) could be detected if they inspiralled together:

$$(4.5) \quad F_{7/3} = \int_{f_l}^{f_h} [f^{7/3} h^2(f)] df,$$




Figure 4.6: Net effect of squeezing on inspiral range integrand. Scientific benefit from squeezing is evident at the few hundred Hz ‘bucket’ where initial LIGO is most sensitive, and although low frequency noise is slightly worse, this in an already noisy spectral band. Enhanced LIGO unambiguously benefitted from squeezing, proving the technique’s efficacy.

$$(4.6) \quad R = \Theta \left(\frac{5c^{1/3} M_{\text{ch}}^{5/3} F_{7/3}}{96\pi^{4/3} \text{SNR}_0} \right)^{1/2} L,$$

where R is inspiral range in kiloparsec, f is frequency, h the strain sensitivity as a function of frequency, Θ is a correction factor of 1.77 for orientation and sky location, c is the speed of light, $M_{\text{ch}} = Gc^{-2}M_{\text{sun}}(M_1M_2)^{3/5}(M_1 + M_2)^{-1/5}$ is the chirp mass of two stars with masses $M_1 = M_2 = 1.4M_{\text{sun}}$, $F_{7/3}$ is as defined in Equation 4.5, SNR_0 is typically 8, and L is the interferometer length in meters. The range integrand in Equation 4.5 can itself be multiplied by the external coefficients in Equation 4.6 to yield a plot of the contribution of each frequency band to the scientific performance of the interferometer, as in Figure 4.5. Taking the difference of before- and after-squeezing shows where squeezing yielded improvement – and demonstrates that it did not seriously degrade interferometer performance¹³.

¹³For additional plots in the log, <http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?>

From the improved high frequency shot noise, we can see that squeezing bought Enhanced LIGO a megaparsec of inspiral range. This number is impressive in several respects: our goal was to achieve a squeezing factor of perhaps as much as 3 dB, but to do it in the shot noise-limited region, at high frequencies, where the inspiral range equations count for much less. The range integrand shows squeezing down to 150 Hz, which is the lowest yet achieved for a gravitational wave interferometer, as can be seen in the Nature Photonics paper [4]. Squeezing at such low frequencies brings most of the inspiral range improvement. In all, the squeezer experiment added about another five percent to H1's range, but also something more significant: prospects for enhancing gravitational wave interferometers beyond the standard quantum limit.

4.3 Squeezing large interferometers: success and prospects

Over 2 dB of squeezing were achieved by the LIGO Hanford squeezing experiment on H1 [4]. Building on the previous success of GEO600 [162], squeezing is now a mature technique that we hope to incorporate into LIGO permanently as soon as it is feasible. Publications by Dwyer [74] and Chua [62] detail the remaining hurdles. Dwyer notes that quadrature phase noise will become increasingly significant as losses are reduced, particularly for filter cavities. Chua proposes a path toward reduced backscattered light noise. As we develop an understanding of these two principal problems, work at MIT is progressing (led by Tomoki Isogai) on realizing a filter cavity suitable for the LIGO frequency band. With it, future interferometers will be able to introduce frequency-dependent squeezing to finally reduce, not just shot noise, but radiation pressure as well, below the standard quantum limit.

group=detector&date_to_view=11/30/2011&anchor_to_scroll_to=2011:11:30:12:54:18-gmeadors
 feedforward MICH correction was attempted but unneeded, as detailed also, http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=12/02/2011&anchor_to_scroll_to=2011:12:02:15:06:38-gmeadors.

CHAPTER V

TwoSpect: Search for a Simulated Scorpius X-1

5.1 Neutron stars in binary systems

Continuous waves from neutron stars in binary systems are distinctive. As Section 2.3.2 comments, binary systems constitute 211 of 379, or 55%, of known pulsars with rotation frequency above 10 Hz and thus in the LIGO band (see the ATNF catalog, currently containing 2328 pulsars total [119]). CW search methods prove powerful, probing well below the LIGO amplitude spectral density (ASD) noise floor given year-length science runs, yet computationally-demanding. A non-trivial possibility exists that there are CWs from neutron stars (NS) in binary systems, buried in existing or forthcoming data, that could be seen with a fast, tractable search method.

TwoSpect [91, 90, 7] is such a method. Developed as a search for neutron stars in binary systems with unknown sky location, orbital period, projected semi-major axis or frequency, it has been tested and run on S6 data. In such a mode it accepts some degradation of sensitivity in order to scan the entire parameter space. If some parameters are known, a deeper search becomes possible. This chapter describes the development of this ‘directed’ search and its application to a simulated data set (a ‘Mock Data Challenge’ or MDC). The next chapter summarizes the preliminary results when applied to data from the last science run, S6.

5.1.1 Continuous gravitational waves from neutron stars

Choosing to pursue binary sources is itself a choice of search direction. Neutron stars can be in isolated systems as well, for which many searches have been conducted (summarized in Section 2.3.2). Yet isolated stars have several issues that make them challenging to detect. Isolated stars have a finite lifetime (albeit many millions of years) over which they are hypothesized to emit detectable GWs. During this detectable lifetime, they are losing energy to GWs – this energy loss results in a continually slowing frequency, denoted by the spindown rate $\dot{f} = df/dt$. GWs are not the only contributor to spindown: non-gravitational energy loss due to friction or magnetic interactions with an accretion disk will also sap the rotational energy of the star and thus its GW emission amplitude. Non-GW losses broaden the uncertainty in spindown rate. Spindown in turn makes GW searches more computationally challenging: templated-phase mismatch over a science run of many months is often significant enough to require additional templates to search over putative spindown values. Since many searches are already computationally limited, a GW source without these challenges would arouse interest.

Accreting binary systems generate such interest. LMXBs, low-mass X-ray binaries, of which Scorpius X-1 (Sco X-1) is an exemplar, are at the focus. As understood from the work of Papaloizou & Pringle [139] and Wagoner [168], LMXBs might exhibit several key properties:

- Longer GW lifetime than isolated sources, due to accretion recycling
- Ellipticity & hot spots, due to accretion
- Hypothetical *torque balance* frequencies

Not only are LMXBs an abundant fraction of the neutron star population, they

also constitute a disproportionate quantity of fast period, *millisecond* pulsars – and interestingly may be capped at a speed limit (hypothesized by Chakrabarty [59]) below the relativistic breakup limit. If so, LMXBs in particular would constitute an ideal search target. Let us review the points of lifetime, ellipticity, and torque balance in further detail.

5.1.2 Binary spin-up and detectable lifetime

As stated, isolated, elliptical, rotating NS will lose energy to GWs. This intrinsically makes their GW radiation short-lived. Binary systems, in contrast, can be recycled by accretion. Not necessarily a steady process, recycling of infalling matter from an orbital partner can raise the spin frequency of a neutron star by conservation of angular momentum. Energy is released as heat and radiation on impact with the surface of the NS, in particular as X-rays in the case of an LMXB. The work of Papaloizou, Pringle and Wagoner elucidates the implications for GW detection. As infalling matter hits an NS, the matter stream can not only increase the spin frequency – and thus GW radiation amplitude – of the NS, but it can also increase NS ellipticity from matter accumulated at the hot spot. Equation 2.25 shows that this further increases radiated GW amplitude. Insofar as the infall stream continues, the NS can be ‘spun-up’ to rotation frequencies with non-negligible GW emission, potentially for a far longer detectable lifetime than an isolated star.

Spin-up might proceed, as Wagoner derives, only up to a specific frequency. In this *torque balance* hypothesis, the matter infall delivers a torque offset by angular momentum loss from GW emission; since GW emission rises with frequency, an NS that is spinning up due to accretion will eventually reach a limiting frequency beyond which GW emission prevents it from spinning faster. This idea, developed by Bildsten [47], connects accretion rate with GW emission. Since accretion rate can

be measured electromagnetically, in X-ray flux $F_{\text{X-ray}}$, this lets us estimate what level of GW emission at frequency f could be expected for a star rotating at frequency $\nu = f/2$ (Equation 5.1 in general, Equation 5.2 for Scorpius X-1 [91, 47]):

$$(5.1) \quad h_0 \approx 5 \times 10^{-27} \left(\frac{300\text{Hz}}{\nu} \right)^{1/2} \left(\frac{F_{\text{X-ray}}}{10^{-8}\text{erg cm}^{-2}\text{s}^{-1}} \right)^{1/2},$$

$$(5.2) \quad h_0 \approx 2.8 \times 10^{-26} \left(\frac{600\text{Hz}}{f} \right)^{1/2}.$$

Torque-balanced neutron stars would remain at their torque-balance frequency: accretion would counteract spindown. This frequency stationarity is computationally useful.

Although the torque-balance equation predicts GW strains for Scorpius X-1 and similar sources that are below even the Advanced LIGO noise floor, there is cause for optimism. Scorpius X-1, although the brightest steady extrasolar X-ray sources, is not (see Section 5.1.3) necessarily the only object of its kind. Particular if other sources emit their X-ray flux in beams, those other sources might be underestimated or unknown to us – and they might be closer, a motivation that drives an all-sky search ???. For Scorpius X-1, there is much incentive to connect extensive electromagnetic observations with GWs. As astrophysical ephemeris data from the source becomes available, search techniques can become more targeted, which when combined with analytical techniques designed to take advantage of that information (and a science run of several months) could push substantially below the Advanced LIGO noise floor. Under these conditions, the torque-balance limit may be attainable. Reaching that limit – whether or not GWs are detected – would tell us about the astrophysics of LMXBs.

5.1.3 Detection rate projections

Equation 5.1 is incentive to develop searches especially for LMXBs. Although the predicted h_0 is still low, it is conceivable that high-luminosity sources such as Scorpius X-1 (see Section 5.8.1 for full parameters) or possible transients might just be within reach of advanced detectors such as aLIGO. Unknown sources with potentially higher GW flux are a motivation for making this search capable of scanning the entire sky. It is difficult to infer how many sources might exist, but over 50 accretion-powered X-ray pulsars were already known as of 2005 [119]. Known, especially-promising sources such as Scorpius X-1 warrant scrutiny with a deeper, *directed search* over their known parameter space. Such searches are the purpose of TwoSpect.

5.2 TwoSpect searches

TwoSpect offers a way to calculate a statistic and infer detection probability for a putative template waveform of a neutron star emitting continuous GWs in a binary system. LIGO CW searches generally use matched filters, like inspiral searches. Although the CW filter is easier to calculate than inspiral filters for merging black holes, the net computational cost is much higher, since integration covers a much longer timespan, compensating for weaker signals. Matched filtering a binary pulsar search in the time-domain, with complete phase information, appears to be computationally intractable for terascale computer clusters in the early 2010s. By switching the problem to frequency domain power instead of amplitude, the search becomes feasible, albeit at a cost in sensitivity.

TwoSpect performs two transforms: it first parcels a science run (year-scale) into overlapping short Fourier Transforms (SFTs, hour-scale or less). SFT frequency bins for a given data stretch are then adjusted to the frequency at the solar system

barycenter, accounting for Doppler shifting that would be induced at the sky location, frequency, and time under investigation. Each row of constant-barycentered frequency bins is then treated as a time series. That time series is then Fourier transformed too – the two in TwoSpect. This transform yields a plane of pixels: the frequency-prime f' vs frequency f plane. Goetz and Riles [90,91] developed this technique in the context of GWs. The plane of pixels provides a tractable data set on which binary CW templates can be tested.

Detectable strain sensitivity h scales poorly with time, $T^{-1/8}$, due to the double transform, instead of $T^{-1/2}$ for a coherent search or $T^{-1/4}$ for semi-coherent, but matched templating techniques are powerful enough that being able to use them is promising. Goetz developed a noise-subtracted R -statistic to measure the cumulative power in each binary CW template. This statistic follows a distribution as the sum of many χ -squared distributions (see Taylor [160] as well as Casella and Berger [56] for statistical reference) and yields consistent p -values. These p -values can be extrapolated from Davies' method, having been vetted by Monte Carlo [90]. By providing these statistics, TwoSpect yields information on a large fraction of parameter space hereto unsearchable. The code is open source and available online [161].

5.2.1 Two spectra: a double Fourier transform

The double-Fourier transform is illustrated in Figures 5.1 and 5.2. Existing LIGO routines perform the steps needed to generate the first figure: SFTs are made in pre-processing and then barycentered. TwoSpect then re-invokes the Fast Fourier Transform routines to transform the frequency vs time plane into a frequency vs frequency-prime plane. Although GW phase information is lost (and, for now, templates ignore orbital phase), the resulting search domain is robust against spin wandering.




Figure 5.1: After Doppler-shifting the frequencies into the solar system barycenter, TwoSpect analyses begin on this first, time-frequency plane. A simulated signal at 100.015 Hz and $asini = 1.44$ is injected with $h_0 = 4 \times 10^{-21}$ into 10^6 seconds of Gaussian noise at $Sh = 4 \times 10^{-24}$ (the projected minimum Advanced LIGO noise level); the signal period is 68023.8259 seconds, as with Scorpius X-1.




Figure 5.2: Fourier-transforming along the ‘rows’ (constant frequency bin, variable time bin) generate a second plane, the frequency-frequency plane. The power of each bin in this transform is plotted as a pixel. By aggregating power, this second Fourier transform enhances signal so that matched templates can be applied for a search. The x -axis label is in fact *bin number* and each value is the 2nd Fourier frequency in Hz times 1800 s.

5.2.2 Inferring neutron stars with companions

With the search domain prepared, templates can be tested. The modulation induced by LMXB partners on their NS companions is typically fractions of a Hertz. Modulation depth is, more precisely [90],

$$(5.3) \quad \Delta f = \frac{2\pi f(a \sin i)}{cP},$$

where f is GW emission frequency, $a \sin i$ is projected semi-major axis, and P is orbital period. Equation 5.3 specifies the amplitude of the sinusoid seen in Figure 5.1, though it must be noted again that it is the power of the transform of that sinusoid in Figure 5.2 that is actually template-tested.

5.2.3 TwoSpect algorithm detection statistic

Template-testing proceeds from a given test frequency f , modulation depth Δf (via astrophysical $a \sin i$), and period P using data prepared for some sky location. Having a time series n SFTs long with k frequency bins, a template weight can be computed for a number of pixels $M < n * k$. Applying these weights yields the R -statistic:

$$(5.4) \quad R = \frac{\sum_{i=0}^{M-1} w_i [Z_i - \lambda_i]}{\sum_{i=0}^{M-1} [w_i]^2},$$

where

- R : template detection statistic
- w : template weight
- i : pixel index of M pixels

- Z : spectral power (after barycentric correction)
- λ : expected noise power

Testing various template models against simulated neutron stars in binary systems reveals the benefit of the R -statistic. Figure 5.3 shows that the statistic responds markedly when the input template values match the simulated ‘injection’ sufficiently well. The sharp response assures accurate parameter estimation. Furthermore, Monte Carlo simulations by Goetz quantify the probability of high R statistics arising from Gaussian noise. Using generating functions and Davies’ method, this probability has already been incorporated into a single-template p -value, seen in figure 5.4. Further, multi-trial Monte Carlo studies done by the author have validated the probability of R - and p -values arising from Gaussian noise.

5.3 Directed TwoSpect’s greater sensitivity

Searching over a wide range of right ascension and declination can multiply computational times for gravitational wave searches by many orders of magnitude. Each GW frequency needs to be corrected for Doppler shift and antenna pattern unique to each sky location. The number of distinct sky points needed is discretized by the allowed mismatch in sensitivity between points. Typically, we search over the parameter space with a grid that allows a 0.2 relative mismatch in R statistic; the space is smooth enough that this grid can be smooth and rectangular. Looking across all LIGO frequencies requires over 10^{18} templates, of order $\mathcal{O}(10^9)$ more templates [90] to do an all-sky analysis than a search at a single location and known period. Thus the all-sky search is, in practice, only feasible when the R -statistic is the last in a stage of hierarchical statistics for candidate GW signals. The initial stage of this hierarchy, an incoherent harmonic sum, is known [90] to reduce potential sensitivity.




Figure 5.3: Exact templates for putative signals weight the pixels in the frequency-frequency plane to generate R statistic for a simulated pulsar (note: not the same as pulsar 40 in the Scorpius X-1 mock data challenge) at 100.015 Hz and $asini = 1.44$. The resulting R values are heatmap-plotted on the modulation depth vs frequency plane.

False alarm probability, absolute value of base 10 logarithm:
 $\log_{10} p$ vs parameters for pulsar 040 at H1
 maximum value: $3.3462236003 \times 10^{18}$ at $(df, f) = (0.0133034, 100.015)$ Hz
 Number of bins in data arrays (df, f) : (73, 73)




Figure 5.4: The Davies algorithm translates R statistic values for exact templates into (single-template) p -values, plotted on the modulation depth vs frequency plane.

A *directed* search could be narrowly focused enough to calculate R -statistics for all interesting binary CW models for a point in the sky. With a known electromagnetic counterpart, such as an LMXB or X-ray transient (XTE), the parameter space can often be reduced to a particular sky location (known to a few milliradians or better) and period (known to fractions of a second). Frequency may (as with XTE J1751-305 [120]) or may not (as with Scorpius X-1 [88]) be known¹. Let us consider a search for an object such as Scorpius X-1 ($P \approx 68023.7$ s, $a \sin i \approx 1.44 \pm 0.18$ light-s). The number of templates needed to cover the parameter space at a mismatch of 0.2 is known from studies that find a spacing of $1/(2T_{\text{coh}})$ in frequency and $1/(4T_{\text{coh}})$ in $a \sin i$ provided sufficient coverage, given coherence time T_{coh} . Here, assume a search over $6\sigma_{a \sin i}$, that is, $\pm 3\sigma$ around the known orbital parameter $a \sin i$. Then,

$$(5.5) \quad N_{\text{template}} = [1 + 2f_{bw}T_{\text{coh}}] \left[\sum_{j=1}^{\frac{f_{\text{max}} - f_{\text{min}}}{f_{bw}}} 1 + 2\pi (f_{\text{min}} + jf_{bw}) \frac{4T_{\text{coh}}}{P} 6\sigma_{a \sin i} \right],$$

simplified,

$$(5.6) \quad N_{\text{template}} = 2 \left(T_{\text{coh}} + \frac{1}{f_{bw}} \right) \left[1 + \frac{4\pi T_{\text{coh}}}{P} (6\sigma_{a \sin i}) (f_{\text{max}} + f_{\text{min}} + f_{bw}) \right] (f_{\text{max}} - f_{\text{min}})$$

for a single interferometer.

f_{bw} is the width of a single analysis band. At present, we use 0.1 Hz bands. N_{template} is $\mathcal{O}(10^8)$ for 3 interferometers over a 500 Hz band given the Sco X-1 orbital parameters. Such a search is tractable, since a single template test requires on the order of a second. This search promises to be significantly more sensitive than the all-sky search. We chose to test this new method first in a Mock Data Challenge

¹At our 8.5 kiloparsec galactic radius [107] both these sources are located in the direction of the galactic center.

(MDC). This MDC lets us ascertain our sensitivity relative to other GW-search algorithms. TwoSpect is 1 of up to 6 existing programs looking for 50 “open” and 50 “closed/blind” Sco X-1-like “pulsars” (LMXBs). Full MDC results are the subject of a forthcoming paper. The rest of this chapter expounds on the development and testing of methods in the course of this MDC.

5.3.1 Sky maps using exact templates

At the beginning of the MDC, this author’s work on TwoSpect played a key role in verifying that the simulation was correctly set up. Although the author played no role in the data generation – and was blinded to the parameters of the closed pulsars – TwoSpect is sensitive both to relatively weak signals and to sky location. Thus it was able to confirm that injected pulsars, as seen in Figure 5.5, were in fact in the expected location of a signal from Scorpius X-1.

5.4 Scorpius X-1 mock data challenge

5.4.1 Fully-templated search for Scorpius X-1

Scorpius X-1 MDC data necessitated an efficient means of searching over a few hundred million putative templates using similar data streams from three interferometers (Hanford H1, Livingston L1, and Virgo V1). Ideally, all one hundred 5-Hz search bands would be illuminated in the manner of Figure 5.6. Achieving this goal required some development by the author.

5.4.2 Narrow-band heat maps in parameter space

Initially, TwoSpect was configured for all-sky searches with the capability to bypass the incoherent harmonic sum and test a single point in parameter space. It was possible to test multiple templates only by rerunning the entire TwoSpect pipeline. Initial input/output loaded and Doppler-shifting made this highly wasteful, although

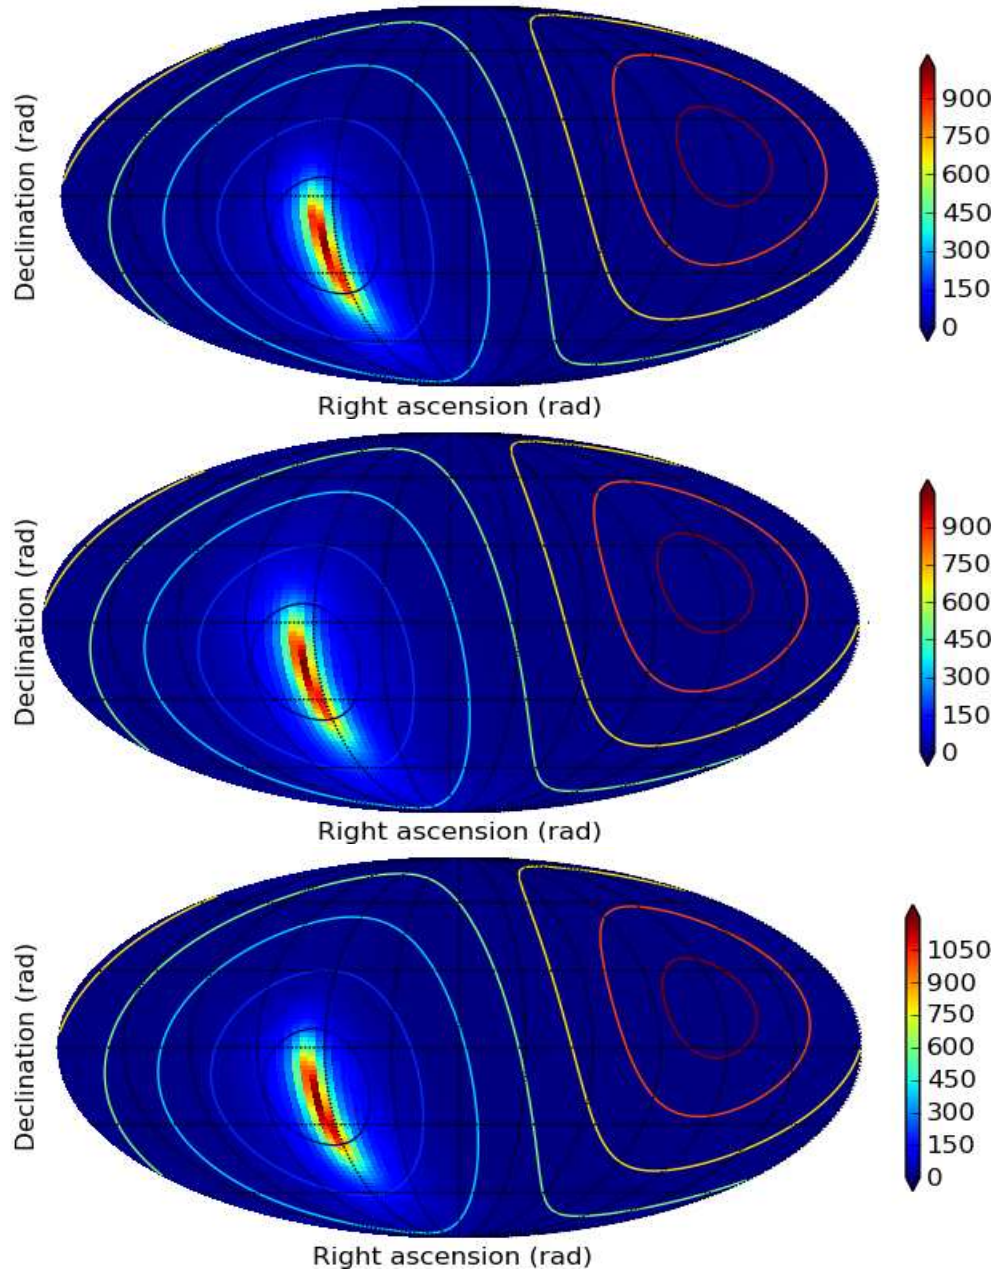


Figure 5.5: All-sky maps, {H1, L1, V1} interferometer analysis from top to bottom, for template tests varying right ascension and declination. Scorpius X-1 mock data challenge pulsar 16 (101x101 templates), showing $\log_{10} p$ -value on a Mollweide projection. Contour lines at 1-radian great-circle distance intervals from the intended injection location of Sco X-1. The results match the intended injection and confirm that the simulation is accurately representing the known sky location of Sco X-1.




Figure 5.6: Scorpius X-1 Mock Data Challenge (MDC) pulsar 40 {H1}: 5 Hz band. The p -value (single-template, applying Davies' Method to the R statistic) is shown in this heatmap, peak in red. All templates are plotted on the (frequency, modulation depth) plane. This is a relatively broadband view.




Figure 5.7: Heatmaps {H1, L1, V1} of 11x11 templates centered around Scorpius X-1 MDC pulsar 8. This is a relatively narrowband view.

it was possible to obtain results in narrow regions around which the unblinded, open MDC pulsars were specified. Figure 5.7 illustrates one such result.

5.4.3 Wide-band heat maps in parameter space

The author’s first contribution to the pipeline was streamlining searching over arbitrary-width frequency bands. Input/output is now done only once for a given search band, as is Doppler-shifting, expediting the testing of many templates for the R -statistic by several orders of magnitude. While testing so many templates is arguably unnecessary – the results are highly correlated – it is computationally straightforward and yields the best gain in sensitivity over previous searches. Figure 5.8 shows that these results match up with prior work, from Figure 5.7.

Whereas much existing TwoSpect post-processing to data had been focused on follow-up of all-sky results, parts of the directed search post-processing have needed to be re-invented. The MDC validated these methods.

5.4.4 Revisiting & refining detection criteria

Running TwoSpect as a directed-search algorithm involves calculating the R -detection statistic across the probable parameter space. A search is conducted with templates for each grid point in the Scorpius X-1 parameter space; period is known sufficiently well to restrict the search to the two dimensions of signal frequency and frequency modulation. The grid spacing, inversely proportional to spectrum coherence time, was chosen to allow a mismatch no more than 0.2 in the detection statistic. Because of known period and sky location, the incoherent harmonic sum stage of TwoSpect, used for the all-sky search, was bypassed entirely.




Figure 5.8: Scorpius X-1 MDC pulsar 8 {H1}: 5 Hz band. This heatmap shows 3.6×10^5 templates, 10 to 22 mHz modulation depth, 120-125 Hz frequency. The peak signal at about ($df = 0.019$, $f = 121.9$) Hz.




Figure 5.9: Scorpions X-1 MDC statistics. These histograms of the R statistic and p -value distribution helped in understanding noise, temporal gap & spectral leakage, as well as establishing a threshold p -value \sim false alarm probability of 1%. These p -values are for single templates, appropriate to the all-sky search but not to a dense templated search with a large trials factor. Here, histograms show statistics in the absence of a signal. The left-skew of the p -values is associated with gaps in the data (as is right-skew, with different gaps). Diurnal bias in barycentering is assumed, but the correlation is not fully understood.

5.5 Mock Data Challenge procedure

Each interferometer in the data challenge was analyzed individually for the detection statistic and corresponding single-template p -value. A set of highest p -value outliers in 5 Hz bands was produced for each interferometer, subject to a p -value threshold inferred from Gaussian noise. These sets were compared in pairwise coincidence (H1-L1, H1-V1, or L1-V1), where coincidence required proximity within a few grid points in the parameter space. Any surviving outliers were classified as detections.

The highest p -value outlier in a single interferometer in that band yielded the estimated parameters. Uncertainties in these parameters were also determined from unblinded injections, using method of moments for signal frequency and modulation depth and confidence intervals for signal amplitude. Upper limits were declared from the best estimate of the 95% confidence level of non-detected, unblinded, injected signals. The largest uncertainty in upper limits and signal amplitude estimation derives from the ambiguity between true h_0 signal and $\cos \iota$ inclination. This ambiguity cannot be resolved with the present algorithm and depends partially on the assumed prior distribution of signal amplitudes; the uncertainty was estimated by simulation.

Put another way, TwoSpect in its directed search mode, tests templates with a model of f , $a \sin i$, and P , as well as sky location. The latter two are fixed for Scorpius X-1, as they are well-known. The R -statistic is not sensitive to time of ascension. If a coincident detection is made between any interferometer pair in a 5 Hz band, model parameters are read off from the extremal p -value template at any one interferometer; h_0 is proportional to the quarter-root of the test statistic.

Uncertainty in f and $a \sin i$ is determined from the standard deviation of known injections; it is on the scale of the template grid except for marginally detected pulsars. The h_0 uncertainty is largely due to the uncertainty in $\cos \iota$.

5.5.1 Detection Claims

After studying the Gaussian noise in the Scorpius X-1 MDC open data set, we were able to set thresholds for detection claims.

TwoSpect’s R -statistic and p -value space on the frequency vs modulation depth plane showed significant structures, particular around loud injections. These structures corresponded to the distribution of power into pixels by way of modulation depth, Earth’s Doppler motion, and possibly spectral leakage. These regions of the open data set parameter space were excised before proceeding with the Gaussian noise study.

TwoSpect also found later-explained differences in the noise between the 360-s SFTs, used for pulsar bands above 360 Hz, and the 840-s SFTs, used for those bands below. The shorter SFTs were ostensibly noisier, requiring a more extreme $\log_{10} p$ of -12.0 , rather than -7.75 as in the longer SFTs, to cut single-IFO outliers and prevent any false alarms from surviving the coincident test between IFOs. On revisiting the issue after the MDC, it was found that a misconfiguration bug explains the discrepancy (along with the larger number of templates and imprecise order statistics). With the bug-contaminated Gaussian noise excluded, removing 3 million of 21 million Gaussian noise templates, the required $\log_{10} p$ for 360-s SFTs would only be -8.80 . This conclusion is consistent with our expectation that the p -value calculation should be independent of coherence time.

After studying the effect of pairwise coincidence requirements on surviving Gaussian noise outliers, we were satisfied that we would achieve a false alarm rate of 0.01

or better by setting the following detection criteria. Note that the $\log_{10} p$ values refer to single-template p -values; nearby templates are correlated in the presence of signal.

5.5.2 Detection criteria

- single-IFO candidates are the up-to-200 most extreme p -value outliers in a 5-Hz band that had a $\log_{10} p \leq \text{threshold}$, where $\text{threshold} = -7.75$ if $f < 360.0$ Hz (those that used 840-s SFTs) or -12.0 if $f \geq 360.0$ Hz (those that used 360-s SFTs).
- each candidate must survive at least one double-IFO coincidence test, involving a pairwise comparison of single-IFO candidates to see whether they are within $1/T_{\text{SFT}}$ in both frequency (f) and modulation depth (df).

→ if there is any candidate surviving these criteria in a 5 Hz band, we mark detected, else not detected.

5.5.3 Parameter Estimation

MDC data allowed checking TwoSpect’s parameter estimation on the 31 pulsars detected in the open data set.

Note that the h_0 reported in this section had not yet been recalibrated for either the $\cos \iota$ ambiguity due to assumed circular polarization (see subsequent section, factor of 1.74) or a systematic rescaling endemic to TwoSpect (factor of 1.11). Instead, the first step was to rescale the known h_0 -injected from the MDC open data table into an h_0 -effective. This h_0 -effective equaled,

$$(5.7) \quad h_{0\text{-effective}} = \frac{1}{\sqrt{2}} \sqrt{\left(\frac{1 + \cos^2 \iota}{2}\right)^2 + (\cos \iota)^2} \times h_{0\text{-injected}},$$




Figure 5.10: Error in strain estimation versus circular-effective injected strain. Higher injected strain results in higher absolute errors.

the rescaling necessary to convert the strain into effective units of detected strain. Any pipeline that assumes circular polarization should require a similar procedure.

We plotted the error (h_0 -effective, as inferred, minus h_0 -effective, as calculated from the true injection parameters) in the h_0 reported by TwoSpect versus h_0 -effective for the 31 open pulsars detected.

This same error was also plotted vs p -value and frequency, for h_0 in Figure 5.11, f estimation in Figure 5.12, and $a \sin i$ estimation in Figure 5.13 These values are

reported in the header of the graphs.

Blue lines indicate the overall uncertainty. Red lines (on plots vs p -value) indicate a least-squares power-law regression.

5.5.4 Upper Limits and Detection Efficiency

Upper limits and detection efficiency were also calculated using data in the open pulsar set.

For detection efficiency, we calculated the h -effective for the 31 detected and 19 non-detected pulsars and found the average detection rate in bins according to h -effective. These bins were non-uniform in size due to the interest in finding the 95% detection efficiency point despite the paucity of statistics (only 50 pulsars total). Binomial uncertainty was also calculated and each bin's $1-\sigma$ worst case was graphed in Figure 5.14.

Consequently we plotted the distribution of recovered h_0 versus injected h_0 -effective (the error of which is shown above, for detected pulsars) in Figure 5.15. Further injection studies should show how this upper limit varies with frequency, for a given injected h_0 , but at the time of the MDC, we did not feel confident in extrapolating this relationship.

5.5.5 $\cos \iota$ Ambiguity

The cosine of the inclination angle of the pulsar, $\cos \iota$, casts an ambiguity over the determination of h_0 . For TwoSpect, which assumes circular polarization, the approximate true value of h_0 will indeed be as reported if $|\cos \iota| = 1$, but will be greater for smaller $|\cos \iota|$ is less (i.e., the gravitational wave is elliptically polarized). In the case of linear polarization, h_0 will be about $2^{3/2}$ times larger than reported.

While an analytical calculation of the expectation value of the correction factor




Figure 5.11: Parameter estimation: error in strain and dependence on recovered p -value (top) and frequency (bottom). The strain appears broadly distributed, without any systematic patterns. The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies.




Figure 5.12: Parameter estimation: error in frequency and dependence on recovered p -value (top) and frequency (bottom). The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies.

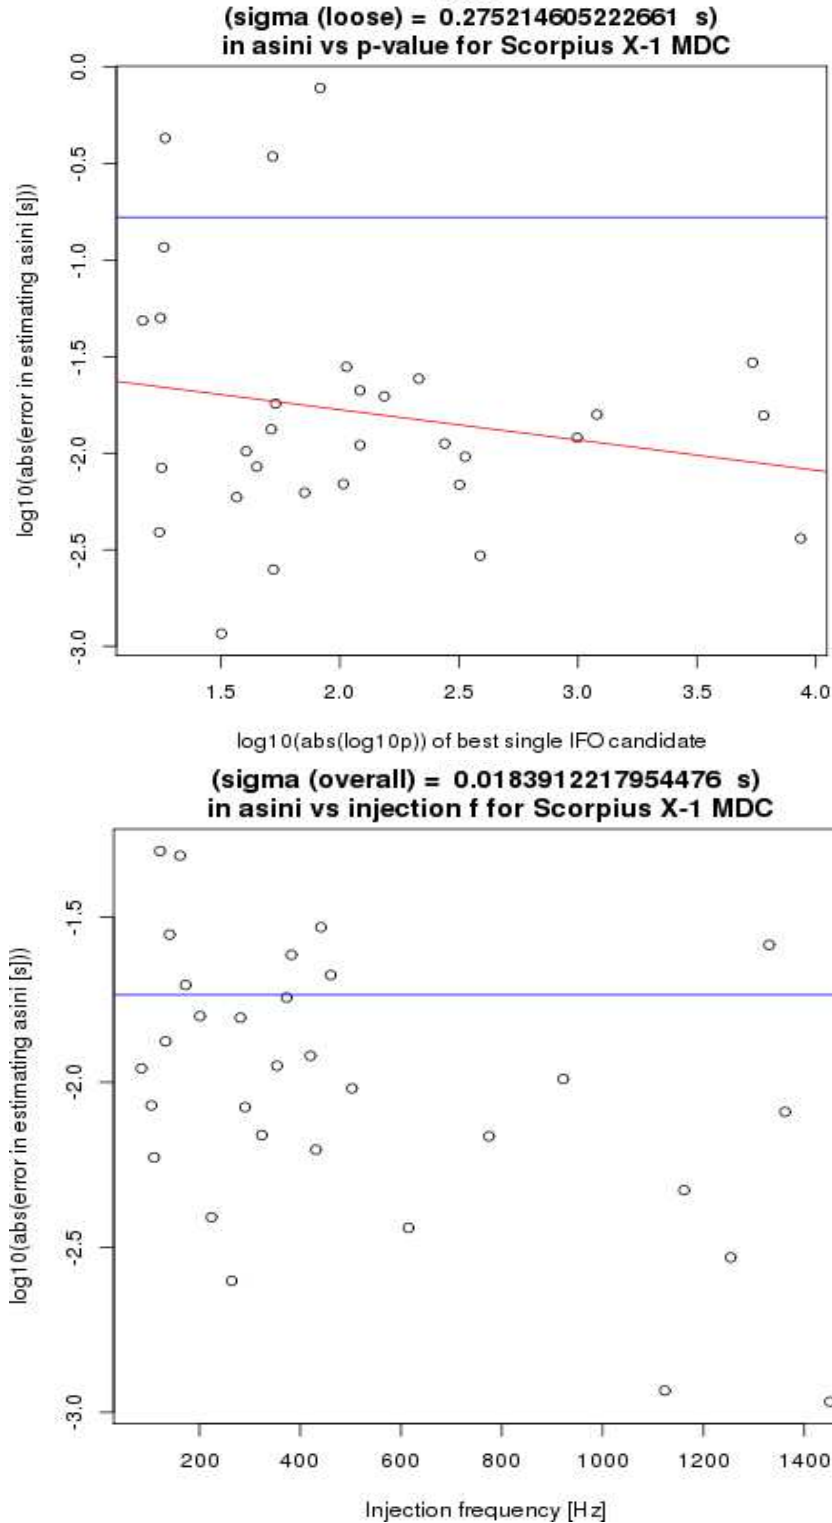


Figure 5.13: Parameter estimation: $a \sin i$ (projected semi-major axis; directly proportional to modulation depth for a given frequency and period) and dependence on recovered p -value (top) and frequency (bottom). The overall error vs frequency is shown at bottom after a rerun to fix a misconfiguration where inadequate data was read in at high frequencies.




Figure 5.14: Open pulsar detection efficiency curve. Because only 50 pulsars were in the open set, this curve is relatively-poorly defined – the binning has been chosen to give the most accurate representation based on the chosen thresholds. The 95% level is approximately about 3×10^{-25} (again, without the corrective factors of 1.74 and 1.11) but is too imprecise to judge using this method.

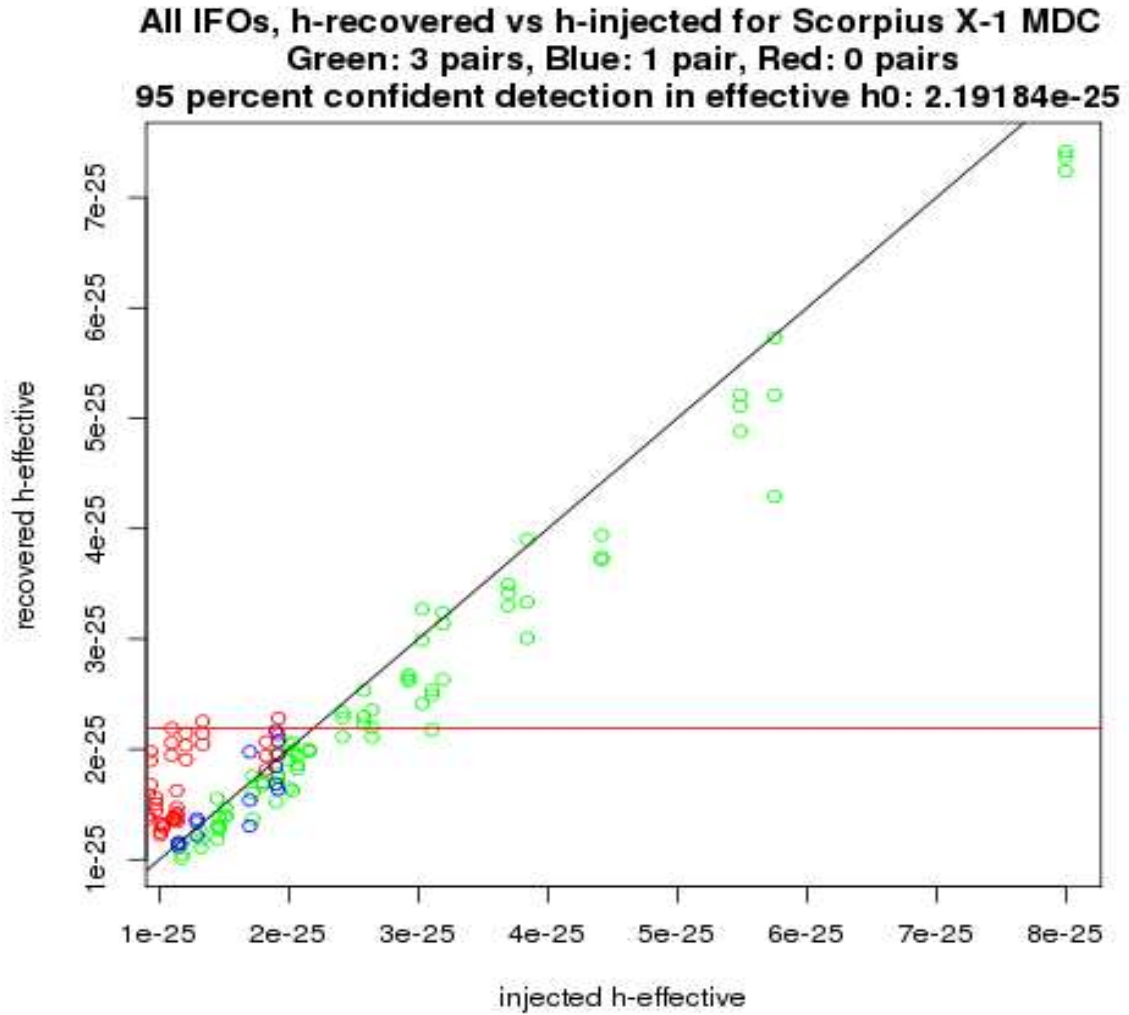


Figure 5.15: Detections and upper limit determination. Depending on whether a injection was seen in three, one, or no detector pairs, it was assigned a color-coded circle and plotted in recovered strain versus effective circular strain injected. (There are no injections seen with two detection pairs, because this plot only shows the loudest outlier from each 5 Hz band; if some injection were seen in two and not three pairs, it would mean two distinct coincidences were seen, only one of which would be the loudest). Color-coding red pulsars as non-detected, blue as single pairwise detection, and green as triple pairwise detection, we identified a shelf of non-detected pulsars that was 95% contained by an upper limit about 2.19×10^{-25} . This number, when corrected, yielded the upper limit of $1.74 \cdot 1.11 \cdot 2.19 \times 10^{-25} = 4.23 \times 10^{-25}$ for TwoSpect. The unity-slope line is shown to ascertain whether a further empirical rescaling factor was needed (it was: constant 1.11). The zero-slope line is shown to indicate the ninety-five percent confidence upper limit in the absence of detection.

is easy, it will not easily take into account the circular bias of detected signals. That is, a pipeline will tend to see a slightly greater proportion of signals that are more circularly polarized, because the effective h_0 of those signals is greater. This “circularizes” the correction factor in a way dependent on the detection efficiency of the pipeline and on the assumed prior distribution of pulsars. Although the effect is relatively minor, we decided to simulate it because the size of the effect was unknown at the time.

In this simulation, 2 million pulsars were generated with h_0 between 3×10^{-26} and 3×10^{-24} with a distribution of $1/h_0$.

We made a toy model of our detection efficiency, assuming no pulsars were detected below 1×10^{-25} effective, all were above 3×10^{-25} , and the fraction detected was linear in h_0 between those values (Figure 5.16).

Together with a uniform $\cos \iota$ distribution on $[-1, 1]$, this led to a trapezoidal distribution of recovered, detected h_0 values with a curved lower (left) edge (Figure 5.17).

The upper end of the distribution (right side of the trapezoid, Figure 5.17) was excluded because we are trying to find the average bijective mapping (slope) f : (detected h_0) \rightarrow (true h_0), and including detected $h_0 > 1 \times 10^{-24}$ meant that we were failing to see the complete injected h_0 space. There was f^{-1} : (true h_0) \rightarrow (detected h_0), but not f . More plainly, suppose we looked at a detected h_0 reported as 1.5×10^{-24} , and that our average corrected factor had been calculated to be 2.5 (it was not) – this would imply that the true h_0 was 3.75×10^{-24} – but this would be outside the domain of the simulation, so there would be no way to check it. The analogous problem should not happen at the lower end of the distribution (left side of the trapezoid).


(Detection efficiency - uncertainty) for Scorpius X-1 MDC mini

Figure 5.16: Simulated detection efficiency curve. Because the $\cos \iota$ ambiguity simulation require a priori model of detection efficiency, we described it simply. Here, no detections were claimed below 1×10^{-25} , all were detected above 3×10^{-25} , and the probability of detection rose uniformly on the intervening interval.




Figure 5.17: Distribution of 2 million simulated stars, strain between 3×10^{-26} and 3×10^{-24} under a log-uniform distribution, following application of $\cos \iota$ and detection efficiency cuts.

In turn, we looked for the relationship between the recovered h_0 of this “detected” distribution and the corresponding original, true h_0 . The slope would give us the conversion factor. The first attempt was to grid the (detected h_0) \times (true h_0) space into 2D pixels. This was suggestive, and yielded the regressed slope in Figure 5.18.

There is a systematic bias in the grid method, both by one pixel (hence why the mean was adjusted downward to 1.74) and in the associated uncertainties. Plotting these uncertainties on the distribution of detected h_0 vs true h_0 shows how wide those error bars are, in Figure 5.19.

This bias in Figure 5.19 is likely due to sampling: numerical fluctuations in the grid method made it unstable, especially toward the high h_0 end of the distribution. Instead, we manually adjusted a $\pm\sigma$ until the CDF encompassed the appropriate 68% confidence interval, finding a σ in the slope of 0.37 with a mean slope of 1.74. The reason for the aforementioned restriction of the plot to h_0 -effective $< 1\times 10^{-24}$ can be seen in the distortion at levels above that in the following plot:

The chosen $1.74 \pm 0.37\sigma$, however, yielded the necessary correction factor.

Finally, we tested all of our calibration factors for h_0 with the associated confidence intervals and found the fraction of open data estimated h_0 , f and $a \sin i$ within their 1σ error bars. The results were conservative:

h_0 : 77.4%, f : 74.2% $a \sin i$: 67.7%, Period: 100%, [*n.b.*, we tested only one period, 68023.8259 s].

These error bars were then used without modification for claiming uncertainties on the closed pulsars.




Figure 5.18: Regression using grid. By binning the simulated stars on the true strain vs detected (recovered) strain plane, an accurate mean slope for the $\cos i$ correction was ascertained. It had to be modified downwards by the equivalent of one bin, to 1.74. However suggestive, the $1\text{-}\sigma$ thresholds proved inaccurate, probably due to noise fluctuations.




Figure 5.19: Simulation with fit lines as given by the bin-method regression.




Figure 5.20: Confidence intervals with final fit. After manual optimization of the cumulative distribution function, and constraint to the region with a full bijective mapping between injected and recovered strains (below 1×10^{-24}), a $1\text{-}\sigma$ value of 0.37 in the slope was found to give accurate confidence intervals. The reason for the aforementioned restriction of the plot to $h_0\text{-effective} < 1 \times 10^{-24}$ can be seen in the distortion at levels above that. The chosen $1.74 \pm 0.37\sigma$, however, yielded the necessary correction factor.

5.6 Summary of the MDC

TwoSpect competed extremely effectively in the MDC. Comparisons are the subject of a forthcoming paper, but for our own work, the author can report the following successes.

5.6.1 Mock data challenge results

Analyses of the MDC correctly recovered about two-thirds of the simulated stars:

- 34 of 50 closed (and 31 of 50 open) ‘pulsar’ signals detected
- f , $a \sin i$ and h_0 estimated
- 4.23×10^{-25} strain upper limit (UL) in 4×10^{-24} strain $\text{Hz}^{1/2}$ noise declared for the 16 non-detected, closed signals
- with injections to refine UL, applicable to real data

As of the current draft of the comparison paper [52], TwoSpect has detected more pipelines than the other three pipelines with results; the Radiometer [43] method detects slightly fewer, followed by the Sideband [129] and then Polynomial [166] searches.

One challenge in the transition to real data is that the MDC used entirely Gaussian data. Non-Gaussian test injections are the subject of Chapter VI. Although the distribution of h_0 values in the MDC was astrophysically optimistic, we gained knowledge about the transition from low- to high-SNR detections. This MDC also validated our ability to recover orbital and GW parameters accurately, including for blinded simulations.

5.6.2 Binary search summary beyond the MDC

TwoSpect is well-suited to the Scorpius X-1 mock data challenge. With this experience, the author is pursuing Scorpius X-1 (and J1751-305) searches in real data, as described in Chapter VI. It is also believed that directed binary searches can be made more sensitive with straightforward changes (see Section 5.8). In the near term, the author and the LIGO continuous waves group will direct TwoSpect and kindred binary searches toward promising targets such as LMXBs. In the long term, this work will enhance all searches, the bridge of accreting binaries providing a firm link to electromagnetic astronomy as the age of gravitational wave astrophysics begins.

5.7 Plans for improvement

TwoSpect presents a viable option for seeking continuous gravitational waves from neutron stars in binary systems; yet more sensitivity would reveal a richer sky. Indeed, aLIGO designs and the Scorpius X-1 torque-balance limit do not guarantee detection in the coming generation of interferometers. Several improvements can thus be investigated.

- *Coherently combine multiple interferometer outputs:*

Add complex Fourier coefficients (with phase corrections)

to create a multi-detector statistic (under way by Goetz)

- *Elliptical polarization:*

search antenna pattern weightings corresponding to

elliptical polarization – better sensitivity and

parameter reconstruction, including h_0

- *Orbital phase:*

Search over initial orbital phase by coherently combining template and doubly Fourier-transformed data – better sensitivity and discrimination

- *Parameter space patterns:*

Exploiting patterns in the R -statistic parameter space to improve search time, sensitivity, or both

5.7.1 Coherent interferometer synthesis

Coherent interferometer synthesis for H1-L1-V1 is already well-underway by Goetz. Data from multiple interferometers can be added in-phase for a putative signal model, and this technique already appears to be yielding improvements in detection efficiency. To wit, the synthesized h is given by Equation 5.8:

$$(5.8) \quad h(f, t) = \Sigma_j (h_j(f, t)e^{i\phi_j(f, \alpha, \delta)}),$$

$$(5.9) \quad \phi_j(f, \alpha, \delta) = 2\pi f T_j(\alpha, \delta) + \phi_0,$$

where $h_j(f, t)$ is the complex h value in SFT for interferometer j , time t , frequency f and $\phi_j(f, \alpha, \delta)$ is the phase shift for right ascension α , declination δ (an overall phase shift ϕ_0 factors out because TwoSpect computes its test statistic from power, not the complex SFT coefficient). Here, $T_j(\alpha, \delta)$ is the time-of-flight delay between interferometers (projected on vector from α, δ). Further investigations are forthcoming.

5.7.2 Circular & elliptical polarization

The general formula for the polarization components of gravitational radiation is well-known. The equation is stated (for instance, with $\Phi = 2\phi$, in an earlier Scorpius X-1 search [24] or more recently in the TwoSpect all-sky search [7]) in terms of two polarization components, F_{\times} and F_{+} :

$$(5.10) \quad h(t) = h_0 \left(F_{\times}(t, \alpha, \delta, \psi) \frac{1 + \cos^2(\iota)}{2} \cos[\Phi(t)] + F_{+}(t, \alpha, \delta, \psi) \cos(\iota) \sin[\Phi(t)] \right).$$

Presently, TwoSpect effectively searches only for circular polarization, making it most sensitive when $\cos \iota = 1$. If antenna pattern weighting is adjusted to test for F_{\times} and F_{+} according to a general, elliptical polarization model, sensitivity to other values of $\cos \iota$ should be enhanced. It remains unclear whether this test would be sufficient by itself to remove the $\cos \iota$ ambiguity in estimated h_0 if a detection is obtained, but earlier work on the PowerFlux all-sky, isolated star pipeline suggests improved parameter estimation is achievable [65, 128].

The current algorithm calculates pixel powers P for SFT n , bin k :

$$(5.11) \quad \tilde{P}_k^n = \frac{F_n^2(P_k^n - \langle P_k \rangle^n)}{(\langle P_k \rangle^n)^2} \left[\sum_{n'}^N \frac{F_{n'}^4}{(\langle P_k \rangle^{n'})^2} \right]^{-1},$$

$$(5.12) \quad F^2(t, \alpha, \delta) = F_{\times}^2(t, \alpha, \delta) + F_{+}^2(t, \alpha, \delta),$$

where F is antenna pattern polarization weighting. Generalizing to elliptical polarization angle ψ with weights a, b :

$$(5.13) \quad F^2(t, \alpha, \delta, \psi) = aF_{\times}^2(t, \alpha, \delta, \psi) + bF_{+}^2(t, \alpha, \delta, \psi)$$

5.7.3 Orbital phase & beyond

Orbital phase is currently ignored. Only power information in the 2nd Fourier plane is used to inform R -statistic templates. Ephemeris data (*e.g.*, time of ascension or periapsis) to inform the orbital phase is often available for LMXBs.. Orbital phase might be incorporated into the R -statistic calculation, by checking the phase of the 2nd Fourier plane pixels (the 2nd transform obscures GW phase information but should preserve orbital phase). Even for extremely short SFTs, a year-long science run produces only hundreds of thousands of SFTs, not billions of time samples, so a rough orbital phase-match would be attempted. This step might reduce the noise background in a directed search with a known orbital phase, allowing a weak signal to stand out. Conversely, it could be used to infer orbital phase. Depending on the computational costs of this step, it might be aimed toward follow-up searches.

5.7.4 Parameter space patterns

Computational costs for TwoSpect have led to suggestions for ways to exploit the R statistic’s behavior over our parameter space. Simulated annealing² would address the wrong issues in our high-noise, weak-signal analyses. ‘Hill-climbers’³ could explore the parameter space until encountering one of the X -pattern arms. Low SNR compared to available computational resources means this climbing strategy will probably not be necessary. Other features of the X -arms could still be exploited.

These X -patterns follow the lines $df = f_{\text{signal}} \pm (f - f_{\text{signal}})$, intersecting at f_{signal} .

We naïvely expect the X pattern to be defined by templates with turning-point frequencies in the time-frequency plane, with an R proportional to the peak in a way that is as the relative power of that bin compared the power in the whole

²Discussed with Maria Alessandra Papa, Sam Finn and others.

³Proposed by Ethan Obie Romero-Severson.

sinusoid. Since the template touches only once per period, the ratio would naïvely be thought to be $T_{\text{coh}}/P_{\text{signal}}$. For Scorpius X-1 ($P = 68023.8259$ s), that would lead to ratios of roughly $0.026 \approx 1/38$ for 1800 s SFTs, $0.01 \approx 1/81$ for 840 s SFTs, and $0.0053 \approx 1/189$ for 360 s SFTs. In actuality, the wrong putative template and the true signal are both stationary at the same turning-point frequency in the time-frequency plane, where the derivatives vanish:

$$(5.14) \quad \frac{df_{\text{true}}(t)}{dt} = \frac{df_{\text{template}}(t)}{dt} = 0,$$

at which point there is only gradual deviation of df from the true modulation depth, governed by the frequency bin resolution, $1/T_{\text{coh}}$. This stationarity means that the R -statistic in the X -pattern is much larger than naïve predictions. With this understood, we expect to exploit the observed X -patterns in Figures 5.3 and 5.4.

Point spread function deconvolution in astronomy presents a parallel to the patterns observed in TwoSpect. While an *ansatz* deconvolution proved counterproductive⁴, future work could be fruitful, especially for ‘proper’ maximum likelihood estimation of orbital parameters.

5.7.5 Relevance to follow-up

TwoSpect currently estimates f and $a \sin i$, limited by the accuracy of its grid-spacing and some noise fluctuations. More refined grid spacing is possible in the event of a detection, although this has not been studied. Determining h_0 with greater accuracy should be possible if we know $\cos \iota$ by another means. Section 5.7.2 could yield knowledge of gravitational wave polarization angle, informing $\cos \iota$, and coherent combinations of detector data within TwoSpect might reduce noise in detected LMXBs, as well as making quieter sources detectable.

⁴Summing up pixels in the X -pattern with guesswork weights yielded worse SNR.

Coherent synthesis, elliptical polarization, and orbital phase improvements need implementation, validation and testing. It can be hoped that they will provide additional sensitivity for TwoSpect to yield more detections, and possibly to add new parameters to the list of what TwoSpect can estimate. Although these enhancements will increase computational cost, we can also attempt to offload some of that cost to distributed computing, such as the Einstein@home project. Together with the quieter noise floors of advanced detectors, we may reasonably hope to detect gravitational waves from neutron stars in binary systems.

Since data from the initial detectors exists now, it is, in fact, prudent to see whether, despite the astrophysical predictions, a signal might already be seen. Chapter VI presents this search.

5.8 Addendum

5.8.1 Scorpius X-1 parameters

Many Scorpius X-1 parameters are known to high accuracy [88]. The first LIGO search was published in 2007 [24].

- Distance: 9000 light-years (2.8 kpc)
- Eccentricity: $< 3 \times 10^{-3}$
- Sky location: $\alpha=16\text{h}19\text{m}55.1\text{s}$, $\delta=-15\text{d}38\text{m}24.9\text{s}$
- X-ray luminosity: $2.3 \times 10^{31}\text{W}$, $60000 L_{Sol}$ ($2.5 \times 10^{-10} \text{Wm}^{-2}$ at Earth)
- Period: $68023.70 \pm 0.04 \text{ s}$

Sco X-1 torque-balance limit (from X-rays) [47, 91]:

$$h_0 \approx 2.8 \times 10^{-26} \left(\frac{600\text{Hz}}{f} \right)^{1/2} .$$

Note that Chakrabarty [59] has hypothesized the existence of an LMXB speed limit; given a fastest known millisecond pulsar of 716 Hz, this could imply that a search up to $2\nu \approx 2 \times 716\text{Hz} = 1432$ Hz, plus a margin of error, should suffice to cover the astrophysical parameter space.

CHAPTER VI

Directed TwoSpect: Scorpius X-1 and XTE J1751-305

Results from this chapter are preliminary and have not been reviewed yet by the LIGO Scientific Collaboration.

6.1 Directed TwoSpect

TwoSpect performed well in Chapter V's Mock Data Challenge, which warrants using the program to analyze the best existing GW data. At the time of this writing, the best consistent stretch of GW data remains that of Science Run 6 (S6) taken from 2009 July 09 to 2010 October 20 by LIGO Hanford Observatory (LHO) and LIGO Livingston Observatory (LLO). Advanced LIGO (aLIGO) commissioning is already surpassing the sensitivity of S6 for short periods of time. While early LIGO continuous wave (CW) searches used short science runs such as S2 [24], aLIGO data duration is for now too short for an analysis competitive with present upper limits, in particular the 2011 Radiometer [12] S5 high frequency results and 2014 TwoSpect [7] S6 low frequency results. The increased sensitivity anticipated from aLIGO Observing Run 1 (O1), which is planned to last several months in summer 2015, could yield interesting results, and the 9-month O3 is planned for 2017. Preliminary Scorpius X-1 broadband upper limits from S6 are presented in this chapter.

6.1.1 Targeted, directed and all-sky search sensitivity

Several approaches exist toward a search for GWs from known objects such as Scorpius X-1¹. When seeking continuous GWs, these methods are classified as all-sky, directed, or targeted, in order of increasing focus of the search. More focused searches make sense as more prior information is known, such as sky location, neutron star frequency and binary system orbital parameters. Additional information lets some CW searches refine, for instance, their template models of GW signals. If the search is designed for minimal information, in particular if sky location is unknown, it is generally called an *all-sky search*. Some information – such as a low-uncertainty sky location, less than a square arcminute – helps what are called *directed searches* gain sensitivity. Sources with comprehensively documented parameters, such as rotation frequency and spindown rate, lend themselves to *targeted searches* that search over a very narrow range of parameters, *e.g.*, putative GW phase and orientation angles.

TwoSpect was designed [91, 90] as an all-sky search for unknown neutron stars. The voluminous parameter space of that search, as described in Chapter V, prompted tradeoffs in sensitivity vs computational cost, in particular the limitation of test statistic calculation to only those outlier candidates that survived an incoherent harmonic sum stage. For objects with constrained sky location and NS parameters, thorough calculations of the R test statistic become feasible.

For Scorpius X-1, many parameters are known (updated ephemerides were determined in 2014 by Galloway [88]), but critically, rotation frequency is not. Sky location is known, and the period is 0.7873114 ± 0.0000005 days, *i.e.*, 68023.70 ± 0.04 s with $1\text{-}\sigma$ uncertainty². The projected semimajor axis, $a \sin i$, is 1.44 ± 0.18 s with 1-

¹Discovered by Riccardo Giacconi in 1962 with a sounding rocket [89].

²Access to a preliminary ephemeris in the MDC led us to use $P = 68023.8259$ s in for searches in the MDC and S6; MDC simulations justify assuming that this variation has negligible impact on TwoSpect, which would only be able to discriminate between different periods around 68023 s at a resolution of about 40 s or greater, depending on

σ uncertainty³. Rotation frequency uncertainty drives the cost of the search. While the Chakrabarty speed limit [59] and neutron star breakup limit 2ν , the GW emission frequency of an NS quadrupole, to $\mathcal{O}(2 \text{ to } 3)$ kHz, the dominant high frequency limit f_h is driven as much by the noise floor of the LIGO detectors. This noise floor increases linearly with frequency and at 2 kHz is an order of magnitude worse than its most sensitive, about 2×10^{-23} strain $\text{Hz}^{1/2}$ between 150 and 200 Hz. Photon shot noise, that is, quantum vacuum fluctuations are the limiting noise source at those frequencies. The low-frequency limit, f_l , is also driven by the noise floor of the detector, which becomes contaminated by seismic noise below about 40 Hz.

A search for Scorpius X-1 in S6 data should thus take place between about 40 and 2000 Hz. TwoSpect requires shortened Short Fourier Transforms (SFTs) [90] to fully capture the spectral power if a GW source has a high frequency or large $a \sin i$. This suggests dividing the main search: 40 to 360 Hz can be searched in 840-s coherence time SFTs and 360 to 2040 Hz in 360-s coherence time SFTs. These sets constitute the primary search. Although MDC and simulation experience suggests spectral power is lost only slowly as GW frequency exceeds optimal coherence time, we have also prepared 260 Hz to 360 Hz SFTs with 360-s coherence time and 1400 Hz to 2040 Hz SFTs with 300-s coherence time; these are most useful for high $a \sin i$ signal models. These additional sets constitute the ‘overlap’ search. Calculating the number of templates in the primary search (the overlap search is smaller because of shorter coherence times) over these frequencies and $\pm 3\sigma_{a \sin i}$, if analyzed in 0.1 Hz computational bands⁴ using Equation 5.6,

Fourier transform coherence time.

³Orbital parameters interpreted in correspondence between C. Messenger and D. Galloway.

⁴0.1 Hz computational bands fit efficiently on cluster memory in under 2 GB of RAM; for convenience, templates are (redundantly) tested at both the lower & upper bounds of the 0.1 Hz.

(6.1)

$$N_{\text{template}} = 2(840.1) \left[1 + \frac{3360\pi}{68023.8} 400.1 \right] (320) + 2(360.1) \left[1 + \frac{1440\pi}{68023.8} 2400.1 \right] (1680),$$

or 3.392×10^7 templates in 840-s SFTs and 1.9434×10^8 templates in 360-s SFTs, per interferometer. Searching both LHO's H1 interferometer and LLO's L1 interferometer, the grand total is 4.5652×10^8 templates.

Each of these templates takes between 0.3 and 3 s to run on late-2000s to early-2010s CPUs, depending on vector extensions (such as SSE) and clock speed. Given approximately two thousand cores at a cluster such as the LIGO Data Grid at the California Institute of Technology and the LIGO observatories or Atlas at the Albert Einstein Institute in Hannover, Germany, a fully templated TwoSpect search for Scorpius X-1 in S6 data can be completed in roughly a month. This has been carried out by the author and is the subject of the remainder of this chapter. Because of the enhanced sensitivity and wider frequency range of the directed search, this analysis can improve on the TwoSpect all-sky limits by Goetz [7] despite consuming much fewer computational resources (albeit for a single, promising source).

6.1.2 Enhancements enabled by directed searching

Directed TwoSpect as used for the search in this chapter closely resembles the all-sky TwoSpect pipeline, except for post-processing. The post-processing is discussed below. The bulk of the processing, skipping the incoherent harmonic sum stage that reduces the number of templates search, involves generating the test statistic, R , for each point in a rectangular grid spaced at $1/(2T_{\text{coh}})$ in frequency and $1/(4T_{\text{coh}})$ in $a \sin i$, which keeps mismatch between putative signal and template to within 20% of the peak R value [90]. Since sky location and period are fixed, the search is

two-dimensional on the f and $a \sin i$ plane. Because the detector-frame realization of $a \sin i$ is as a frequency modulation, df , this search space can be describes as the ‘ df vs f ’ plane. Each pixel in the ‘ df vs f ’ plane has an h_0 and $\log_{10} p$ -value associated with its R statistic, as noted in Chapter V. The h_0 is proportional to $R^{1/4}$. Estimating $\log_{10} p$ is more complicated, but it loosely scales with R in Gaussian noise and has been calibrated into an accurate, single-template p -value using Davies’ method, a computational implementation of the Gil-Pelaez formula⁵. Our search uses $\mathcal{O}(10^8)$ templates. This number is in contrast to the few, distant-and-uncorrelated, templates used for outlier follow-up in the all-sky search. Thus the directed search necessitates enhancements to post-processing.

In the near future, enhancements enabled by the directed search will be possible. Sensitivity from full templating is the first step. The next steps can involve a search over orbital phase and polarization with additional computation steps and templating techniques, as discussed in Section 5.7. Before these extra dimensions, however, the author has had to validate new approaches to analyzing this dense 2D parameter space.

6.2 Quantifying directedness: sensitivity studies in real data

Gaussian noise studies, of which the Mock Data Challenge was one, led to the development of most of the code for detection and parameter estimation, as well as a set of upper limits (ULs). UL determination in the MDC was severely constrained by the sample size of 50 unblinded ‘open’ injections. While detection criteria and parameter estimation were sound, ULs demanded more simulated injections. With the MDC complete, resources were free to conduct these injections into several fre-

⁵Gil-Pelaez lets us solve R statistic as a sum of weighted χ^2 variables. For such sums, the joint probability distribution can be derived using characteristic functions [91] or generating functions [147].

quency bands of S6 data, at 142 Hz as well as 162 and 222 Hz (the latter two are consistent, detailed in Appendix B). A Feldman-Cousins approach [79] (to accommodate the possibility of a detection) then let us set one-sided confidence intervals: the upper limits.

6.2.1 Real S6 data: detection efficiency

Figure 6.1 plots the *detection efficiency* of TwoSpect in a 0.1 Hz test band starting at 142 Hz. Five hundred signals are injected into each LIGO interferometer (H1 and L1). For each injection, the astrophysical parameters were the same, but antenna pattern detector responses were applied per interferometer. A separate TwoSpect analysis was run for each injection simulation to avoid cross-contamination. The results of corresponding analyses were compared between IFOs. Templates meeting the detection criteria of the Mock Data Challenge (as discussed in Chapter V, having a single-template $\log_{10} p \leq -7.75$) were, if coincident in both interferometers by being in adjacent df , f pixels, counted as detections.

Every injected signal had some amplitude h_0 , but we expect the weakest injections to be swamped by detector noise. Detection efficiency curves illustrate when the signal becomes detectable, a certain fraction of the time, as injection amplitude is varied. Some signal parameters vary as well, such as polarization (via astrophysical $\cos \iota$ and ψ), frequency and $a \sin i$, and in-band differences in the noise realization make the same h_0 detectable or not, depending on these nuisance parameters. The detection efficiency curve is marginalized, that is, ignores these extra parameters to show only detection rate vs h_0 . TwoSpect begins to detect roughly 95% of injections when h_0 is about 10% of the strain amplitude spectral density (in $1/\sqrt{\text{Hz}}$) of the detector.

Relative detection efficiency of simulated signals vs h_0 at [142.0, +0.1] Hz
 lower of two max h_0 reported between H1, L1: $5.78485e-25$




Figure 6.1: Detection efficiency of 500 injections (each at H1, L1) into S6 data at 142 Hz, given threshold $\log_{10} p = -7.75$. The least-squares curve fit is to a symmetric sigmoid, which matches the data well at high detection efficiency. Since the 95% efficiency region is the most interesting, the relatively poor fit at low efficiency is not much concern, although the fit could be improved with additional parameters.




Figure 6.2: Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 142 Hz (injections also done at 162, 222 Hz)

6.2.2 Real S6 data: h_0 recovered vs injected

Using the same set of injections as for detection efficiency, one can also recover estimated h_0 values, as Figure 6.2 illustrates. TwoSpect returns an h_0 proportional to $R^{1/4}$, but the proportionality constant was found to be slightly off in the MDC from that used in all-sky searches, prompting an empirical rescaling by 1.11⁶. Moreover, studies of the $\cos \iota$ uncertainty had led us to correct h_0 estimates further when as-

⁶This scale factor appears the same across declinations, simply an update to the conversion factor from R to h_0 .

suming random polarization: a uniform distribution of $\cos \iota$ values near the detection threshold induced a 1.74 ± 0.37 factor, so we multiplied the S6 injections by 1.74. The real data injections were necessary to clarify how these h_0 levels corresponded to a given confidence level. In other words, the reported UL should be above the injected h_0 at least 95% of the time.

Figure 6.2 reports the fraction of the total injection set for which our final formula yields a UL greater than the injected h_0 , which is higher than 95% because the formula must hold locally over the entire range. When attention is restricted to $h_0 > 2 \times 10^{-24}$, we find that the confidence level multiplier must be 2.0. Assembled, a UL for each 0.1 Hz band is reported by the formula,

$$(6.2) \quad h_{0, \text{UL, band}} = 2.0 \times 1.74 \times 1.11 \times \sup(\{h_{0, \text{reported, band}}\}).$$

With coincident analysis of interferometers, the lowest UL reported for a band is taken as the joint UL. This criterion and Equation 6.2 supersede the single MDC UL with a frequency-dependent UL, calculated for each 0.1 Hz band between 40 and 2040 Hz.

See Appendix B for more details on injection studies.

6.3 Scorpius X-1 search using Directed TwoSpect in S6

6.3.1 S6: Scorpius X-1 search plan

MDC studies completed, the author has conducted a 2 kHz, $\pm 3\sigma_{a \sin i}$ search over all S6 data from H1 & L1. Stated again briefly, we analyzed the following:

- 40 to 360 Hz with 840-s SFTs,
- 360 to 2040 Hz with 360-s SFTs,

- Overlapping band verifying 260 to 360 Hz with 360-s SFTs,
- Overlapping band verifying 1400 to 2040 Hz with 300-s SFTs

As of this writing, all of these production searches are complete. The similarity in results between primary (840-s SFT) and overlapping (360-s SFT) bands between 260 and 360 Hz suggests that the high frequency overlap will also yield concordant results. Thus the author derives preliminary ULs for 40 to 360 Hz from the 840-s SFTs and 360 to 2040 Hz from the 360-s SFTs, as described below.

6.3.2 S6: Scorpius X-1 heatmaps

As the S6 search was run on the Atlas cluster at AEI Hannover, individual bands could be checked for consistency with known artifacts. Heatmaps plotting the intensity of R , $\log_{10} p$ and estimated h_0 (as reported, not corrected by Equation 6.2) were generated. Figure 6.3 shows one set of 50 bands, spanning 5 Hz, on H1. This band clearly shows the presence of the 46.7 Hz *calibration line*⁷. Seeing known spectral features at their expected frequencies confirms that TwoSpect is reading data accurately. The heatmaps can also be used to follow up on interesting outliers.

6.3.3 S6: Scorpius X-1 upper limits, random polarization

Most important are the upper limits (ULs). TwoSpect has been configured to analyze only ‘usable SFTs’, which must pass a Kuiper’s test. Non-Gaussian or otherwise extremely noisy data is not used. The 60 Hz and first three harmonic lines (120, 180, 240) Hz harmonic lines, as well as frequencies near the violin modes around 340 to 350 Hz, are excluded by these tests. Altogether, the 40 to 360 Hz H1 search excluded 16.4 Hz and L1 16.2 Hz, while the 360 to 2040 Hz H1 search excluded 21.4

⁷These lines help calibrate the h_0 amplitude spectral density [127], although their TwoSpect reported amplitudes do not match the intended h_0 because the lines are not modulated like neutron stars in binary systems.




Figure 6.3: S6 h_0 heatmap shows real data features, such as 46.7 Hz cal line




Figure 6.4: Joint 95% upper limits for Scorpius X-1 in S6 data, using a confidence interval given by reported $h_0 \times (1 + 1.0) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. This spectrum covers 40 to 2040 Hz using the lower upper limit from either interferometer (H1 or L1) when both yielded data. A total of 28.8 Hz were in bands that yielded no real upper limit (because the quarter root of the test statistic was imaginary) in either interferometer, generally due to excessive noise in that band. Bands were left-closed and right open, e.g., $[40.0, 40.1)$, $[40.1, 40.2) \dots [2039.9, 2040.0)$. These results are promising, given that the Advanced LIGO noise floor is expected to be up to 5 times lower at the most sensitive current frequencies of a few hundred Hz. Advanced LIGO will also make even larger relative improvements from 10 to 40 Hz, where the torque-balance equation allows higher GW strain than at high frequencies. Together with recent and anticipated improvements in TwoSpect analyses, order-of-magnitude better upper limits should be anticipated for full-sensitivity Advanced LIGO results.

Hz and L1 16.9 Hz of search bands. For generating upper limits, only 28.8 Hz out of 2000.0 Hz could not be determined from either interferometer.

A 95% confidence level requires (Equation 6.2) a correction factor: $h_0 (1 + 1.0) \dots \times [\cos \iota \text{ factor}]$. The entire spectrum, with the final 95% confidence UL for random polarization, is shown in Figure 6.4.

Previous Radiometer searches using S5 data [12] may appear to show a better, lower UL, but the results are not directly comparable. The Radiometer UL is calculated for 90% and is for circular (optimal) rather than random, polarization. In addition, signal leakage across the 0.25 Hz bins used in the Radiometer search leads to UL degradation by as much as 70% at a frequency of 1500 Hz⁸. The complete transformation from Radiometer h_{RMS} to an upper limit is believed to be $h_{\text{UL}} = 2.43 \times 1.2 \times 1.7 \times h_{\text{RMS}}$ at 1500 Hz, based on current understanding of the differences in algorithm. TwoSpect's S6 ULs are then an improvement. Moreover, given an outlier, TwoSpect could do something not yet possible with other comparable pipelines: measure the projected semimajor axis, $a \sin i$.

6.3.4 S6: Scorpius X-1 outliers

Table 6.1 presents a list of outliers present in both interferometers between 40 and 360 Hz. Tables C.1 to C.3 in Appendix C present a list of outliers present in both interferometers between 360 and 2040 Hz. Follow-up of 40 to 2040 Hz outliers is beginning, but the search is not yet complete. Some coincident outliers present in the initial analysis include lines at powers of 2 and a probable trace of a hardware-injected simulated pulsar. As noted earlier, power line frequencies and violin modes were automatically dismissed by TwoSpect, using a Kuiper's Test on SFTs. More exhaustive study of the candidates in Tables 6.1 and C.1 through C.3

⁸Confirmed in private communication and will be noted in the forthcoming Bulten *et al* Sco X-1 paper [52].

Outlier Number	Frequency (Hz)	Explanation
1	42.00	2 Hz comb
2	64.00	Power of 2 line
3	108.10	–
4	108.85	Injected pulsar 3
5	109.50	–
6	111.02	–
7	128.00	Power of 2 line
8	139.52	–
9	154.04	2 Hz comb
10	156.82	–
11	157.99	2 Hz comb
12	158.36	–
13	158.87	–
14	190.86	–
15	192.54	Injected pulsar 8
16	200.53	–
17	200.60	–
18	209.21	–
19	209.28	–
20	223.66	–
21	256.02	Power of 2 line
22	268.13	–

Table 6.1: List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 40 to 360 Hz.

will continue, although GW detection would be a rather optimistic prospect given the torque-balance limit for Scorpius X-1 [47] and estimated detection efficiency.

GW emission is informed by torque-balance, but enough uncertainties exist to justify careful searches, particular when Advanced LIGO data will become available: detection may be plausible. These uncertainties are challenging to quantify. The significance of the torque-balance limit is that accretion should drive the neutron star to spin at the frequency of GW emission. Any slower, and accretion would spin it up faster. Any faster, and GW emission would sap its energy. Yet the implicit assumption is that rotational energy is not being depleted by friction or other means. X-ray flux must also be accurately estimated, which for other systems may be complicated by X-ray beams, collimated for instance by the accretion disk – indeed, all-sky searches [7] assume many LMXBs are unseen. Moreover, the frequency of

Scorpius X-1 is unknown, and this means that the torque-balance limit ranges from a strain of 1.1×10^{-25} at 40 Hz to 1.5×10^{-26} at 2040 Hz. The overall uncertainty on the torque-balance limit is thus at least this large.

Two further notes: we have assumed that the orbital tilt of Scorpius X-1 leads to a projected semi-major axis of $a \sin i = 1.44 \pm 0.18$ light-seconds. This is manifest in the search range of frequency modulations, as greater $a \sin i$ leads to greater relative modulation depth. Ephemeris data for $a \sin i$ may improve in the future, which would ease future investigations. Additionally, ephemeris data for $\cos \iota$, the neutron star inclination angle, is becoming available. Knowing $\cos \iota$ would make our upper limits more meaningful, and possibly more constraining. If a GW detection is eventually made, $\cos \iota$ would connect the observed GW strain to the actual GW strain emitted. This would be crucial in connecting observations of Scorpius X-1 and similar systems to astrophysics. When actual GW strain emitted is known, the torque-balance equation can be tested, and we can probe the process of accretion in LMXBs and beginning probing the surface, and perhaps interior physics [117], or neutron stars.

The search for GWs from Scorpius X-1 continues, and the imminent observational runs of aLIGO give hope that its signature may soon be detected. Although the S6 upper limit is about 1.3×10^{-24} , we saw in Chapter V that at the minimum of the Advanced LIGO noise floor, the expected upper limit is 4.23×10^{-25} . This is within a factor of 4 of the 40 Hz torque-balance estimate, which itself has comparable uncertainty, and we expect TwoSpect improvements to be of order of a few. With continued improvements to CW algorithms and the enhanced sensitivity of future interferometers, we expect that Giacconi's discovery [89] of the first and brightest extrasolar X-ray source will be paralleled with a GW discovery.

6.4 XTE J1751-305 search using Directed TwoSpect in S6

6.4.1 S6: XTE J1751-305 background

Discovered by Markwardt *et al* in 2002 [120], the X-ray transient (XTE) J1751-305 has indications it too could emit GWs [158]. As detailed by Strohmayer and Mahmoodifar, J1751 has been searched for signs of non-radial oscillation modes, such as r - (and g -) modes. Preliminary results suggested that r -modes might be consistent with X-ray observations. Subsequently, Andersson *et al* [38] argued that an r -mode would have already spun down below detectable levels, but Lee [114] suggested that a crust-only surface r -mode might still be present. J1751 is an intriguing candidate in its own right and highlights the different kinds of directed GW search that TwoSpect can conduct.

J1751 is the XTE with the shortest known period, $P \approx 2545.3$ seconds or 42 minutes, and $a \sin i \approx 0.010 \pm 0.003$ s. Crucially, its spin frequency, unlike that of Scorpius X-1, is known with microHertz precision: $f = 435.31799$ Hz. A search can thus be fast ($< 10^5$ templates). At an estimated $d > 7$ kpc, near the galactic center, J1751 is distant for a CW source, but the chance of a r -mode motivates a quick look. We expect GW emission to be most likely at the following frequencies:

- ν_{spin} : 435.31799 Hz
- r -mode: $(2-0.5727597) \cdot (435.31799 \text{ Hz}) = 621.3034 \text{ Hz}$
- $2\nu_{\text{spin}}$: 870.63598 Hz

6.4.2 S6: XTE J1751-305 heatmaps

A straightforward search over $+14, -3\sigma_{a \sin i}$ has been conducted, having searched frequency bands of $[434.5, 436.5]$ & $[620.5, 622.5]$ Hz, & $[869.5, 871.5]$ Hz.




Figure 6.5: Quick look at J1751-305, H1 $\log_{10} p$, 435 Hz ν_0 rotation frequency.




Figure 6.6: Quick look at J1751-305, L1 $\log_{10} p$, 435 Hz ν_0 rotation frequency.




Figure 6.7: Quick look at J1751-305, H1 $\log_{10} p$, 621 Hz r -mode.




Figure 6.8: Quick look at J1751-305, L1 log₁₀p, 621 Hz *r*-mode.




Figure 6.9: Quick look at J1751-305, H1 $\log_{10} p$, 870 Hz $2\nu_0$, twice rotation frequency.




Figure 6.10: Quick look at J1751-305, L1 $\log_{10} p$, 870 Hz $2\nu_0$, twice rotation frequency. The dominant feature is believed to be an instrumental line.

A TwoSpect analysis and heatmaps have been made for the ν_{spin} , r -mode, and the $2\nu_{\text{spin}}$ frequency bands. Upper limits and candidate comparison can then be made. Even if no GW signal is seen – the probable outcome absent r -modes – XTE J1751-305 offers a prototypical search strategy for well-characterized, directed CW sources.

6.5 Summary of Directed TwoSpect S6 searches

The author has conducted the initial, primary stage of a 2 kHz, $\pm 3\sigma_{a \sin i}$ directed search for Scorpius X-1 over all S6 data from H1 & L1 using TwoSpect. The analysis from 40 to 360 Hz used 840-s SFTs, and from 360 to 2040 Hz used 360-s SFTs, (to prevent spectral leakage). This primary analysis (4 kHz of data between the two detectors) has yielded preliminary joint upper limits over the 2 kHz from 40 to 2040 Hz. Overlapping bands are being used to confirm the results using 360 s SFTs to validate the 260 to 360 Hz band and 300 s SFTs to validate 1400 to 2040 Hz, but initial output justifies treating the existing results as valid.

Production analysis for this S6 search was 100% done in roughly 30 days. In the process, our post-processing methodology has been adjusted from all-sky techniques [7] to one appropriate for a densely-sampled, correlated ‘ df vs f ’ parameter space. The remaining questions are in the identification of outliers – already ~ 22 coincident features have been noted between 40 and 360 Hz – and explanation of too quiet bands. Too quiet bands, seen in Figure 6.4, appear due to the noise subtraction in the R statistic; further investigation will see whether these presumably non-genuine ‘good’ bands can be systematically rectified.

In addition, a search for XTE J1751-305 has been completed, with its own post-processing underway.

Scorpius X-1 and XTE J1751-305 are among the best potential sources of GWs in the sky. This chapter shows that we can extract an upper limit around 1.3×10^{-24} from Scorpius X-1 in S6 data and demonstrates that these searches will be practicable in the future. Advanced LIGO will turn on for O1 in summer 2015: TwoSpect, one of the prime algorithms for finding continuous GWs from neutron stars in binary systems, is ready.

As previously noted, results from this chapter are preliminary and have not been reviewed yet by the LIGO Scientific Collaboration.

CHAPTER VII

Outreach: Exhibit at the World Science Festival

Science is understanding the world – *θεωρία*; art is affecting the world – *πραξις*. Outreach for scientists is a matter not just of connecting with colleagues in the academy, as C.P. Snow [157] said, bridging the Two Cultures. It is reaching back to our innate humanity and the world from which it came, the universe trying to understand itself. No single exhibit reaches this goal, though we do our best with rubber sheets, lasers, and interactive light and sound sculptures. Diligence nonetheless can spark curiosity. LIGO has made concerted efforts to reach further into the universe with ever more sensitive technology, and in leadup to the World Science Festival (WSF) exhibit on LIGO in 2010 in Manhattan, we tenaciously developed more sophisticated exhibits to connect with a broader audience.

Many kinds of exhibits have been made to showcase aspects of general relativity. Rubber sheets of ‘spacetime’ are familiar to many museum-goers; the WSF had one too, with a central mass to deform the sheet like gravity and a smaller marble that could ‘orbit’ it. More central to the WSF exhibit about LIGO, however, was an interferometer. Requiring a strong laser, stable alignment, and asymmetric-length arms to great a distinct bullseye fringe pattern, this interferometer’s design, installation and commissioning was the responsibility of the author. As a microcosm of

the challenges mentioned in the introduction, the WSF offers humbling insight into both gravitational wave astronomy and its significance to the public.

7.1 Prototypes: travelling kiosks and the Ann Arbor Hands-On Museum

The Michigan Gravitational Wave Group (MGWG) came to the WSF project with experience. Ramon Armen had developed a small interferometer with Goetz and Riles, approximately 50 cm square, that integrated into a kiosk design. This kiosk was developed in collaboration with the Ann Arbor Hands-On Museum exhibits director John Bowditch. Charlie Stout wrote an interactive Flash program for the kiosk.

In the small interferometer, a 0.9 mW Helium-Neon laser beam passes through a beam splitter, travels down two arms of slightly different length, and reflects off retroreflecting mirrors at the end of each arm. The retroreflectors contain three mutually perpendicular planes that redirect the beam parallel and opposite to its incident path, offset by twice the distance from the retroreflector corner. Retroreflectors are more robust against misalignment than flat mirrors, although reflected images cast a six-spoke shadow pattern due to the three intersecting lines between the planes. Slight beam offset away from the center can avoid these spokes. These optical elements are mounted on a portable breadboard with pre-fabricated optical table-style holes. Recombined at the beam splitter, some light goes to a beam dump while the remainder is directed to a screen. As the path length of the beams and consequent phase delay changes – due to thermal expansion of the breadboard, vibrations, and so on – the screen shows a fluctuating fringe pattern that illustrates the constructive and destructive interference of light waves. Mismatched wavefront curvatures, arising from travel along the different-length arms, give this fringe pat-

tern a well-defined circular bulls-eye shape. Speakers, when used, let exhibit-goers hear how vibrations, which they themselves can make by tapping on the kiosk, reverberate through the mirrors and be both seen on-screen and heard as fringes pass the photodiode. In the kiosk, a swinging lever arm, controlled by whomever stands at the kiosk, can block the light from one of the arms, blocking the interference; the user can see that the fringe pattern indeed comes from the wave nature of light.

A full set of parts for the kiosk interferometer was assembled by the author (Table 7.1) according to Figure 7.1 in the process of building another unit. This additional interferometer went on travelling exhibition, organized by Marco Cavaglia of the University of Mississippi, across the United States starting with the 2009 WSF, and it was most recently on display at Stanford University in August 2014. It is durable, but for the LIGO ‘Astronomy’s New Messengers’ 2010 WSF exhibit, the MGWG crafted something more dramatic.

Bill of materials: interferometer parts

Equipment as below builds one interferometer for our Ann Arbor Hands-On Museum outreach design. We include company names, part numbers, quantities, and prices, after the work done in 2008 by Ramon Armen. Assembly follows from the work detailed on

<http://gallatin.physics.lsa.umich.edu/~keithr/outreach/>,

in particular the parts list (plain text file) written by the author,

http://gallatin.physics.lsa.umich.edu/~gmeadors/Interferometer_parts,

and the schematic

http://gallatin.physics.lsa.umich.edu/~keithr/outreach/ifo_schematic.

jpg.

See Table 7.1 for details on parts from each of the manufacturers: CVI Melles-

Description	Part Number ¹	Quantity	Approximate Unit Price
HeNe laser	05 LLR 811-249	1	\$320
Laser stand	07 LHE 001	1	\$40
Description	Part Number ²	Quantity	Approximate Unit Price
-50 lens	NT32-996	1	\$27
-25 lens	NT32-992	1	\$27
Lens holder	NT54-980	2	\$33.00
Description	Part Number ³	Quantity	Approximate Unit Price
Beam splitter	038-0590	1	\$250
Retroreflector	055-2340	2	\$289.00
"" holder	113-0045	2	\$69.00
Description	Part Number ⁴	Quantity	Approximate Unit Price
BS holder	LMR2	1	\$23.50
Rotary mount	RP01	1	\$91.00
1.0" holder	PH1-ST	1	\$7.03
1.5" holder	PH1.5-ST	4	\$7.22
0.75" post	TR075	1	\$4.74
1.0" post	TR1	4	\$4.74
Screw base 1	BA1S	6	\$5.13
Screw base 2	BA1	1	\$5.56
Screw type	Length	Quantity	
1/4"-20 capped	5/4"	8	
1/4"-20 capped	3/4"	7	
1/4"-20 set	3/8"	5	
#8-32 set	1/2"	1	

Table 7.1: Bill of materials (lasers, lenses, mirrors, optic mounts and screws) for interferometer assembly. Manufacturers: ¹ CVI Melles-Griot, ² Edmund Optics, ³ OptoSigma, ⁴ ThorLabs

Griot¹, Edmund Optics², OptoSigma³, and ThorLabs⁴.

Figure 7.1 shows the base structural dimensions of this kiosk interferometer.

7.2 World Science Festival interferometer manufacture

Based on the kiosk interferometer, we collaborated with Cavaglia on a larger interferometer for the WSF. Displayed first in June 2010 in New York City, the WSF interferometer incorporates new features. The projection screen was enlarged, made free-standing, and incorporates an integrated audio-frequency photodiode connected to a new speaker system. Laser power is greater, to be seen better on the bigger

¹CVI Melles-Griot: +1 505 296 9541; press 2 at menu

²Edmund Optics: +1 800 363 1992; press 1 (?) for sales

³OptoSigma: +1 949 851 5881

⁴ThorLabs: +1 973 579 7227; press 1 (?) for sales




Figure 7.1: Kiosk interferometer schematic



Figure 7.2: Interferometer fringes during construction in Ann Arbor

screen, and laser light has been carefully contained with the use of multiple baffles. All are contained inside a plexiglass ‘beam tube’ fastened with security screws. This interferometer stood at the center of the WSF exhibition on LIGO, which received around 2200 visitors in four days.

7.2.1 Laser, optics and display

Starting from the request for a larger interferometer, the MGWG found a brighter laser with longer coherence length. Laser Lab Components, Incorporated (LLCI) manufactured a 60 mW Nd-YAG 532 nm laser rated to a coherence length of 50 meters, plenty for the longer path length. The laser intensity was so bright that we installed a times-8 attenuator to provide just 7.5 mW of visible light. As the

wavelength is green, it is even more visible to the human eye than the red HeNe of the kiosk laser. To protect visitors' eyes, the author designed a system of eight black-painted, wooden baffles to intercept all ghost beams and secondary reflections (both specular and diffuse).

These baffles were manufactured, along with the free-standing display screen, in Ann Arbor by Fingerle Lumber. Additional black, corrugated cardboard baffles were made by the author to shield the multiple reflections immediately around the laser and the retroreflectors. This display screen was hollow in order to accommodate a photodiode, BNC connection, and cable for the aforementioned speaker system. Installation of the photodiode assembly and hookup to a speaker system was performed by the MGWG.

Optical installation and alignment was done by the author. Retroreflector and beam-splitter follow the kiosk model except in size, using a 3-inch diameter beam-splitter to avoid beam clipping. The large interferometer adds a beam-focusing lens between the laser and beam-splitter, which serves with the beam-splitter to steer the beams onto the retroreflectors. In order to avoid clipping the beam-splitter and to avoid the retroreflector spoke shadows, the retroreflector mounts allow for one-dimensional sliding on the shorter X -arm and two-dimensional sliding on the longer Y -arm. These adjustments are typically fine-tuned after an initial image is formed on the output port of the beam-splitter.

A trio of first steering mirror, beam-expanding lens, and second steering mirror project the beam from the beam-splitter output onto the display screen. The steering mirrors are both flat; the first is at a ninety degree angle, in-plane with the other optics, and the second is also ninety degrees in the main optical plane but tilted upwards to shine toward the screen. In between, the beam-expanding lens is chosen

to maximize the width of the image on-screen, constrained by the clipping of the second steering mirror, Unfortunately, the beam path after the second steering mirror is too high for a conventional lens mount to insert the lens there. In practice, the arrangement highlights the fringe pattern.

While the fringe pattern fluctuates considerably, with an amplitude sometimes of several wavelengths, this is to be expected in a meter-scale instrument with path-length in air and without an insulation from ground motion.

Initial estimation of optics locations was done in Matlab with a Gaussian beam, ABCD matrix model, followed by physical fine-tuning, both done by the author. The BOM of large interferometer core optics (superceding Table 7.1's beamsplitter and retroreflectors) includes the following:

- (1 unit) OptoSigma, 112-2855 (50 mm gimbal, 360 degree rotation): \$305
- (2 units) OptoSigma, 055-2370 (50 mm retroreflector, 425-675 nm): \$315 each
- (2 units) Edmunds Optics, NT63-114 (50 mm mirror, 532 nm): \$220 each
- (1 unit) Newport Optics, Fused Silica Broadband Dielectric Beamsplitter (76.2 Dia, 12.45mm, $\lambda/20$, 480-700nm): \$1190.00
- (8 units) Thorlabs, 50 mm lens holders

7.2.2 Aluminum baseboard

All optics are mounted on a purpose-made aluminum baseboard. Prototyping was conducted on an optical table, which would be too large to ship conveniently. No breadboards were known of adequate size, and typical honeycomb construction would also likely be heavy. Alro Steel of Jackson, Michigan manufactured the baseboard used in the exhibit. Although the author tried using neoprene rubber sheets as a

x-position (in)	y-position (in)
10	80
04	18
04	12
04	10
08	10
12	12
12	10
10	12
10	10
10	08
10	06
12	04
10	02
44	10
46	10

Table 7.2: Hole locations (in inches from origin) for the WSF interferometer aluminum baseboard, plotted with suggested alterations on Figure 7.5.

form of seismic isolation for the baseboard, this was found to offer no improvement in the fringes.

Table 7.2 and Figure 7.5 reference the positions of holes in the aluminum baseboard that must be tapped. These holes permit the attachment of plexiglass blocks with perpendicular screw tappings that in turn allow the attachment of the plexiglass enclosure described in Section 7.2.3.

7.2.3 Plexiglass enclosure

To protect the interferometer and visitors from each other, the laser and optics were included on a plexiglass enclosure, also mounted to the aluminum baseboard. Figure 7.3 presents an AutoCAD view of the enclosure as a whole, with dimensions shown in Figure 7.4. The AutoCAD design file can be found online:

`http://gallatin.physics.lsa.umich.edu/~keithr/outreach/meadors_wsf_draft_08-unified.dwg`

To manufacture the enclosure, the author visited and contracted with Plex Lab




Figure 7.3: AutoCAD corner view of plexiglass and aluminum interferometer enclosure. Plexiglass in light gray, aluminum in dark gray. Plexiglass is attached to aluminum with small plexiglass blocks (a refinement to the design), placed three to the long *Y*-arm, two to the short *X*-arm, and three to the corner station. The triangular prism is not anchored to the aluminum block. It is open to the bottom so that a projection screen can be inserted, typically atop a flexible rubber sheet. In the World Science Festival exhibit, this projection screen included a photodiode and attached signal cable.




Figure 7.4: Front view of interferometer aluminum and plexiglass enclosure, with dimensions.




Figure 7.5: Dimensions and hole locations of the aluminum plate (dark gray) relative to plexiglass (light gray), with proposed hole locations given by Table 7.2. These hole locations were as tapped.

Corporation (formerly of Warren, Michigan). The enclosure was fabricated in pieces: the arms were made from joined half-tube ceilings together with straight walls. A laser ‘end station’ and triangular prism enclosure for the projection screen were joined to it. Joints endured significant stresses during transport and the arms became detached; in the future it would be better to fabricate the enclosure with detachable arms from the start. Plexiglass enclosure pieces attached to the baseboard using small rectangular plexiglass blocks. Blocks were fastened to the tapped baseboard with screws, then the enclosure lowered overhead (requiring at least two people to handle the prism section). Through-holes through the enclosure allowed security screws to clamp into tapped, perpendicular holes in the blocks from the outside.

A unified view of the plexiglass and aluminum, with aluminum screw-holes indicated for the optics, is seen in Figure 7.5, and a view of the finished product (minus the projection screen) in Figure 7.6.



Figure 7.6: Interferometer assembly in Ann Arbor, Michigan. Projection screen not yet installed. Note the outer end of the Y-arm (bottom left of photo), with perforated-and-covered plexiglass to allow sound passage. Humans (*e.g.*, Evan Goetz at left, Keith Riles at right) can be protected during alignment processes by laser safety curtains, seen in back.




Figure 7.7: World Science Festival interferometer installed in the New York City exhibition hall, June 2010. Optics aligned and baffles installed to protect visitors from scattered light. Also visible: left center, an initial LIGO input mirror and suspension.

7.3 Exhibitions: New York City, Portsmouth, Fort Wayne

Built in Michigan, our interferometer was transported by U-Haul truck (driven by Riles) for exhibition at the WSF.

According to Equation 2.36, the interferometer, if limited only by shot noise, could transduce GW strain to power P :

$$(7.1) \quad P = \left[4(1\text{m}) \frac{2\pi c}{532\text{nm}} \frac{(60/8)\text{mW}}{c} \sin \left(\frac{2\pi}{532\text{nm}} (L_s + L_0) \right) \right] h_+.$$

If operated where the sine term equalled one (not done in LIGO for reasons noted in Chapter II), then this would give yield 3.5×10^8 Watts per unit strain, or 7

attoWatts for a gravitational wave at the initial LIGO noise floor, $h_0 = 2 \times 10^{-23}$. Though we never measured the photodiode dark current, it seems dubious that such a signal would be detectable. In the audio band the interferometer was visibly noisy, as the plexiglass enclosure and aluminum baseboard afforded little acoustic isolation. Though no measurements were taken, it is conceivable that it was quieter at several kiloHertz, where the shot noise argument would be valid. Were isolation better, the full 60 mW laser power used, and photodiodes calibrated, it would not be so strained to compare the WSF larger interferometer to Robert Forward's [82], which had 35 to 55 mW single-mode laser power and a 8.5 m folded arm length. Scaling down from Forward, we could expect $10^{-15}(\text{strain})/\sqrt{\text{Hz}}$.

While rather worse scientifically than LIGO, our WSF large interferometer at least could travel to the masses. The exhibit has seen four locations and several thousand visitors, including 2200 in the June 2010 WSF itself. After its debut in New York City, it visited Portsmouth, Ohio in April 2011 and Fort Wayne, Indiana in September 2011; it now resides in the LIGO Livingston Science Center. Figure 7.7 gives the scale of the interferometer in the whole exhibition.

7.3.1 Exhibit overview

One part of a larger exhibit, the interferometer sat at the center. Next to it, on the same table, sat an actual small mirror from one of the LIGO interferometers. Surrounding it was a curved fabric wall, nestled into which was a complete kiosk interferometer, on the left. Two video screens played *Einstein's New Messengers*, a documentary about LIGO, on the walls, along with a third, silent video screen outside showing a shorter feature. Also snug in the exhibit, on the right, was a rubber sheet of deformable spacetime. At back, two computers were set up for playing *Black Hole Hunter*, a program developed at Cardiff for participants to try to hear a gravitational

wave inspiral buried in noise. In the center, another computer, directly behind the large interferometer, linked to a vast light sculpture (built by Leni Schwendinger); the computer could illuminate the sculpture's LEDs according to the sounds of an inspiral, a burst, the stochastic background, a continuous wave (audio simulations generated by Riles), or the user's own voice⁵. Solid-backed posters hung from several points along the fabric wall. Initial setup in New York required about three days, although the Fort Wayne installation, without the light sculpture, was complete in one. The author re-commissioned the interferometer, and, at the WSF, explained the exhibit to visitors.

7.3.2 World Science Festival 2010

Our exhibition, situated in the Broad Street Ballroom in Lower Manhattan, received a spectrum of visitors. All ages, from preschoolers to grandparents, found something with which to interact. Many expressed curiosity about the meaning of gravitational wave astronomy, how soon GWs will be seen, and what the implications mean for humans and the universe. The exhibit was hosted with the help of Cavaglia and Brooke Rankins from Mississippi, as well as Szabi & Zsusa Marka and their students from Columbia University, including Max Factourovich and Maggie Tse, along with Michael Hartman from the University of Florida. Figure 7.8 shows the author with our interferometer.

7.3.3 Portsmouth and Fort Wayne

Following the success of the WSF, the entire exhibition went into temporary storage. Ten months later, it was erected again at Shawnee State University by Cavaglia, the author, and Shawnee's Tim Hamilton. There it stood in a student lounge. Five months after that, Science Central, a museum in Fort Wayne, showcased

⁵As demonstrated by an opera singer in attendance, although all participants made beautiful patterns.



Figure 7.8: Helping to host the exhibit at the World Science Festival. Optics aligned, with fringe pattern, but an edge of the beam-splitter reflection is visible in the projected image. Additional features: the reflected blue lights of the sculpture above and posters of the exhibit.

the exhibition (save for the light sculpture) in their travelling exhibit hall. Although the plexiglass is slightly worse for wear, the interferometer endures, just requiring re-alignment at new locations.

7.4 Future LIGO outreach

Epilogue: as of June 2014, the exhibition is safely on display in the Livingston Science and Education Center at LIGO Livingston Observatory. Cavaglia reports that the interferometer still works as well as in New York, with some fixes to the plexiglass. LIGO has not yet been cast into the public's vision. Dedicated individuals have created exhibits, both travelling and stationary, tangible and digital. Most recently, data from the fifth science run, S5, has been released to the public. Though the author contributed directly to these efforts only by building interferometers and giving talks, we can expect more to be asked of us when Advanced LIGO and its peers make their first discoveries. We must convey the meaning of a new astronomy.

CHAPTER VIII

Conclusion

8.1 Cycles of science

Science connects us to the world. From *scientia* to *Wissenschaft*, science has ever connoted both sense and wisdom. It is how humanity takes in from, and the arts are how humanity gives back to, nature. Yet the neat duality of modern identity is no more mysterious than those in modern physics. Experimental physics unites insight and craft: it makes tools of discovery. As physicists, our aspiration to know the cosmos is realized if and only if that knowledge is both true and understood.

Gravitational waves will give new information about the universe. For the present, we do strive for gravitational wave astronomy to be understood, as in Chapter VII. Above all, we must know whether or not it is true. The author's projects detailed in Chapters II through VI explain how we may yet learn what is. This thesis covers both instrumental and analytical improvements to the experimental & observational science of gravitational wave astrophysics.

8.1.1 Improvements to observatories

Gravitational wave interferometry is a young science. The astrophysical potential is even more nascent. First, we must have a view. Techniques for interferometry, as in LIGO, are continually evolving: this thesis details my engagement with efforts to

make the interferometers more stable and more sensitive, including in going beyond the standard quantum limit imposed by vacuum fluctuations. These methods will be key to improving on Advanced LIGO.

Chapter II details how the LIGO Scientific Collaboration is reaching out to our colleagues in the astronomical community to unite our view of the universe. The author also discusses contributions made to characterizing the LIGO Hanford Observatory, building a phase camera for use in future interferometers, and studying detector glitches. Chapter IV shows that more fundamental improvements are possible. By way of quantum optical squeezing, the standard quantum limits of the electromagnetic field can be surpassed. Using squeezing *in lieu* of laser power, the observatory reached unprecedented sensitivity.

8.1.2 Understanding instruments, refining data

LIGO interferometers are complicated instruments with many interacting parts. The study of these parts is known as Detector Characterization. Characterizing the detector tells us the difference between earthbound glitches and spaceborn signals. Detector characterization has many enhancements to give to LIGO. My own contribution has been in translating characterization into correction.

Chapter III explains how the couplings between the auxiliary length controls and the gravitational wave strain introduce unwanted noise – and how it can be subtracted. Great care is taken to ensure that these fixes do not add any noise or unphysical signal. These methods can both be applied *post facto* and, without the surrounding machinery, in real-time. The author’s work began at the end of the last science run, S6, and similar techniques may make such challenges a smaller obstacle in the future.

Advanced observatories are becoming ever more complex: as Section 8.2.1 will

summarize, we need means to see and disentangle their ties. This thesis shows how automated techniques can loosen some knots in our servo loops.

8.1.3 Searching deep-space

Data analysis itself must be sophisticated. Even the quietest, most sensitive gravitational wave interferometer is uninteresting if we cannot understand it. Here we present the interpretation of S6 in the context of low-mass X-ray binaries (LMXBs).

There are four key ways to listen to the gravitational wave sky: inspirals, bursts, stochastic, and continuous wave. Our search in Chapters V and VI has been for the last: continuous waves from neutron stars in binary systems. TwoSpect is already a capable search for unknown systems across the entire sky; the author modified the search to focus on Scorpius X-1 and XTE J1751-305, having honed methods on simulated data in collaboration (and friendly competition) with fellow searchers.

Preliminary TwoSpect results from Sco X-1 and J1751, as Section 8.2.2 will conclude, will soon be presented and appear to be best so far for frequencies above 500 Hz. Given Advanced LIGO data, these techniques should take gravitational wave analyses to the verge of the torque-balance limit and start to tell us about the astrophysics of LMXBs.

8.1.4 Reaching out, looking up

The sky excites the mind. Just as astronomy with light and X-rays, radio waves and neutrinos has caught the popular imagination, it is the hope of projects like those in Chapter VII to inspire others to ponder gravitational wave astronomy. While increasingly confident in the sources we see, the adventure lies in the unknown. The author has shown a simple and effective tool, a model interferometer, for showing a new kind of antenna, a new sort of observatory, to search out that unknown. Over

two thousand people saw the model interferometer on exhibit at the World Science Festival. As interest in gravitational waves grows, we hope to communicate our findings even more broadly.

8.2 Scientific merit: filtering and analysis

8.2.1 Feedforward improvement to LIGO data

Quantitative enhancements to LIGO are hard. The LIGO Scientific Collaboration is composed of over a thousand scientists, and most contributions are indirect. Feedforward subtraction of auxiliary channel noise is a rare case where an individual can directly improve the scientific power of a gravitational wave interferometer operating at full sensitivity after the data taking. This algorithm enhanced inspiral range by about four percent in S6 (4.14% H1, 3.60% L1), potentially allowing 12% more inspiral events to be detected.

More exciting still is the community forming behind a family of related techniques. Automated subtraction of a large class of noise sources – not only length control but glitches, gravitational gradients, and Schumann resonances of the Earth’s magnetic field – could be cancelled, according to work underway by fellow LIGO researchers. One of the author’s chief interests in this project was in proving that such a technique can be both effective and safe for the data. We believe that this has been shown.

Although the 12% increase in inspiral events detectable in S6 will soon be eclipsed by Advanced LIGO, Advanced LIGO in turn is built of the success of techniques akin to the feedforward this chapter proves valid. We hope that commissioning silences most noise sources, but we are glad to offer this filtering method for consideration in case it is needed.

8.2.2 TwoSpect directed search for neutron stars in binary systems

Neutron stars are the densest laboratories in the universe. This thesis does not detect gravitational waves from neutron stars, but it does help validate a method for finding them as they emanate from LMXBs.

Detecting gravitational waves is not easy, or else it would have been done long ago. Nonetheless, our algorithms, like our observatories, are better than they have ever been. The TwoSpect search has evolved into a form suitable for searching Scorpius X-1 and XTE J1751-305. On Scorpius X-1, we set a 95% confidence upper limit (UL) of GW strain for random polarization GW from 40 Hz to 2040 Hz. This UL appears to be the best (lowest) so far, particularly for shot noise-limited frequencies. The UL constrains possible GW emission to be lower than 1.3×10^{-24} in the most sensitive LIGO bands.

TwoSpect is already prepared for better data to arrive when Advanced LIGO begins observing runs in 2015. Indeed, enhancements are possible between now and then that may make it yet more capable. Neutron stars, the heaviest compact objects entirely in our universe – black holes in some sense having left the universe behind – remain opaque to our scrutiny for now. Our work complements X-ray and radio observations. Gravitational waves will unveil some of the secrets of neutron stars.

With TwoSpect's sensitivity demonstrated in this thesis, we think that seeing gravitational waves from LMXBs is a serious possibility in the Advanced LIGO era.

8.3 Entering the advanced detector era

Second-generation interferometers, Advanced LIGO and its peers, are almost here. Both LIGO Hanford and LIGO Livingston have completed installing all their components, have sealed their vacuum, and are actively commissioning. Advanced Virgo



Figure 8.1: LIGO Hanford Observatory sunset, initial detector era. Photo by author. Like the Hanford desert wiped clean by the Missoula floods, the gravitational wave sky may relate cosmic tales of cataclysm and rebirth in the distant past.

is also underway. Tunnels for KAGRA have been completed. Gravitational waves may not yet be seen for several years, and perhaps only faintly at first. As a fellow scientist¹ noted, Kepler saw two supernovae in the Milky Way in his lifetime; there have been none since. Perhaps we will be profoundly unlucky. Yet we have planned with circumspection and care for what we think we can expect: if two neutron stars merge within 200 Megaparsecs of Earth, then our observatories, by decade's end, should hear that inspiral. The bursts of supernovae and bumps on neutron stars will be sensible too, and, in time, the Big Bang. Even if we fail to see these (it would require much inspection and introspection before we are certain there is nothing to be seen), something will be learnt. Should more be there than we expect, then all our curiosity will be justified.

8.4 Vision of a dark sky

Why gravitational wave astronomy at all? What could be out there?

At the dawn of the twenty-first century, the Standard Model of physics, and quantum field theory in general, can be studied on curved spacetime, but the spacetime

¹Brian O'Reilly of Livingston.

itself remains scarcely better understood than when Einstein first proposed it general relativity. After almost a century, the fabric of the universe still keeps its secrets. Gravitational wave astronomy will perceive ripples in that fabric.

In this thesis, we have made inroads to this new astronomy. Superficially separate, the common thread is the pursuit of fundamental issues by skillful choice of perspective. With squeezing, the optical experimenter views a gravitational wave interferometer as a quantum system and sees how fluctuations in the vacuum, not just in the laser, create noise – which can be cancelled with a beam of no-light. With feedforward, the noise due to intrinsic couplings between the interferometer servos is found by coherence in the frequency domain – and cancelled with subtraction that can take place either in real-time or long afterward. With TwoSpect, signals buried beneath noise are uncovered by comparing multiple instruments – cancelling noise, in effect, with the build-up of signal in other observatories. Communicating these advances to a wider world is the final question of fundamental issues and choice of perspective. Gravitational wave observations have not yet seen a signal, yet we find ways to make our research meaningful.

Until we can understand gravity, we will be ignorant to the range of forces present in the cosmos. LIGO is a way to hear the echoes of gravitation. For long-lasting signals, we can even ‘see’ them, just as our ears can triangulate sounds. From the windswept, tumbleweed-coated plains of Hanford and the pine forests of Livingston may emerge our first visions of this thus-far dark sky. From it may come unexpected sources. Even if, though, we see only what we imagine will be seen, the insight into the hearts of neutron stars, the explosions of giant suns, the collisions of black holes, and the earliest, as-yet opaque instants of the primordial universe will be wonder enough. The author hopes to have contributed in some small way to this project.

APPENDICES

APPENDIX A

Fields and curvature

Gravitation as described by general relativity is complex, so let us start with a simpler theory: electromagnetism. The same mathematics that predicts electromagnetic waves (light), which are our means of detecting gravitational waves in LIGO, can then be extended by analogy to gravitational waves. This derivation is detailed so that the gravitational analogue can be expressed more simply. Definitions follow Carroll [55]; for a primer in physical mathematics, see Boas [50].

A.1 Mathematical conventions

Let us set geometric, natural units where $1/4\pi\epsilon_0 \rightarrow 1$, $c \rightarrow 1$, $G_C \rightarrow 1$. Greek indices indicate four dimensions, Latin indices three, unless specified otherwise. Assume the Einstein summation convention, *e.g.*, $x_i y^i \equiv \sum_{i=0}^3 x_i y^i$. Anti-symmetrization of indices can be indicated by subscripted square brackets (n.b., all differential forms are antisymmetric), symmetrization by parentheses.

Vectors are typically expressed by reference to index, *e.g.*, a vector v^μ , although implicitly a geometrical vector includes its basis vectors, e^μ such that $\mathbf{v} = v^\mu e_\mu$. Note that when the vector indices are contravariant, the basis vectors are covariant. Let subscript commas indicate ordinary partial derivatives, i.e., $x^j_{,i} \equiv \partial_i x^j$. Semicolons indicate covariant derivatives in general relativity, which reduce to commas (partial

derivatives) in flat space. Suppose we work in a space supplied with a metric tensor, $g_{\mu\nu}$, which acts as a bilinear operator (generalizing the inner product) to produce infinitesimal arclength ds according to $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$ for infinitesimal lengths dx^μ in the direction of, and dual to, basis vectors e_μ . The metric allows index raising and lowering, *e.g.*, $g_{\mu\nu}x^\mu = x_\nu$, $g^{\mu\nu}x_\mu = x^\nu$.

Vectors and matrices can be indicated by boldface, *e.g.*, $\mathbf{v} = v^\mu e_\mu$ is a vector, $\mathbf{M} = M^{\mu\nu} e_\mu e_\nu$ is a matrix, whereas differential forms (dual to vectors) can be indicated by *sans serif*, *e.g.*, $\mathbf{v} = v_\mu dx^\mu$ is a 1-form, $\mathbf{M} = M_{\mu\nu} dx^\mu \wedge dx^\nu$ is a differential 2-form. Tensors generalize vectors and differential forms, including both co- and contra-variant components, and can be built from tensor products \otimes , such as $T_{\rho\sigma}^{\mu\nu} = (\mathbf{M}) \otimes (\mathbf{N}) = (M^{\mu\nu}) \otimes (N_{\rho\sigma}) = (\mathbf{u}^\mu \otimes \mathbf{v}^\nu) \otimes (\mathbf{x}_\rho \otimes \mathbf{y}_\sigma)$.

The wedge \wedge is the exterior or Grassman product, the antisymmetric operator on differential forms that generalizes the cross product, *e.g.*, for the wedge of two 1-forms, $\mathbf{v} \wedge \mathbf{w} = -\mathbf{w} \wedge \mathbf{v}$. Generally, for p -form \mathbf{A} , q -form \mathbf{B} , $(\mathbf{A} \wedge \mathbf{B})_{\mu_1 \dots \mu_{p+q}} = (p+q)! / (p!q!) A_{[\mu_1 \dots \mu_p} B_{\mu_{p+1} \dots \mu_{p+q}]}$. The Levi-Civita symbol is $\epsilon_{\mu_1 \dots \mu_n}$, in \mathbb{R}^n , $+n$ for even permutation, -1 for odd, 0 otherwise. The Hodge star maps differential k -forms in an n -dimensional space to $(n-k)$ -forms, *e.g.*, in 3-space, $\star dx = dy \wedge dz$, $\star dy = dz \wedge dx$, $\star dz = dx \wedge dy$, or from a scalar f , $\star f = (f)dx \wedge dy \wedge dz$, generally in n -space, $(\star \mathbf{A})_{\mu_1 \dots \mu_{n-p}} = 1/(p!) e^{\nu_1 \dots \nu_p}_{\mu_1 \dots \mu_{n-p}} \mathbf{A}_{\nu_1 \dots \nu_p}$.

Index raising and lowering creates a homology between vectors and differential one-forms, $e_\mu \leftrightarrow dx^\mu$ (for coordinate-free math, this is implicit). By itself, d indicates the exterior derivative; df of a scalar (0-form) f is (modulo homology) its gradient ∇ , and $d\mathbf{a}$ maps a k -form to a $(k+1)$ -form, but d^2 is always 0 . Precisely, $(d\mathbf{A})_{\mu_1 \dots \mu_{p+1}} = (p+1)\partial_{[\mu_1} \mathbf{A}_{\mu_2 \dots \mu_{p+1}]}$.

The divergence $\nabla \cdot$ of a vector field \mathbf{X} is homologically given by $\star d \star \mathbf{X}$, the curl

$\nabla \times$ by $\star d\mathbf{X}$, and the Laplacian ∇^2 (in 3-space) or d'Alembertian \square (in 4-space) of a scalar function are given by $\star d \star f$. Covariant exterior derivatives will be noted by $D\mathbf{v} = d\mathbf{v} + \omega \wedge \mathbf{v}$ for a differential form \mathbf{v} and a *connection*; this derivative reduces to $d\mathbf{v}$ in flat space.

The usual, coordinate basis vectors set $e_\mu = \partial_\mu$.

A.2 Electrodynamics

Modern physical theories are usually described as extremizing an action, \mathcal{S} . The action is defined as the integral, over space-time, of a Lagrangian density \mathcal{L} . To integrate, the integral generally requires a metric, $g_{\mu\nu}$. Writing the determinant of the metric as $|g| = \det(g_{\mu\nu})$, the action to be extremized (setting $\delta\mathcal{S}$ to 0) is

$$(A.1) \quad \mathcal{S} = \int \mathcal{L} \sqrt{-|g|} d^4x.$$

The minus sign accounts for the -1 in the Minkowski metric, which physically corresponds to the hyperbolic transformation between space and time in special relativity. Whereas the usual \mathbb{R}^3 Euclidean metric below is g_{ij} , the Minkowski $\mathbb{R}^{1,3}$ metric is $\eta_{\mu\nu}$:

$$(A.2) \quad g_{ij} = \mathbb{I}^3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Given a Minkowski metric $\eta_{\mu'\nu'}$, a curved 4-space $\mathbb{R}^{1,3}$ metric can be written in *vielbein* terms: $e_\mu^{\mu'}$, $e_\nu^{\nu'}$ such that $g_{\mu\nu} = e_\mu^{\mu'} e_\nu^{\nu'} \eta_{\mu'\nu'}$.

The question then becomes how we can define the Lagrangian density for different physical theories. In both electromagnetism and gravitation, the quantity is a kind of curvature, contracted from respectively the Maxwell and Ricci tensors. This curvature, analogous to an energy term, is then extremized subject to a conserved source term: 4-current in electromagnetism, stress-energy and the cosmological constant in general relativity.

As educated coordinate transformations make a great number of problems simpler – a theme of the both the feedforward regression and the data analysis presented in later chapters of this thesis – let us proceed with the derivation in a way that is manifestly covariant under transformations. If desired, one can then extend this treatment to electromagnetism in curved spacetime, such as under the influence of a gravitational wave.

Electromagnetism is theoretically defined by a 1-form $\mathbf{A} = A_\mu dx^\mu$ where A^μ is a vector potential [131] given by $A_\mu = (\phi, -A_x, -A_y, -A_z)$ [92]. The scalar electric potential is ϕ and the magnetic vector potential is A_i . In a given gauge ϕ , the theory is *gauge symmetric* (ϕ invariant) such that the addition of the differential $d\phi$ to \mathbf{A} does not change the theory. This changeless symmetry, in the group $U(1)$, arises because the physics of the theory can be described through the derived Maxwell field tensor $F_{\mu\nu}$, or curvature form \mathbf{F} :

$$(A.3) \quad \mathbf{F} = d\mathbf{A} = (\partial_\mu A_\nu) dx^\mu \wedge dx^\nu = \frac{1}{2} F_{\mu\nu} dx^\mu \wedge dx^\nu,$$

$$(A.4) \quad F_{\mu\nu} = A_{\nu,\mu} - A_{\mu,\nu},$$

By the usual flat-space definitions of potential for an electric field E and magnetic field B , $E_i = -\phi_{,i} - A_{i,0}$, $B_i = \epsilon_{ijk} \partial^j A^k = \epsilon_{ijk} A^{k,j}$. Electrodynamics in curved space

are well-studied; many equations can be translated by replacing partial with covariant derivatives. Therefore, the Maxwell tensor in explicit coordinate form is

$$(A.5) \quad F_{\mu\nu} = \begin{bmatrix} 0 & E_x & E_y & E_z \\ -E_x & 0 & -B_z & B_y \\ -E_y & B_z & 0 & -B_x \\ -E_z & -B_y & B_x & 0 \end{bmatrix}.$$

Maxwell's theory then corresponds to the minimization of the action, finding $\delta S = 0$, for a Lagrangian \mathcal{L} with a stress-energy term for the field itself constrained by an interaction with the current 1-form $J = J_\mu dx^\mu$:

$$(A.6) \quad \mathcal{S} = \int \left(-\frac{1}{4} F^{\mu\nu} F_{\mu\nu} + J_\mu A^\mu \right) \sqrt{-|g|} d^4x.$$

A Hodge dual \star then defines the dual to the Maxwell form, $G \equiv \star F$. Maxwellian electromagnetism then is concisely described by extremizing the action to derive the following equations:

$$(A.7) \quad dF = 0,$$

$$(A.8) \quad d\star F = 4\pi \star J.$$

Applied in flat space, where $g_{\mu\nu} = \eta_{\mu\nu}$, and substituting in the explicit coordinate form of the Maxwell tensor, the theory is easily expressed as four familiar first-order differential equations of the \mathbf{E} and \mathbf{B} fields in terms of electrical charge density ρ and 3-current J^i (see the textbook by Jackson [105]), in (1 time)+(3 space)-dimensions:

$$(A.9) \quad \star d \star \mathbf{E} = E^i_{,i} = \rho,$$

$$(A.10) \quad \star d \mathbf{E} + \partial_t \mathbf{B} = \epsilon^{ijk} E^{k,j} + B^i_{,0} = 0,$$

$$(A.11) \quad \star d \star \mathbf{B} = B^i_{,i} = 0,$$

$$(A.12) \quad \star d \mathbf{B} - \partial_t \mathbf{E} = \epsilon^{ijk} B^{k,j} - E^i_{,0} = 4\pi J^i.$$

Returning to four full dimensions and specifying the *Lorenz gauge*, where we choose any vector potential for which $\star d \star \mathbf{A} = A^\nu_{,\nu} \equiv 0$,

$$(A.13) \quad \star d \star d \mathbf{A} = -4\pi \mathbf{J},$$

with $\square \equiv \partial_\mu \partial^\mu$, this is often written as $\square A^\sigma = -4\pi J^\sigma$ – a wave equation. This derivation has been laborious in order to establish how differential forms allow a cleaner expression of electromagnetism; the differential form equation also is covariant, accurate in curved spacetime. Our task now is to compare with general relativity to see whether a similar wave equation will emerge.

APPENDIX B

Scorpius X-1 real data detection efficiency

Further injection tests have been conducted in real data. The figures in Chapter VI illustrate the efficiency by which TwoSpect recovers simulated stars in a 0.1 Hz band starting at 142.0 Hz in S6 data. Additional injection tests, here, show bands with starting frequencies of 162.0 and 222.0 Hz. In each band, 500 injections were introduced into S6 data. For each injection, physically-expected corrections are made for location and orientation differences between the two LIGO interferometers (H1 and L1). The 500 injections varied f and $\cos \iota$. The injection was made and TwoSpect analyzed the entire stretch of S6 (2009 July 09 to 2010 October 20), approximately 4×10^7 seconds of data, with gaps in the duty cycle of the observatories. These gaps, arising from trains, logging, wind, earthquakes and other factors disrupting the interferometers, have been a concern, and we wished in part to demonstrate TwoSpect's robustness against them.

B.1 Real S6 data: detection efficiency at 162 Hz

Figure B.1 shows the detection efficiency curve of TwoSpect at 162 Hz, where the noise floor of the interferometers is very similar to that of 142 Hz. Results are consistent with Figure 6.1. The sigmoid trace shows the least-squares best fit to the raw data, which is also plotted at several levels of binning.

Relative detection efficiency of simulated signals vs h_0 at [162.0, +0.1] Hz
 lower of two max h_0 reported between H1, L1: $5.79788e-25$




Figure B.1: Detection efficiency of 500 injections (each at H1, L1) into S6 data at 162 Hz, given threshold $\log_{10} p = -7.75$




Figure B.2: Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 162 Hz (injections also done at 142, 222 Hz)

B.1.1 Real S6 data: h_0 recovered vs injected

Figure B.2 correspondingly shows the h_0 upper limit versus h_0 injected curve of TwoSpect at 162 Hz. Analogously, it resembles Figure 6.2. Note that the lower h_0 injected values, on the left side of the plot, return a higher upper limit than their true value. Injections are recovered more rarely as h_0 diminishes, eventually blending into interferometer noise. One of our concerns has been how non-Gaussian data might affect upper limits, and this figure shows that TwoSpect is again highly

Relative detection efficiency of simulated signals vs h_0 at [222.0, +0.1] Hz
 lower of two max h_0 reported between H1, L1: 7.09483e-25




Figure B.3: Detection efficiency of 500 injections (each at H1, L1) into S6 data at 222 Hz, given threshold $\log_{10} p = -7.75$

robust. On the right side of the plot, with increasing h_0 injected, the recovered upper limit parallels the injection value, with a confidence level multiplier to ensure that the UL is higher than the true value at least 95% of the time even with arbitrary underlying $\cos \iota$ and uncertain $a \sin i$.

B.2 Real data: detection efficiency at 222 Hz

Just like Figure B.1, Figure B.3 shows the detection efficiency curve, albeit in the slightly-higher noise floor at 222 Hz.




Figure B.4: Raw h_0 & tentative 95% confidence UL $> 2 \times 10^{-24}$; 500 injections into S6 data at 222 Hz (injections also done at 142, 162 Hz)

B.2.1 Real S6 data: h_0 recovered vs injected

Injections, as in Figure B.2, are shown with recovered upper limits in Figure B.4. TwoSpect successfully recovers injections with the same confidence level factor in each of the three injection bands: 142, 162, and 222 Hz.

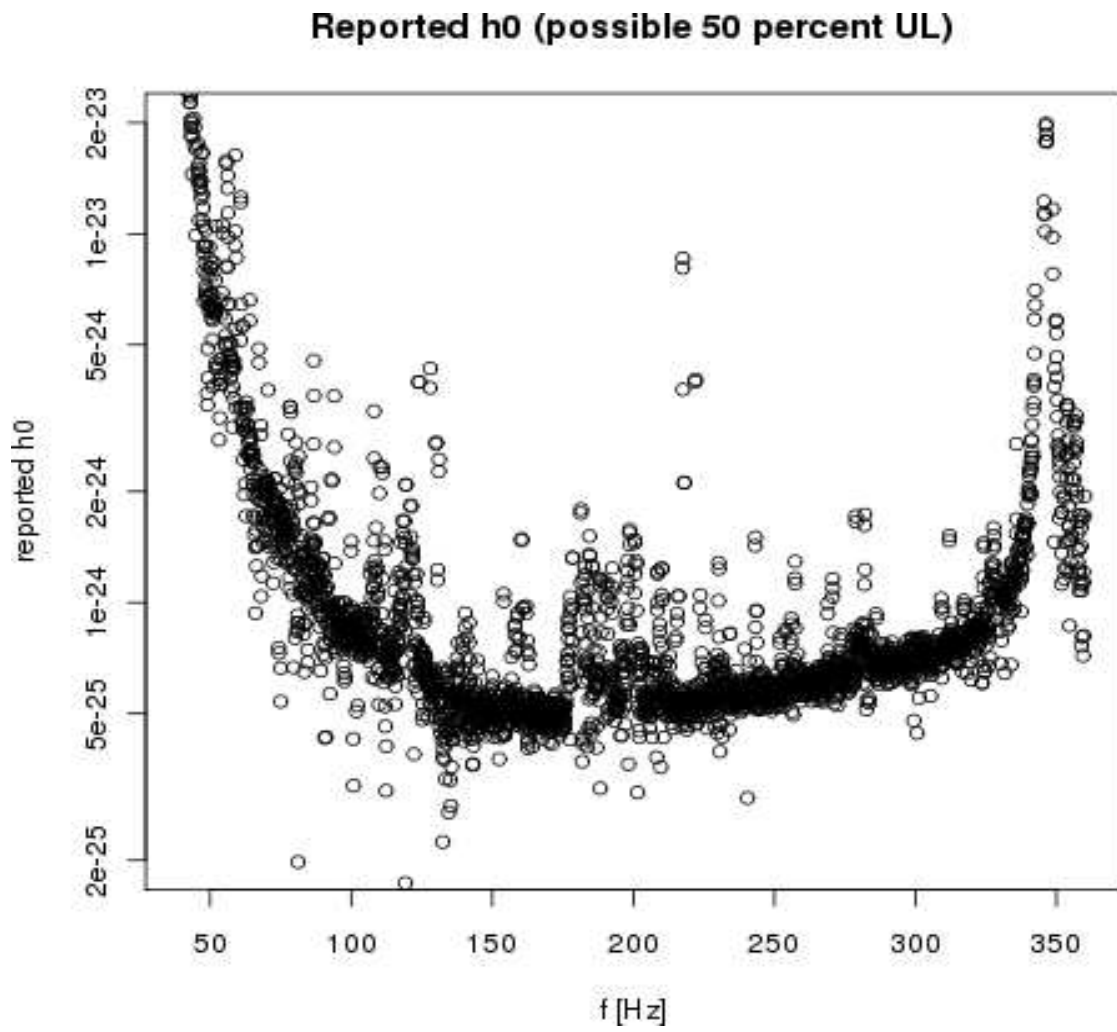


Figure B.5: H1: loudest h_0 in 0.1 Hz bands, effectively circular polarization




Figure B.6: L1: loudest h_0 in 0.1 Hz bands, effectively circular polarization

B.3 S6: Scorpius X-1 upper limits, raw circular output

TwoSpect's h_0 output, proportional to the R statistic, has been tested and is informative for the all-sky search [90, 7]. This h_0 value corresponds to a circularly polarized gravitational wave. Without any corrections, the output is seen in Figures B.5 and B.6. The confidence level of these values in the multi-trial, directed search over the fully-templated df vs f plane required Figures 6.2, B.2, and B.4 to interpret. Corrected by Equation 6.2 and taking the lower of either upper limit when a band is covered by both H1 and L1, the joint upper limit is seen in Figure 6.4. With this knowledge, the author has interpreted the TwoSpect S6 results as a broadband upper limit, from 40 to 2040 Hz, on gravitational wave emission from Scorpius X-1.

APPENDIX C

Scorpius X-1 further preliminary upper limits and outliers

C.1 Upper limits at 90% confidence levels

Figures C.1 and C.2 show tentative ULs that, in the final analysis, correspond to approximately 90% confidence for random polarization. The 90% confidence level requires a correction factor of $h_0 \times (1 + 0.8) [\cos \iota \text{ factor}]$. These figures also illustrate the most sensitive part of the LIGO spectrum.

C.2 Preliminary high frequency outliers

The production analysis of Scorpius X-1 data in Chapter VI yielded the following lists of Stage I outliers for follow-up, detailed in Tables C.1 to C.3, covering coincident outliers (present in both interferometers) between 360 and 2040 Hz. These and the outliers in Table 6.1 are planned to be followed-up with run-averaged spectra and strain histograms, and their relative values in H1 and L1 will be compared. When finished, we will have an indication whether these outliers are artifacts, as expected, or could be signals from Scorpius X-1.




Figure C.1: H1: loudest $h_0 \times (1 + 0.8) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. The 'Tentative h_0 95 percent UL' statement was later revised; the 95% confidence level is higher by 2.0/1.8. This plot shown corresponds to about 90% confidence in random polarization ULs in the final analysis.




Figure C.2: L1: loudest $h_0 \times (1 + 0.8) \times [\cos \iota \text{ factor}]$ in 0.1 Hz bands. The ‘Tentative h_0 95 percent UL’ statement was later revised; the 95% confidence level is higher by 2.0/1.8. This plot shown corresponds to about 90% confidence in random polarization ULs in the final analysis.

Outlier Number	Frequency (Hz)	Explanation
23	360.01	60 Hz mains power line harmonic
24	360.10	60 Hz mains power line harmonic
25	361.47	–
26	375.37	–
27	383.18	–
28	400.10	–
29	403.77	–
30	404.80	–
31	419.79	–
32	419.89	–
33	420.00	–
34	420.10	–
35	420.20	–
36	435.25	–
37	440.10	–
38	448.08	–
39	450.95	–
40	468.10	–
41	479.88	–
42	479.92	–
43	479.95	–
44	480.01	–
45	480.10	–
46	482.21	–
47	500.05	–
48	539.95	–
49	540.00	–
50	540.10	–
51	551.85	–
52	552.02	–
53	568.10	–
54	570.35	–
55	599.78	–
56	599.89	–
57	600.00	–
58	600.06	–
59	646.52	–
60	656.64	–
61	691.00	–
62	691.07	–
63	691.15	–
64	692.14	–
65	719.78	–
66	719.89	–
67	719.97	–
68	720.03	–
69	770.23	–
70	839.99	–
71	870.00	–

Table C.1: List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 360 to 900 Hz.

Outlier Number	Frequency (Hz)	Explanation
72	908.93	–
73	942.74	–
74	957.69	–
75	963.21	–
76	1022.97	–
77	1033.98	–
78	1091.47	–
79	1098.21	–
80	1147.69	–
81	1166.13	–
82	1166.20	–
83	1171.08	–
84	1190.61	–
85	1216.10	–
86	1305.99	–
87	1306.12	–
88	1306.36	–
89	1306.47	–
90	1306.66	–
91	1306.97	–
92	1307.06	–
93	1312.45	–
94	1318.69	–
95	1374.10	–
96	1375.81	–
97	1397.90	–
98	1398.22	–
99	1489.21	–
100	1495.19	–
101	1495.29	–
102	1495.50	–
103	1495.94	–
104	1496.00	–
105	1505.59	–
106	1506.02	L1 detector artifact
107	1514.16	H1 detector artifact
108	1563.03	–
109	1574.28	–
110	1578.44	–
111	1600.32	–
112	1607.70	–
113	1607.90	–
114	1608.00	–
115	1608.08	–
116	1611.51	–
117	1627.74	–

Table C.2: List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 900 to 1700 Hz.

Outlier Number	Frequency (Hz)	Explanation
118	1719.26	–
119	1719.40	–
120	1738.20	–
121	1738.63	–
122	1738.76	–
123	1738.92	–
124	1824.02	–
125	1842.91	–
126	1920.08	–
127	1940.69	–
128	1940.81	–
129	1976.25	–
130	2005.58	–
131	2007.77	–

Table C.3: List of Scorpius X-1 outliers in stage I of the search of S6 data. This list covers 1700 to 2040 Hz.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] Web: <http://web.haystack.mit.edu/SRT/>. Thanks to Haystack Observatory and the Massachusetts Institute of Technology.
- [2] M.G. Aartsen et al. Observation of high-energy astrophysical neutrinos in three years of IceCube data. *Phys. Rev. Lett.*, 113(101101), 2014.
- [3] J. Aasi et al. Directed search for continuous gravitational waves from the Galactic center. *Phys. Rev. D.*, 88(102002), 2013.
- [4] J. Aasi et al. Enhancing the sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photon.*, 7(613), 2013.
- [5] J. Aasi et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown in LIGO-Virgo data from 2009-2010. *Phys. Rev. D.*, 87(022002), 2013.
- [6] J. Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Phys. Rev. D.*, 88(122004), 2013.
- [7] J. Aasi et al. First all-sky search for continuous gravitational waves from unknown sources in binary systems. *Phys. Rev. D*, 90(062010), 2014.
- [8] J. Aasi et al. First searches for optical counterparts to gravitational-wave candidate events. *Astrophys. J.*, 211(7), 2014.
- [9] J. Aasi et al. Gravitational-waves from known pulsars: results from the initial detector era. *Astrophys. J*, 785(119), 2014.
- [10] J. Abadie et al. Calibration of the LIGO gravitational wave detectors in the fifth science run. *NIM-A*, 623:223–240, 2010.
- [11] J. Abadie et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Class. Quant. Grav.*, 27(173001), 2010.
- [12] J. Abadie et al. Directional limits on persistent gravitational waves using LIGO S5 science data. *Phys. Rev. Lett.*, 107(271102), 2011.
- [13] J. Abadie et al. All-sky search for gravitational-wave burst in the second joint LIGO-Virgo run. *Phys. Rev. D*, 85(122007), 2012.
- [14] J. Abadie et al. Search for gravitational waves associated with gamma-ray bursts during LIGO Science Run 6 and Virgo Science Runs 2 and 3. *Astrophys. J*, 760(12), 2012.
- [15] J. Abadie et al. Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys. Rev. D*, 87(082002), 2012.
- [16] J. Abadie et al. Upper limits on a stochastic gravitational-wave background using LIGO and Virgo interferometers at 600-1000 Hz. *Phys. Rev. D*, 85(122001), 2012.
- [17] B. Abbott. *Phys. Rev. D*, 77(022001), 2008.

- [18] B. Abbott, R. Abbott, R. Adhikari, A. Ageev, B. Allen, R. Amin, S.B. Anderson, W.G. Anderson, M. Araya, H. Armandula, and et al. Analysis of LIGO data for gravitational waves from binary neutron stars. *Phys. Rev. D*, 69:122001, 2004.
- [19] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, B. Allen, R. Amin, S.B. Anderson, W.G. Anderson, M. Arain, and et al. Upper limit map of a background of gravitational waves. Technical Report astro-ph/0703234v1, arXiv, 2007-03-09.
- [20] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, B. Allen, R. Amin, S.B. Anderson, W.G. Anderson, M. Araya, and et al. Searching for a stochastic background of gravitational waves with LIGO. Technical Report astro-ph/0608606v1, arXiv, 2006-08-28.
- [21] B. Abbott et al. Detector description and performance for the first coincidence observations between LIGO and GEO. *Nucl. Instrum. Meth.*, A517:154–179, 2004.
- [22] B. Abbott et al. *Phys. Rev. D*, 76(042001), 2007.
- [23] B. Abbott et al. Coherent searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: results from the second LIGO science run. *Phys. Rev. D*, 76(8):082001, 2007.
- [24] B. Abbott et al. Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: results from the second ligo science run. *Phys. Rev. D*, 76:082001, 2007.
- [25] B. Abbott et al. *Phys. Rev. D*, 79(022001), 2009.
- [26] B. Abbott et al. *Phys. Rev. Lett*, 102(111102), 2009.
- [27] F. Acernese et al. Status of Virgo. *Classical and Quantum Gravity*, 22(18):S869, 2005.
- [28] F. Acernese et al. Advanced Virgo baseline design. Technical Report VIR-0027A-09, Virgo, 2009.
- [29] P.A.R. Ade et al. BICEP2 I: detection of B-mode polarization at degree angular scales. Technical Report astro-ph/1403.3985, arXiv, March 2014.
- [30] P.A.R. Ade et al. Planck 2013 results. XVI. cosmological parameters. Technical Report astro-ph/1303.5076, arXiv, March 2014.
- [31] R. Adhikari. *Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [32] S. Adrián-Martínez et al. *J. Cosmol. Astropart. Phys.*, 1306(008), 2013.
- [33] Advanced LIGO Team. Advanced LIGO reference design. Technical Report M060056-v2, LIGO, 2011.
- [34] B. Allen, W.G. Anderson, P.R. Brady, D.A. Brown, and J.D.E. Creighton. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85(12), 2012.
- [35] B. Allen and A.C. Ottewill. Detection of anisotropies in the gravitational-wave stochastic background. *Phys. Rev. D*, 56(2), 1997.
- [36] B. Allen and J.D. Romano. Detecting a stochastic background of gravitational radiation: signal processing strategies and sensitivities. *Phys. Rev. D*, 59:102001, 1999.
- [37] Bruce Allen, Wensheng Hua, and Adrian C. Ottewill. Automatic cross-talk removal from multi-channel data. Technical Report gr-qc/9909083v1, arXiv, 1999.

- [38] N. Andersson, D.I. Jones, and W.C.G. Ho. Implications of an r mode in XTE J1751-305: mass, radius and spin evolution. *Mon. Not. R. Astr. Soc.*, 442:1786–1793, 2014.
- [39] M. Ando. DECIGO and DECIGO pathfinder. *Class. Quant. Grav.*, 27(084010), 2010.
- [40] J.N. Bahcall. Solar neutrinos. i. theoretical. *Phys. Rev. Lett.*, 12(11):300, 1964.
- [41] J.N. Bahcall. Astrophysical neutrinos: 20th century and beyond. *Nucl. Phys. B (Proc. Suppl.)*, 91:9, 2001.
- [42] S. W. Ballmer. *LIGO interferometer operating at design sensitivity with application to gravitational radiometry*. PhD thesis, Massachusetts Institute of Technology, 2006.
- [43] Stefan W Ballmer. A radiometer for stochastic gravitational waves. *Classical and Quantum Gravity*, 23(S179), 2006.
- [44] S.W. Ballmer. A radiometer for stochastic gravitiational waves. *Class. Quantum Grav.*, 23:S179, 2006.
- [45] D. Baumann. Tasi lectures on inflation. Technical Report hep-th/0907.5424, arXiv, 2012-11-30.
- [46] C. Beal and R.M. Wilson. A report on the research and development of the e-bubble collaboration. Web: <http://www.nevis.columbia.edu/reu/2005/ebub.pdf>, August 2005.
- [47] L. Bildsten. Gravitational radiation and rotation of accreting neutron stars. *Astrophys. J. Lett.*, 501(L89), 1998.
- [48] E. Black. Notes on the pound-drever-hall technique. Technical Report T980045-00-D, Laser Interferometer Gravitational Wave Observatory, April 1998.
- [49] E. Black. *Amer. J. Phys.*, 69(79), 2001.
- [50] M. Boas. *Mathematical methods in the physical sciences*. Wiley, second edition, 1983.
- [51] P.R. Brady, T. Creighton, C. Cutler, and B.F. Schutz. Searching for periodic sources with LIGO. *Phys. Rev. D*, 57(4):2101, 1998.
- [52] H.J Bulten, S.G. Crowder, V. Dergachev, D. Galloway, E. Goetz, R.J.G. Jonker, P.D. Lasky, G.D. Meadors, C. Messenger, S. Premachandra, K. Riles, L. Sammut, E.H. Thrane, J.T. Whelan, and Y. Zhang. Gravitational waves from sco x-1: a comparison of search methods. Technical Report P12XXXXX-x0, LIGO, 2014.
- [53] D. Buskulic, L. Derome, R. Flaminio, R. Hermel, J.C. Lacotte, F. Marion, A. Masserot, L. Massonnet, B. Mours, R. Morand, P. Pupp0, D. Verkindt, and M. Yvert. Monitoring and adaptive removal of the power supply harmonics applied to the Virgo readout-out noise. *IJMPD*, 9(1):263–267, 2000.
- [54] J.A.R Caldwell and J.P Ostriker. *ApJ*, 251(61), 1981.
- [55] S.M. Carroll. Lecture notes on general relativity. Technical Report gr-qc/9712019v1, arXiv, 1997-12-03.
- [56] G. Casella and R.L. Berger. *Statistical inference*. Cengage, Stamford, 2nd edition, 2001.
- [57] C.M. Caves. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Phys. Rev. Lett.*, 45(2):75, 1980.
- [58] C.M. Caves. Quantum-mechanical noise in an interferometer. *Phys. Rev. D*, 23(8):1693, 1981.
- [59] D. Chakrabarty et al. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature*, 424:42, 2003.

- [60] S. Chandrasekhar. *Ellipsoidal Figures of Equilibrium*. Yale, New Haven, 1969.
- [61] Sheon S. Y. Chua. *Quantum enhancement of a 4km laser interferometer gravitational-wave detector*. PhD thesis, Australian National University, April 2013.
- [62] S.S.Y. Chua, S. Dwyer, L. Barsotti, D. Sigg, R.M.S. Schofield, V.V. Frolov, K. Kawabe, M. Evans, G.D. Meadors, M. Factourovich, R. Gustafson, N. Smith-Lefebvre, C. Vorvick, M. Landry, A. Khalaidovski, M.S. Stefszky, C.M. Mow-Lowry, B.C. Buckler, D.A. Shaddock, P.K. Lam, R. Schnabel, N. Mavalvala, and D.E. McClelland. Impact of backscattered-light in a squeezing-enhanced interferometric gravitational-wave detector. *Classical and Quantum Gravity*, 31(035017), 2014.
- [63] H. Collins. *Gravity's Shadow: the Search for Gravitational Waves*. Chicago, Chicago, Illinois, 2004.
- [64] H. Collins. *Gravity's Ghost and Big Dog: Scientific Discovery in the Twenty-first Century*. Chicago, Chicago, Illinois, 2013.
- [65] V. Dergachev. Description of PowerFlux algorithms and implementation. Technical Report T050186-00-Z, LIGO, 2005.
- [66] V. Dergachev. *An all-sky search for continuous gravitational waves*. PhD thesis, University of Michigan, 2009.
- [67] Ryan DeRosa, Jennifer C. Driggers, Dani Atkinson, Haixing Miao, Valera Frolov, Michael Landry, Joseph A. Giaime, and Rana X. Adhikari. Global feed-forward vibration isolation in a km scale interferometer. *Classical and Quantum Gravity*, 29(215008), 2012.
- [68] D. Deschrijver et al. *IEEE Microwave Wire. Comp. Lett.*, 18(6), 2008.
- [69] S. Dimopoulos, P.W. Graham, J.M. Hogan, M.A. Kasevich, and S. Rajendran. Gravitational wave detection with atom interferometry. *Physics Letters B*, 678(1):37–40, 2009.
- [70] Katherine Laird Dooley. *Design and performance of high laser power interferometers for gravitational-wave detection*. PhD thesis, University of Florida, 2011.
- [71] R.W.P. Drever et al. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31:97–105, 1983.
- [72] Jennifer C. Driggers, Matthew Evans, Keenan Pepper, and Rana Adhikari. Active noise cancellation in a suspended interferometer. *Rev. Sci. Instrum.*, 83(024501), 2012.
- [73] Jennifer C. Driggers, Jan Harms, and Rana X. Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Phys. Rev. D*, 86(102001), 2012.
- [74] S. Dwyer, L. Barsotti, S.S.Y. Chua, M. Evans, M. Factourovich, D. Gustafson, T. Isogai, K. Kawabe, A. Khalaidovski, P.K. Lam, M. Landry, N. Mavalvala, D. McClelland, G.D. Meadors, C.M. Mow-Lowry, R. Schabel, N. Smith-Lefebvre, M. Stefszky, C. Vorvick, and D. Sigg. Squeezed quadrature fluctuations in a gravitational wave detector using squeezed light. *Opt. Express*, 21(16), 2013.
- [75] Sheila E Dwyer. *Quantum noise reduction using squeezed states in LIGO*. PhD thesis, Massachusetts Institute of Technology, February 2013.
- [76] A. Einstein and N. Rosen. On gravitational waves. *J. Franklin Inst.*, 223:43, 1937.
- [77] P.A. Evans et al. *ApJS*, 203(28), 2012.
- [78] Will M. Farr. *Numerical relativity from a gauge theory perspective*. PhD thesis, Massachusetts Institute of Technology, February 2010.

- [79] G.J. Feldman and R.D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873, 1998.
- [80] L.S. Finn and D. F. Chernoff. Observing binary inspiral in gravitational radiation: one interferometer. *Phys. Rev. D*, 47:2198–2219, 1993.
- [81] É.É Flanagan and S.A. Hughes. The basics of gravitational wave theory. *New J Phys.*, 7(204), 2005.
- [82] R.L. Forward. Wideband laser-interferometer gravitational-radiation experiment. *Phys. Rev. D*, 17(2):379, 1978.
- [83] Nickolas Fotopoulos. Searching for stochastic gravitational waves using co-located interferometric detectors. Master’s thesis, Massachusetts Institute of Technology, September 2006.
- [84] T.F. Fricke, N.D. Smith-Lefebvre, R. Abbott, R. Adhikari, K.L. Dooley, M. Evans, P. Fritschel, V.V. Frolov, K. Kawabe, J.S. Kissel, B.J.J. Slagmolen, and S.J. Waldman. DC readout experiment in Enhanced LIGO. *Classical and Quantum Gravity*, 29(6):065005, 2012.
- [85] Tobin Thomas Fricke. *Homodyne detection for laser-interferometric gravitational wave detectors*. PhD thesis, Louisiana State University, December 2011.
- [86] P. Fritschel. Advanced LIGO systems design. Technical Report T010075-v2, LIGO, 2009.
- [87] Peter Fritschel, Rolf Bork, Gabriela González, Nergis Mavalvala, Dale Ouimette, Haisheng Rong, Daniel Sigg, and Michael Zucker. Readout and control of a power-recycled interferometric gravitational wave antenna. *Appl. Opt.*, 40(28):4988–4998, 2001.
- [88] D. K. Galloway, S. Premachandra, D. Steeghs, T. Marsh, J. Casares, and R. Cornelisse. Precision Ephemerides for Gravitational-wave Searches. I. Sco X-1. *Ap J*, 781:14, January 2014.
- [89] R. Giacconi, H. Gursky, F.R. Paolini, and B.B. Rossi. Evidence for X rays from sources outside the solar system. *Phys. Rev. Lett.*, 9(11), 1962.
- [90] E. Goetz and K. Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Class. Quant. Grav.*, 28(215006), 2011.
- [91] Evan A. Goetz. *Gravitational wave studies: detector calibration and an all-sky search for spinning neutron stars in binary systems*. PhD thesis, University of Michigan, 2010.
- [92] D.J. Griffiths. *Introduction to Electrodynamics*. Prentice Hall, Upper Saddle River, N.J., third edition, 1999.
- [93] L.P. Grishchuk. *Sov. J. Exp. & Theor. Phys.*, 40(409), 1974.
- [94] G. Gurzadyan. *Radioastrofizika*. Akademia nauk armenyanskoi SSR, Erevan, Armenian SSR, 1956. Russian language.
- [95] B. Gustavsen. *IEEE Trans. Power Delivery*, 21(3), 2006.
- [96] B. Gustavsen and A. Semlyen. *IEEE Trans. Power Delivery*, 14(3), 1999.
- [97] G.M. Harry et al. Advanced LIGO: the next generation of gravitational wave detectors. *Class. Quant. Grav.*, 27(084006), 2010.
- [98] G.M. Harry, P. Fritschel, D.A. Shaddock, W. Folkner, and E.S. Phinney. Laser interferometry for the Big Bang Observer. *Class. Quant. Grav.*, 23:4887–4894, 2006.
- [99] S. Hawking. *On the Shoulders of Giants*. Running Press, Philadelphia, Pennsylvania, 2002.

- [100] S. Hild, H. Grote, J. Degallaix, S. Chelkowski, K. Danzmann, A. Freise, M. Hewitson, J. Hough, H. Luck, M. Prijatelj, K.A. Strain, J.R. Smith, and B. Willke. DC-readout of a signal-recycled gravitational wave detector. *Classical and Quantum Gravity*, 26(5):055012, 2009.
- [101] G. Hobbs et al. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Class. Quant. Grav.*, 27(084013), 2010.
- [102] P. Horowitz and W. Hill. *The art of electronics*. Cambridge, Cambridge, 2nd edition, 1989.
- [103] R.A. Hulse and J.H. Taylor. *Astroph. J.*, 195(L51), 1975.
- [104] B. Iyer. LIGO-India: proposal for an interferometric gravitational-wave observatory. Technical Report M1100296, LIGO, 2011.
- [105] J.D. Jackson. *Classical electrodynamics, 3rd edition*. John Wiley & Sons, New York, New York, 1999.
- [106] P. Jaranowski, A. Królak, and B. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: the signal and its detection. *Phys. Rev. D*, 58(063001), 1998.
- [107] F.J. Kerr and D. Lynden-Bell. Review of galactic constants. *Mon. Not. R. Astr. Soc.*, 221:1023–1038, 1986.
- [108] J. Kissel. Mich and prc coupling into darm. Web: http://ilog.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=06/16/2009&anchor_to_scroll_to=2009:06:16:23:51:59-kissel, June 2009.
- [109] Jeffrey S. Kissel. *Calibrating and improving the sensitivity of the LIGO detectors*. PhD thesis, Louisiana State University, December 2010.
- [110] S. Klimenko. Line noise investigation. Web: <http://www.phys.ufl.edu/LIGO/LINE/>, 2000.
- [111] M.J. Koop and L.S. Finn. Gravitational wave detector response in terms of spacetime riemann curvature. Technical Report gr-qc/1310.2871, arXiv, 2013-10-10.
- [112] K. Kuroda et al. *Classical and Quantum Gravity*, 27(084004), 2010.
- [113] J.M. Lattimer and M. Prakash. Neutron star observations: prognosis for equation of state constraints. *Physics reports*, 442:109, 2007.
- [114] U. Lee. Excitation of a non-radial mode in a millisecond X-ray pulsar XTE J1751-305. *Mon. Not. R. Astr. Soc.*, 442:3037–3043, 2014.
- [115] R.B. Leighton, C.D. Anderson, and A.J. Seriff. The energy spectrum of the decay particles and the mass and spin of the mesotron. *Phys. Rev.*, 75(9):1432, 1949.
- [116] LIGO Scientific Collaboration. Matapps software repository. Web: <https://www.lsc-group.phys.uwm.edu/daswg/projects/matapps.html>, 2013.
- [117] L. Lindblom and G. Mendell. *Astrophys. J.*, 444(804), 1995.
- [118] M. Maggiore. Gravitational wave experiments and early universe cosmology. Technical Report gr-qc/990901v4, arXiv, 2000-02-06.
- [119] R.N. Manchester, G.B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility pulsar catalogue. *Astronom. J.*, 129(4), 2005.
- [120] C.B. Markwardt et al. Discovery of a second millisecond accreting pulsar: XTE J1751-305. *Astrophys. J.*, 575:L21–L24, 2002.

- [121] N. Mavalvala. *Alignment issues in laser interferometer gravitational-wave detectors*. PhD thesis, Massachusetts Institute of Technology, 1997.
- [122] C.F. McKee and J.P. Ostriker. *ApJ*, 218(148), 1977.
- [123] G.D. Meadors. Study of the LHO 4K recycling cavity sideband and carrier response to excitation and perturbations. LIGO technical report: LIGO-T050276-00-I (advisor: Richard Gustafson), September 2005.
- [124] G.D. Meadors. E-bubble: seeing the sun's center. Columbia University technical report, web: http://www.nevis.columbia.edu/reu/2006/Meadors_Report.pdf (advisor: Jeremy Dodd), August 2006.
- [125] G.D. Meadors. Searching for a gravitational wave stochastic background. LIGO technical report (advisors: Vuk Mandic, Stefan Ballmer), October 2007.
- [126] G.D. Meadors. Re-searching galactic structure with Reed's radio telescope. Reed College bachelor's thesis (advisor: Robert Reynolds), May 2008.
- [127] Grant David Meadors, Keita Kawabe, and Keith Riles. Increasing LIGO sensitivity by feedforward subtraction of auxiliary length control noise. *Classical and Quantum Gravity*, 31(105014), 2014.
- [128] G. Mendell and K. Wette. Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves. *Class. Quant. Grav.*, 25(144044), 2008.
- [129] C. Messenger and G. Woan. A fast search strategy for gravitational waves from low-mass x-ray binaries. *Classical and Quantum Gravity*, 24(S469), 2007.
- [130] A.A. Michelson. *The velocity of light*. (U of Chicago P, Chicago, 1902).
- [131] C.W. Misner, K.S. Thorne, and J.A. Wheeler. *Gravitation*. Freeman, New York, 1970.
- [132] P.J. Mohr and B.N. Taylor. CODATA recommended values of the fundamental physical constants: 1998. *Rev. Mod. Phys.*, 72(2):351, 2000.
- [133] P.F. Muciaccia, P. Natoli, and N. Vittorio. Fast spherical harmonic analysis: a quick algorithm for generating and/or inverting full-sky, high-resolution cosmic microwave background anisotropy maps. *Astrophys. J.*, 488:L63, 1997.
- [134] R. Narayan and J.P. Ostriker. Pulsar populations and their evolution. *Astrophys. J.*, 352:222–246, 1990.
- [135] C.D. Ott. *Classical and Quantum Gravity*, 26(063001), 2009.
- [136] B.J. Owen. Maximum elastic deformations of compact stars with exotic equations of state. *Phys. Rev. Lett.*, 95:211101, 2005.
- [137] B.J. Owen. How to adapt broad-band gravitational-wave searches for r -modes. *Phys. Rev. D*, 82:104002, 2010.
- [138] B.J. Owen et al. Gravitational waves from hot young rapidly rotating neutron stars. *Phys. Rev. D*, 58:084020, 1998.
- [139] J. Papaloizou and J.E. Pringle. Gravitational radiation and the stability of rotating stars. *M.N.R.A.S.*, 184(501), 1978.
- [140] P.C. Peters and J. Matthews. Gravitational radiation from point masses in a Keplerian orbit. *Phys. Rev.*, 131(435-440), 1963.
- [141] R. Prix. *Neutron stars and pulsars*, page 651. Springer, Berlin, 2009.

- [142] M. Punturo et al. The Einstein Telescope: a third-generation gravitational wave observatory. *Class. Quant. Grav.*, 27:194002, 2010.
- [143] M. Rakhmanov and R. Savage. Measurement of dynamic responses of the 4-km ligo fabry-perot arm cavities. Technical Report G040045-00-Z, Laser Interferometer Gravitational Wave Observatory, February 2004.
- [144] M. Rakhmanov, R.L. Jr. Savage, D.H. Reitze, and D.B. Tanner. Dynamic resonance of light in fabry-perot cavities. Technical Report P010013-00-D, Laser Interferometer Gravitational Wave Observatory, January 2001.
- [145] J.S. Read et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys. Rev. D*, 79:124033, 2009.
- [146] K. Riles. Gravitational waves: sources, detectors and searches. *Prog. in Particle & Nucl. Phys.*, 68(1), 2013.
- [147] Ethan Obie Romero-Severson. *Sexual Volatility and the Spread of HIV*. PhD thesis, University of Michigan, April 2011.
- [148] B.S. Sathyaprakash and B.S. Schutz. *Living Rev. Relativity*, 12(2), 2009.
- [149] P. Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, Singapore, 1994.
- [150] P.R. Saulson. If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves? *Amer. J P*, 65(501), 1997.
- [151] B.F. Schutz. Gravitational wave sources and their detectability. *Class. Quant. Grav.*, 6:1761, 1989.
- [152] P. Shawhan. Gravitational-wave astronomy: observational results and their impact. *Class. Quant. Grav.*, 27(084017), 2010.
- [153] I. Shklovsky. *Cosmic radio waves*. Harvard, Cambridge, MA, 1960. Translated by Richard B. Rodman and Carlos M. Varsavsky.
- [154] D. Sigg. Frequency Response of the LIGO Interferometer. Technical Report T970084, LIGO, 1997.
- [155] R. Simpson. *Introductory electronics: for scientists and engineers*. Allyn and Bacon, Upper Saddle River, N.J., 1987.
- [156] Nicolás de Mateo Smith-Lefebvre. *Techniques for improving the readout sensitivity of gravitational wave antennae*. PhD thesis, Massachusetts Institute of Technology, June 2012.
- [157] C.P. Snow. *The two cultures; with an introduction by Stefan Collini*. Cambridge, Cambridge, 1993.
- [158] T. Strohmayer and S. Mahmoodifar. A non-radial oscillation mode in an accreting millisecond pulsar? *Astrophys. J*, 784(72), 2014.
- [159] H. Takahashi. *Adv. Commun. Syst.*, 1(227), 1965.
- [160] J.R. Taylor. *Introduction to error analysis: the study of uncertainties in physical measurements*. University Science Books, second edition, 1997.
- [161] The LIGO Scientific Collaboration. LALApps repository. Web: <http://www.lsc-group.phys.uwm.edu/daswg/>.

- [162] The LIGO Scientific Collaboration. A gravitational-wave observatory operating beyond the quantum shot-noise limit. *Nature Phys.*, 7(12):962–965, 2011.
- [163] The LIGO Scientific Collaboration and The Virgo Collaboration. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, 460:990–4, 2009.
- [164] The MathWorks, Inc. *MATLAB and Signal Processing Toolbox Release 2012a*. The MathWorks, Inc., Natick, Massachusetts, 2012.
- [165] K.S. Thorne. 300 Years of Gravitation. In S.W. Hawking and W. Israel, editors, *300 Years of Gravitation*. Cambridge, Cambridge, 1987.
- [166] S. van der Putten, H. J. Bulten, J. F. J. van den Brand, and M. Holtrop. Searching for gravitational waves from pulsars in binary systems: An all-sky search. *Journal of Physics Conference Series*, 228(1):012005, May 2010.
- [167] S. Vitale. Space-borne gravitational wave observatories. *Gen. Relativ. Gravit.*, 46(1730), 2014.
- [168] R.V. Wagoner. Gravitational radiation from accreting neutron stars. *Ap. J.*, 278(345), 1984.
- [169] S. Waldman and S. Chua. eligo hanford output mode cleaner repair. Technical Report T1100562, LIGO, 2011.
- [170] D.F. Walls. Squeezed states of light. *Nature*, 306:141–146, 1983.
- [171] J. Weber. Detection and generation of gravitational waves. *Phys. Rev.*, 117(1):306, 1960.
- [172] J.M. Weisberg, D.J. Nice, and J.H. Taylor. Timing measurements of the relativistic binary pulsar PSR B1913+16. *ApJ*, 722:1030–1034, 2010.
- [173] Joel M. Weisberg and Joseph H. Taylor. Relativistic binary pulsar B1913+16: Thirty years of observations and analysis. In *Aspen Winter Conference on Astrophysics: Binary Radio Pulsars*, Aspen, July 2004.
- [174] R. Weiss. The LIGO interferometers: how they work and how well they work. Technical Report G030024-00-D, Laser Interferometer Gravitational Wave Observatory, February 2003.
- [175] P.D. Welch. *IEEE Trans. on Audio Electroacoustics*, AU-15:70–73, 1967.
- [176] L.B. White and B. Boashash. Cross spectral analysis of non-stationary processes. *IEEE Transactions on Information Theory*, 36(4):830–835, 1990.
- [177] N. Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. MIT, Cambridge, Massachusetts, 1949.
- [178] C.M. Will. *Theory and experiment in gravitational physics*. Cambridge, Cambridge, 1993.
- [179] C.M. Will. The confrontation between general relativity and experiment. *Living reviews in relativity*, 4(4), 2001.
- [180] B. Willke et al. *Classical and Quantum Gravity*, 19:1377–1387, 2002.
- [181] M. Zimmermann and E. Szedenits, Jr. Gravitational waves from rotating and precessing rigid bodies: simple models and applications to pulsars. *Phys. Rev. D*, 20:351, 1979.