

Improving the Performance and Sensitivity of Gravitational Wave Detectors

Thesis by
Eric Antonio Quintero

In Partial Fulfillment of the Requirements for the
degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2018
Defended September 25, 2017

© 2018

Eric Antonio Quintero
ORCID: 0000-0002-4269-3445

All rights reserved

ACKNOWLEDGEMENTS

My time at Caltech has benefitted immensely from those who have supported and accompanied me. These words are a poor remuneration for this service; my sincere gratitude goes out to all those with whom I have crossed paths over these years.

I could hardly have conceived of taking this path were it not for the relentless support and encouragement of my parents and grandparents, tireless champions for my academic pursuits and passions. My sister was also a reliable trail-blazer and companion, from reading stories to reading thesis drafts. I hope to deserve their efforts by paying it forward.

It has been a distinct privilege to be part of the LIGO Scientific Collaboration and LIGO Laboratory; witnessing groundbreaking discoveries first-hand, along with the years of tireless work that brought these to fruition. I continue to be astounded by the dedication and ability of those I am fortunate enough to call colleagues, including M. Abernathy, A. Brooks, C. Cahillane, T. Chalermongsak, K. Dooley, J. Eichholz, E. Gustafson, E. Hall, J. Kanner, Z. Korth, H. Miao, X. Ni, N. Robertson, J. Smith, M. Thirugnanasambandam, S. Vass, G. Venugopalan, A. Wade, and H. Yamamoto. I would particularly like to thank those who I have come to consider mentors of one form or another who have taught me so much, both professionally and personally: K. Arai, J. Driggers, J. Rollins, N. Smith, R. Smith, G. Vajente, S. Waldman, and D. Yeaton-Massey. I thank my advisor, R. X. Adhikari, for helping me learn from the past, helping me see what matters, and showing me how to be a scientist.

My deepest thanks to the friends who tolerated my madness throughout these years. A shout-out to those who hail from the dark side: D. Bickerstaff, L. Tenorio, J. Motes, K. Rose (plus little K and H!). Thanks to J. Duarte and I. Saberi for sharing the road, some roofs, and some Ernie's.

In truth, finishing my studies was the lesser development during this time. I do not have the words to thank my family, Evelyn and Sofia, for the joy, love, and support they have provided me during these years. Thankfully, I have some time to make up for that.

ABSTRACT

The field of observational gravitational wave astronomy has begun in earnest, starting with the detection of the strain signal from the binary black hole merger GW₁₅₀₉₁₄ by the Laser Interferometer Gravitational-wave Observatory (LIGO) in 2015. The current incarnation of the LIGO observatories, known as Advanced LIGO, has achieved strain sensitivities on the order of $10^{-23} / \sqrt{\text{Hz}}$ in the hundreds of Hz region, which has enabled unambiguous detection of astrophysical gravitational wave signals. Nevertheless, the scientific output from the LIGO observatories is constrained by the instrumental performance and sensitivity, as there remain many more distant and exotic sources to be observed.

This thesis describes a few topics in experimental gravitational physics, broadly unified by the desire to improve the performance and sensitivity of gravitational wave interferometers. First, it describes an experimental effort to search for a novel form of nonlinear mechanical noise that may be relevant for the ultimate performance of the mirror suspension systems used throughout the instrument. Next, it summarizes work done at the CalTech 40m LIGO controls prototype to realize its fully operational state, and a novel automated controls algorithm developed and tested there that may be useful in simplifying the control of current and future interferometers. Finally, it describes work done on a system to identify and subtract unwanted noise couplings out of recorded aLIGO strain data in an automated fashion. The noise subtraction system applied to GW₁₅₀₉₁₄ is demonstrated to reduce the uncertainties of the black hole mass parameters by about 10 %.

This thesis has the internal LIGO document number [Pr1700380](#).

PUBLISHED CONTENT AND CONTRIBUTIONS

Chapter 2 is adapted from VAJENTE, QUINTERO, et al. [1]. E. A. Quintero performed the construction and characterization of the initial measurement prototype, developed the demodulation analysis, assisted in the conceptual design and assembly of the improved measurement system, and contributed to the writing of the manuscript for publication.

¹G. VAJENTE, E. A. QUINTERO, et al., “An instrument to measure mechanical up-conversion phenomena in metals in the elastic regime”, *Review of Scientific Instruments* 87, 065107 (2016) 10.1063/1.4953114.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	vi
List of Illustrations	vii
List of Tables	ix
Chapter I: Introduction	1
1.1 Gravitational Waves in General Relativity	2
1.2 Making Gravitational Waves Observable	4
1.3 Gravitational Wave Interferometer Design	7
Chapter II: Nonlinear Noise in Terrestrial GW Detector Suspensions	17
2.1 Background	17
2.2 Measurement method	20
2.3 The initial prototype of the measurement system	29
2.4 The improved measurement system	32
2.5 Discussion and outlook	40
Chapter III: The CalTech 40m Prototype Interferometer	41
3.1 Introduction	41
3.2 Length Sensing	44
3.3 Performance Upgrades	47
3.4 Lock Acquisition Procedure	51
3.5 Angular Dynamics and Control	58
3.6 Interferometer Sensitivity and Characterization	60
3.7 Future Work	66
Chapter IV: Automating Interferometer Control	69
4.1 Interferometer locking techniques	70
4.2 The CESAR Algorithm	73
4.3 Testing and Results	80
4.4 Future Work	86
Chapter V: Offline Noise Subtraction	87
5.1 Introduction	87
5.2 Foundations	91
5.3 Methods	98
5.4 Noise Subtraction Applied to GW150914	101
5.5 Nonlinear Noise and Regression	112
5.6 Future Work	120
Appendix A: Delay-line Frequency Discriminator Analysis	121
Bibliography	125

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Michelson Interferometer	5
1.2 Schematic of aLIGO DRSE optical configuration	8
1.3 Angular cavity modes modified by radiation pressure torque	10
1.4 Schematic representation of the ALS subsystem	14
2.1 aLIGO Quadruple Pendulum Test Mass Suspension	19
2.2 Schematic of Michelson Interferometer blade motion measurement	21
2.3 Qualitative illustration of up-conversion noise signals	26
2.4 Simulated results of non-linear noise demodulation analysis	27
2.5 Loaded test blade assembly	30
2.6 Prototype Noise Budget	31
2.7 Limits of nonlinear noise in aLIGO blade springs from prototype experiment.	31
2.8 Rendering of improved measurement apparatus	33
2.9 Optical schematic of improved apparatus	34
2.10 Seismic isolation system of the improved apparatus	37
2.11 Noise performance of the improved apparatus	38
3.1 Schematic of 40m Length Sensing System	46
3.2 ALS Sensitivity Improvement	48
3.3 Effects of PRC Angular Feed-forward	50
3.4 Simulated response of IR transmitted light signals to CARM fluctuations at the 40m prototype	52
3.5 Frequency dependent CARM error signal blending	54
3.6 CARM open loop gain at different stages of lock acquisition	55
3.7 Evolution of Circulating Power during Lock Acquisition	56
3.8 Error Signal Blending During Lock Acquisition	57
3.9 40m Prototype Fundamental Noises	62
3.10 Achieved displacement sensitivity of the 40m prototype interferometer	63
3.11 40m DRFPMI Sensing Matrix	64
3.12 AS port camera during DRFPMI Lock	65
3.13 REFL port during DRFPMI Lock	65
3.14 POP port during DRFPMI Lock	65

3.15	Signal Recycling vs. Signal Extraction for the 40m Prototype	68
4.1	Feedback topology considered for the CESAR algorithm	74
4.2	40m Arm Cavity Signals	81
4.3	Time Domain Simulation of a Fabry-Pérot cavity scan using the CESAR algorithm	84
4.4	IR Lock acquisition of a 40m arm cavity using the CESAR algorithm . .	85
5.1	Representative O1 aLIGO Noise at Low Frequencies	88
5.2	Simple additive signal model	92
5.3	Results of Greedy Ranking Search for H1	103
5.4	Results of Greedy Ranking Search for L1	104
5.5	PSDs of H1 Wiener Filter Validation around GW ₁₅₀₉₁₄	105
5.6	PSDs of L1 Wiener Filter Validation around GW ₁₅₀₉₁₄	106
5.7	Comparison of Hardware Injection Posteriors	107
5.8	Comparison of Filtered GW ₁₅₀₉₁₄ Timeseries	109
5.9	Improvement of GW ₁₅₀₉₁₄ Mass Posteriors	110
5.10	Improvement of GW ₁₅₀₉₁₄ Final Black Hole Parameters	111
5.11	ASDs of mock nonlinear noises	115
5.12	Schematic of nonlinear regression ANN	116
5.13	Neural Network Learning Loss	118
5.14	Regression of mocked bilinear Noise	119
A.1	Schematic of Delay-line Frequency Discriminator	122

LIST OF TABLES

<i>Number</i>	<i>Page</i>
1.1 Summary of desired interferometric conditions for the aLIGO DRSE scheme.	11
1.2 Summary of optical heterodyne signals used for steady-state interferometric length sensing in aLIGO [13, 15]	12
1.3 Summary of optical heterodyne signals used for the third harmonic demodulation technique in aLIGO	15
3.1 Comparison of aLIGO and 40m Prototype Parameters	43
3.2 Sensors used for Interferometer Length Controls	45
5.1 Time intervals used for GW ₁₅₀₉₁₄ noise subtraction analysis	101
5.2 Channels identified by the greedy ranking algorithm.	102
5.3 Comparison of GW ₁₅₀₉₁₄ Parameter Estimation	108

CHAPTER I

INTRODUCTION

Einstein's theory of General Relativity is our current best understanding of the physics of extremely massive objects. [1] The defining characteristic of General Relativity is the concept that gravitation can be understood as the fundamentally geometric effect of space-time curvature. According to MISNER et al. [2], "space-time tells matter how to move; matter tells space-time how to curve." Furthermore, Einstein found that his theory supported wave solutions in the linearised weak-field regime[3], though there was debate over their physical reality and skepticism concerning the feasibility of their detection.

On human scales, however, space-time is extremely stiff. Our only means of investigating systems massive enough to exhibit strongly relativistic behavior is through the observations of astrophysical systems. The Advanced LIGO (aLIGO) observatories[4], located in Hanford, WA and Livingston, LA, were constructed with the goal of observing gravitational waves of astrophysical origin, especially those arising from the dynamics of black holes and neutron stars.

Operating at a space-time strain sensitivity on the order of $10^{-23} / \sqrt{\text{Hz}}$ during their first observing run [5], the aLIGO observatories made the first direct observation of gravitational waves [6], from the merger of two black holes. This discovery has begun the era of gravitational wave astronomy, which will enable many new kinds of observations that will inform and expand our understanding of the universe.

This chapter will briefly summarize the manifestation of gravitational waves within the framework of General Relativity and their effect on matter, the measurement principle and conceptual design of the aLIGO observatories, and the instrumental systems that are relevant for the content of this thesis.

I.1 GRAVITATIONAL WAVES IN GENERAL RELATIVITY

The metric tensor $g_{\mu\nu}$, or simply the metric, is a fundamental quantity of space-time in General Relativity, as it allows us to quantify intervals in space and time as we measure them, via $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$. To find the metric for a given configuration of matter and energy defined by the stress-energy tensor $T_{\mu\nu}$, one must solve the Einstein field equation:

$$G_{\mu\nu} = \frac{8\pi}{c^4} T_{\mu\nu} \quad (I.1)$$

Here, $G_{\mu\nu}$ is the Einstein tensor, which is a compact tensor representation of space-time curvature and a function of the metric.

Distant from any massive objects, $T_{\mu\nu} \approx 0$, and the space-time curvature is small. In this weak-field limit, we may solve the Einstein field equation to find that the metric can be written as

$$g_{\mu\nu} \approx \eta_{\mu\nu} + h_{\mu\nu} \quad (I.2)$$

where $\eta_{\mu\nu}$ is the Minkowski metric of flat space-time, and $h_{\mu\nu}$ is a small perturbation of much smaller magnitude. In the transverse-traceless gauge, $h_{\mu\nu}$ can be understood as a *strain* in space-time itself, causing a relative change in the distance between two points in space-time.

In this weak field limit, it can be shown that the following wave equation holds for $h_{\mu\nu}$ [2]:

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0 \quad (I.3)$$

Thus, empty space-time can support *gravitational waves*, transverse fluctuations of the space-time metric that propagate at c . Plane wave solutions to the wave equation traveling along the z axis can be written in the transverse traceless gauge as:

$$h_{\mu\nu}(z, t) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & -h_+ & h_\times & 0 \\ 0 & h_\times & h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \sin \left(\omega \left(t - \frac{z}{c} \right) \right) \quad (I.4)$$

where h_+ and h_\times are the two possible orthogonal polarizations, and ω is the angular frequency of the wave.

The simplest conditions necessary for the generation of gravitational waves is a time-varying mass quadrupole moment [7]:

$$h_{\mu\nu}(t, r) = \frac{2G}{rc^4} \ddot{I}_{\mu\nu} \left(t - \frac{r}{c} \right) \quad (I.5)$$

For instance, a pair of point masses in circular orbit about their common center of mass with constant orbital frequency ω_o has a quadrupole moment with terms proportional to $\sin^2(\omega_o t)$. This in turn causes the emission of gravitational waves at a frequency of $2\omega_o$.

Naturally, gravitational waves carry energy, and thereby would cause the orbit of two massive objects to decay. In TAYLOR et al. [8] a pulsar consisting of binary neutron stars was observed, and its orbital decay matches the prediction of energy loss via gravitational radiation. This was a key observation which confirmed the physical reality of gravitational waves.

1.2 MAKING GRAVITATIONAL WAVES OBSERVABLE

In order to record a gravitational wave signal with some scientific apparatus, the wave must be transduced into some physical quantity that is more conveniently measured. One of the threads of thought that contributed to the early conceptions of LIGO was that a gravitational wave would affect the time of flight of light traveling between two freely falling test masses. Following the derivation in SAULSON [7], let us consider an incident sinusoidal plane wave in the + polarization with angular frequency ω_{GW} , amplitude h , and wave vector $\mathbf{k} = \frac{\omega_{GW}}{c}\hat{\mathbf{z}}$. Since light follows trajectories satisfying $ds^2 = 0$, the time to travel a distance L along the x -axis in the transverse-traceless gauge is, to first order in h :

$$\tau = \frac{1}{c} \int_0^L dx \sqrt{g_{x,x}} \quad (1.6)$$

$$\approx \frac{1}{c} \int_0^L dx \left(1 + \frac{1}{2}h \cos(\omega_{GW}t)\right) \quad (1.7)$$

$$= \tau_0 \left(1 + \frac{hc}{2} \text{sinc}(\omega_{GW}\tau_0)\right) \quad (1.8)$$

where $\tau_0 := \frac{L}{c}$ is the travel time absent a gravitational wave. A ray of light simultaneously propagating along the y -axis will experience a similar shift in travel time due to the gravitational wave, but with opposite sign, since $h_{x,x} = -h_{y,y}$. Thus, the two rays will experience a *differential* phase shift due to the incident gravitational wave.

Given this differential character, the perpendicular arms of a Michelson laser interferometer (see Figure 1.1) lend themselves naturally to the transduction of a gravitational waves, as the recombination of the beams that have traveled to the end mirrors and back converts the phase shift into intensity fluctuations at the output (a.k.a. anti-symmetric) port of the interferometer.

In this case, the resultant difference in the light travel times corresponds to a differential phase shift according to:

$$\Delta\phi(t) = \frac{2\pi c}{\lambda} \tau_0 \text{sinc}(\omega_{GW}\tau_0) h \cos(\omega_{GW}t) \quad (1.9)$$

This phase shift changes the interference condition of the recombined beams at the output of the interferometer. If the interferometer is illuminated with a plane wave

$$E_{\text{in}} = E_0 e^{i\frac{2\pi c}{\lambda}t} \quad (1.10)$$

The power at the output of the interferometer follows

$$P_{\text{AS}} = \frac{E_{\text{in}}^2}{2} (1 - \cos(\Delta\phi(t))) \quad (1.11)$$




Figure 1.1: Schematic representation of a Michelson interferometer for the transduction of gravitational waves to light intensity at the interferometer output. Laser light of wavelength λ is incident on a beam-splitter, which sends half of the power down the two perpendicular arms of length L . The light reflects off of the end mirrors, and recombines at the beam-splitter. Any differential phase shift experienced by the light while in the arms will manifest as a change in the interference condition at the output.

Thus, the Michelson interferometer gives us a means of transducing gravitational waves into optical power fluctuations, which are readily measurable. Furthermore, this is conceptually similar to how Michelson interferometers are used for measurements of differential length fluctuations of its arms via $\Delta\phi = \frac{2\pi}{\lambda}\Delta L$. Thus, the apparent differential length change due to an incident gravitational wave follows

$$\Delta L(t) = L \operatorname{sinc}(\omega_{GW}\tau_0) h \cos(\omega_{GW}t) \quad (1.12)$$

These expressions can be simplified if we assume that $\omega_{GW}\tau_0 \ll 1$, which physically means that the metric perturbation is effectively constant during the time any given wavefront is propagating up and down an arm. Then, $\operatorname{sinc}(\omega_{GW}\tau_0) \approx 1$ and we recover a simple linear transduction of

$$\frac{\Delta L(t)}{L} = h(t) \quad (1.13)$$

which reinforces the conception of a gravitational wave manifesting as space-time strain.

We can see from these expressions that the sensitivity of the Michelson interferometer unsurprisingly depends explicitly on its baseline length. However, for a gravitational wave signal such as GW150914, with a peak strain on the order of 10^{-21} [6], the required baseline for a simple Michelson would have to be enormous in order to overcome the shot

noise of the laser light. This is incredibly impractical, and furthermore would conflict with the time-of-flight approximation discussed above. Thus, while optical transduction of a gravitational wave with a Michelson interferometer is conceptually simple, we must introduce more sophisticated optical techniques to make it work in practice.

1.3 GRAVITATIONAL WAVE INTERFEROMETER DESIGN

We will now briefly examine the optical configuration and the control systems necessary to operate the aLIGO interferometers. We will make several simplifications, as a complete description is beyond the scope of this work. For fuller treatments, see AASI et al. [4], MIZUNO et al. [9], WARD [10], ADHIKARI [11], STALEY et al. [12], MARTYNOV [13], HALL [14], and IZUMI and SIGG [15].

1.3.1 ENHANCING INTERFEROMETER SENSITIVITY WITH RESONANT CAVITIES

One way to enhance the response of a gravitational wave interferometer would be to use resonant Fabry-Pérot cavities in place of the long baseline Michelson arms. The resonant cavity effectively provides a greater phase response to length fluctuations than a Michelson arm, due to the increased light buildup and storage time. This does, however, decrease the response bandwidth of the interferometer.

However, there are still several practical drawbacks. Even if one separates the light reflected back through the input (symmetric) port via polarization rotation or Faraday isolators to prevent illumination of the laser light source, one is effectively wasting a substantial amount of energy in light that carries a gravitational wave signal. One solution is to introduce an additional resonant cavity through the introduction of a partially transmissive mirror before the beamsplitter, referred to as *power recycling*. Furthermore, one can arrange the differential distance of the input test masses from the beamsplitter to operate on the dark fringe of the Michelson, thereby resonating the maximum available power in the power recycling cavity (PRC). This can greatly increase the power buildup in the arm cavities, further enhancing the sensitivity; this was the basic configuration of the initial iterations of LIGO [16].

A practical drawback of increasing the light power stored in the arm cavities via increased finesse and power recycling gain is the narrowing of the interferometer bandwidth; while the low frequency response is enhanced, the interferometer loses the ability to sense high frequency sources. Since there are many technical noise sources at low frequencies that limit the utility of sensitivity gains in that band, it is desirable to broaden the frequency response of the instrument without reducing the power stored in the arms.

One way of achieving this goal is *resonant sideband extraction* (RSE), first described in MIZUNO [18]. In this scheme, an additional partially transmissive mirror is placed at the




Figure 1.2: Schematic representation of the aLIGO DRSE optical configuration, including the output mode cleaner (OMC). Adapted from IZUMI and SIGG [17].

anti-symmetric port (see Figure 1.2). However, in contrast to a power recycling cavity, the length of this signal recycling cavity formed by the addition of this mirror can be tuned to be anti-resonant for the carrier light. This effectively increases the linewidth of the differential arm mode of the interferometer without reducing the power buildup in the arms. This addition of power recycling and signal extraction cavities to a Fabry-Pérot Michelson interferometer is called *dual-resonant sideband extraction* (DRSE). We will also refer to the vertex recycling cavities as the dual-recycled Michelson, or DRMI. This is the essential optical configuration of the Advanced LIGO detectors.

1.3.2 ANGULAR CAVITY DYNAMICS

The test masses are subject to torque from the radiation pressure of the resonant field depending on the spot position, and the resultant angular motion will move the spot position in turn. Working out the dynamics of this system, two orthogonal modes of radiation pressure induced cavity tilt are evident for each angular degree of freedom (pitch and yaw): a “hard” and “soft” mode [19] (see Figure 1.3). While the hard mode exerts a restoring torque, the soft mode on its own is unstable. When the soft mode’s radiation pressure angular spring constant exceeds that of the restoring mechanical spring constant of the mirror suspensions, the angular motion of the cavity mirrors becomes unstable. This occurs at the power given by:

$$P = \frac{\kappa c(1 - g)}{2L\lambda_a} \quad (1.14)$$

where κ is the mechanical angular spring constant, c is the speed of light, g is the cavity g-factor, L is the cavity length and λ_a is the soft mode eigenvalue of the transformation from pitch and yaw into the hard and soft angular modes.

1.3.3 INTERFEROMETRIC LENGTH SENSING

It is generally necessary to use active feedback control to maintain the desired resonant operating condition in the interferometer, a complex affair due to the existence of multiple degrees of freedom. Modulating the phase of the input laser light field provides a manner of measuring the relative phase shifts of different field components, which in turn carry information about the dynamics of the interferometer lengths. By detecting and demodulating the light fields at various points in the interferometer, it is possible to obtain linear feedback signals for the necessary degrees of freedom.

The length degrees of freedom that lend themselves to interferometric sensing and control are made up of combinations of the lengths labeled in Figure 1.2. Specifically, they are: the differential and common arm cavity lengths (DARM and CARM, respectively), the power recycling and signal recycling cavity lengths (PRCL and SRCL, respectively),




Figure 1.3: Depiction of the common and differential angular cavity motion mode, whose stiffnesses are modified by the torque induced by radiation pressure. Adapted from DOOLEY [19].

and the short Michelson differential displacement (MICH). These are defined as follows:

$$\begin{aligned} \text{DARM: } L_- &\equiv \frac{L_x - L_y}{2} & (1.15) \\ \text{CARM: } L_+ &\equiv \frac{L_x + L_y}{2} \\ \text{PRCL: } l_p &\equiv l'_p + \frac{l_x + l_y}{2} \\ \text{SRCL: } l_s &\equiv l'_s + \frac{l_x + l_y}{2} \\ \text{MICH: } L_- &\equiv \frac{l_x - l_y}{2} \end{aligned}$$

Sideband fields created by frontal phase modulation are a critical component of sensing and controlling the state of the interferometer. In the aLIGO DRSE scheme, two modulation frequencies are used, which we will refer to by $f_{1,2}$ (and $\omega_{1,2}$ for angular frequency

as convenient). In aLIGO, f_1 is approximately 9 MHz and $f_2 \equiv 5f_1$, which was chosen to be an exact multiple so that one RF oscillator can serve to synthesize all of the necessary signals. The desired interferometric conditions for the carrier field and the first order sideband fields is shown in Table 1.1, which is achieved through the choice of macroscopic cavity lengths that create the desired free spectral range and resonance spectrum.

Table 1.1: Summary of desired interferometric conditions for the aLIGO DRSE scheme.

Frequency Offset [Hz]	Arm Cavity	PRC	SRC
0 (Carrier)	Resonant	Resonant	Anti-resonant
$\pm f_1$	Off Resonance	Resonant	Anti-resonant
$\pm f_2$	Off Resonance	Resonant	Resonant

In addition to these resonance conditions, the MICH degree of freedom is held at the dark fringe condition, as mentioned previously. Nominally, this would not allow for any light to be present inside of the signal recycling cavity, rendering signal extraction inoperable. For this reason, a small asymmetry is introduced between l_x and l_y , a.k.a./the Schnupp Asymmetry. This breaks the clean separation between the common and differential vertex signals, but seeds the necessary resonating f_2 fields in the SRC, and additionally allows for the sensing of SRCL and MICH signals at the reflected port.

When a cavity length fluctuates at a particular audio frequency, audio frequency sidebands are imparted onto any resonating or reflected carrier or sideband fields — the amplitude of which depends on the resonance condition for that particular field. Optical beats between various fields can be demodulated to recover the audio frequency signal encoded upon them, as with the common Pound-Drever-Hall cavity stabilization technique [20]. Roughly, we seek to demodulate components of optical beats that look something like [15, 21]:

$$S_i = \left(\frac{\partial E_a}{\partial L_i} \otimes E_b + E_a \otimes \frac{\partial E_b}{\partial L_i} \right) \Delta L_i e^{i(\omega_a - \omega_b)t} \quad (1.16)$$

Here, $E_{a,b}$ and $\omega_{a,b}$ are carrier and/or sideband fields' complex field amplitudes and angular frequencies, L_i is the length degree of freedom of interest, \otimes represents an optical beat, and S_i is an RF component of the total field intensity incident on an RF photodetector. These RF photodetectors sample the field at various interferometer ports, each field having different dependency on the length signals themselves. There are two independent signals that may be derived via demodulation, commonly referred to as the

in-phase (I) and quadrature (Q), such that two degrees of freedom can be derived from one optical heterodyne beat. The sampled ports include the reflected field (REFL), the power-recycling cavity pick-off field (POP) and the output field at the anti-symmetric port (AS).

Table 1.2: Summary of optical heterodyne signals used for steady-state interferometric length sensing in aLIGO [13, 15]

Length	Port, Demod. Frequency	Field Products
DARM	AS f_2	$dE_{\text{Carrier}} \otimes E_{\pm f_2}$
CARM	REFL f_1	$dE_{\text{Carrier}} \otimes E_{\pm f_1}$
PRCL	POP f_1	$dE_{\text{Carrier}} \otimes E_{\pm f_1}, E_{\text{Carrier}} \otimes dE_{\pm f_1}$
SRCL	POP f_2	$E_{\text{Carrier}} \otimes dE_{\pm f_2}$
MICH	POP f_2	$E_{\text{Carrier}} \otimes dE_{\pm f_2}$

IZUMI and SIGG [15] describe the dependence of all of the RF photodetector responses to the various length degrees of freedom for the aLIGO interferometers. The ports and demodulation frequencies for the optical heterodyne readout of the length degrees of freedom is shown in Table 1.2. In practice, signals from the AS port are only used to measure and control DARM, in part because it would be undesirable to have any competing feedback effects from other length control loops. All of the other length degrees of freedom have their length signals derived from the REFL and POP sensors. While the CARM signal generally dominates the content of the REFL and POP signals, using high-bandwidth analog laser frequency control suppresses the CARM influence on the residual POP and REFL signals, allowing for their use to control the DRMI.

Another enhancement designed for aLIGO was the use of an output mode cleaner at the anti-symmetric port, a ring cavity that only transmits the fundamental spatial mode of the carrier field, which is allowed to leak into the SRC and out of the AS port via a small DARM offset. This offset creates a first order coupling of DARM fluctuation to DC power transmitted through the OMC and providing a low noise strain signal, as the OMC rejects unwanted noise couplings due to higher order spatial modes and phase noise of the phase modulation oscillator.

Although all of various cavity mirrors are seismically isolated via passive mechanical suspensions and active isolation platforms, their residual motion is larger than the linear operating regime of the interferometer. Thus, it is necessary to engage active feedback control, applying corrective actuation to maintain all of the cavities at their desired operat-

ing point. The mirror suspensions themselves are equipped with electromagnetic and/or electrostatic actuators which adjust the mirror position. The input laser frequency is also actively stabilized to the CARM degree of freedom by adjusting the error point of the laser’s frequency stabilization system via an additive offset (AO); this in turn uses a combination of broadband phase modulation, and adjustment of the input mode cleaner length.

Apart from the interferometric length stabilization loops, there are hundreds of auxiliary control loops that are necessary for the optimal operation of the detectors, including thermal control of the test mass optics, angular control of the suspended optics, the active seismic isolation platforms, the laser frequency control, and more. This is all coordinate through a real-time digital control system, in which signals are digitized, user defined logic is applied to the signals, and actuation signals are synthesized and sent to actuators. The digital nature of this system allows for complex and interconnected logic between the various subsystems. The majority of the real-time control system operates at a sampling rate of 16 384 Hz, which in practice affords digital control loops with control bandwidths of up to 100–200 Hz.

1.3.4 ADVANCED LIGO LOCKING

As will be discussed in some more detail in Section 4.1, bringing the interferometer to its operating point — where all cavities are at the desired resonant condition and all the length degrees of freedom are under stable feedback control — is a very nontrivial affair, as the various cavities interact with each other in complex nonlinear ways. Here we will briefly describe two strategies employed in aLIGO that make the locking acquisition process more deterministic. Substantial information about development of the aLIGO locking procedures can be found in STALEY et al. [12] and MARTYNOV [13].

The primary complication that these strategies seek to simplify is the separation of initial control of the arm cavities and the DRMI cavities. The effective reflectivity of the arm cavities for the carrier light changes dramatically if the arm is flashing in and out of resonance, which makes the usual optical heterodyne signal unsuitable for stable control, as they depend directly on the carrier field amplitude (see Table 1.2).

ARM LENGTH STABILIZATION

The Arm Length Stabilization (ALS) system, as developed in MULLAVEY et al. [22] and IZUMI et al. [23], is a scheme employing frequency-doubled green laser light from auxil-




Figure 1.4: Schematic representation of the ALS subsystem. The green PDH locking components and second arm cavity are excluded for simplicity. Adapted from IZUMI [24].

ary IR lasers that is PDH frequency locked to a single arm cavity to derive a cavity length signal that is independent from the main IR laser light that is present in the DRMI and subject to complex cavity couplings. A dichroic mirror coating on the test masses creates a lower finesse cavity for the green light. The transmitted green light is then interfered with frequency doubled light derived from the main input laser, as shown in Figure 1.4.

Because of the PDH frequency locking, the green light frequency follows the free swinging arm cavity motion within the PDH control bandwidth. Then, the optical beat between main laser green light and transmitted green light, which can be measured with a phase locked loop, carries information regarding the relative motion between the arm cavity and the input laser frequency, which in turn can be used as a new error signal to stabilize the arm cavity length with respect to the main laser frequency.

Since we desire to control the arms in the CARM/DARM basis, aLIGO uses a second optical beat between the green transmitted light from both arm cavities to create a DARM signal. Given the length of the arm cavity and the test mass seismic isolation, the single arm green beat with the main laser mostly tracks that laser frequency noise, as is needed for CARM control. As long as the optical beat frequencies are accurately tracked, and the auxiliary green light remains stably locked to the arm cavity, the arm cavities can be controlled to remain anywhere along their IR resonance profile through an offset in the CARM control loop. This allows the arms to be stably held off of IR resonance while acquiring control of the DRMI.

THIRD HARMONIC DEMODULATION TECHNIQUE

In order to reduce the influence of carrier light fluctuation, ARAI [25] showed that one could derive signals from the optical beat at three times the modulation frequencies. While we usually only consider the first order sidebands that result from phase modulation, there are in reality many higher order PM sidebands that result, albeit with much reduced magnitude. Due to the FSR of the cavities, the second order PM sidebands can be used as stable local oscillators that will not interact with the cavities, and thereby provide a consistent field product that only depends on the interaction of the first-order sideband fields with the DRMI. (See Table 1.3) In truth, the sensitivity of these signals to the carrier is not zero, but the influence is reduced enough that the DRMI may be maintained stably on resonance regardless of the state of carrier light in the arm cavities.

Table 1.3: Summary of optical heterodyne signals used for the third harmonic demodulation technique in aLIGO

Length	Port, Demod. Frequency	Field Products
PRCL	REFL $3f_1$	$dE_{\pm f_1} \otimes E_{\mp 2f_1}$
SRCL	REFL $3f_2$	$dE_{\pm f_2} \otimes E_{\mp 2f_2}$
MICH	REFL $3f_2$	$dE_{\pm f_2} \otimes E_{\mp 2f_2}$

aLIGO LOCK ACQUISITION SEQUENCE

Ideally, one could acquire control of the arm cavities with the ALS system, acquire control of the DRMI with the $3f$ signals, and directly transition from ALS to IR optical heterodyne signals on resonance. Unfortunately, in part due to the dichroic mirror coatings for the ETMs currently not meeting their specification, it is necessary to slowly approach resonance while the arms are under ALS control and use an intermediate IR signal, as the residual cavity fluctuations while using ALS are wider than the interferometer linewidth [12].

Thus, the aLIGO locking procedure proceeds as follows:

1. The auxiliary lasers at each end station are frequency locked to the arm cavities via PDH reflection locking of the green light.
2. ALS control of the arm cavities is engaged, using the green optical beat signals. The DARM control is held with zero offset, while an offset of a few single arm

linewidths is introduced in the CARM control to avoid IR resonance of the main laser light.

3. Control of the DRMI is acquired using the $3f$ length signals.
4. The CARM offset is reduced to an intermediate point, where some IR light begins to resonate in the arm cavities.
5. CARM control is transitioned to use the DC power transmission through the arms as an error signal.
6. DARM control is transitioned to the ASf_2 RF optical heterodyne signal.
7. The CARM offset is further reduced until the REFL f_1 is viable, upon which it is used as the error signal, and the CARM offset is reduced to zero.
8. The interferometer is now fully resonating and is now transitioned to the lower noise DRMI length signals from the $1f$ POP RFPDs.
9. A small DARM offset, on the order of a few picometers, is introduced to allow some carrier light to exit the AS port. The OMC is locked to this carrier light, and used as the final, low-noise, DARM length signal.

CHAPTER 2

NONLINEAR NOISE IN TERRESTRIAL GW DETECTOR SUSPENSIONS

2.1 BACKGROUND

The Advanced LIGO detectors are ground-based laser interferometers intended to observe gravitational waves [4]. To be successful, the LIGO detectors must reach an extreme displacement sensitivity in the audio frequency band. At the low frequency end of this band (10–20 Hz), the horizontal motion of the 40 kg fused silica mirrors, acting as test masses, must be only about $10^{-19} \text{ m}/\sqrt{\text{Hz}}$. Since the detector is located on the ground, it employs complex seismic isolation systems to reduce the contamination of the sensitivity by local seismic activity. The Advanced LIGO test mass suspension system [26, 27], shown schematically in Figure 2.1, consists of a quadruple pendulum for horizontal isolation and incorporates three stages of 50 cm long cantilever spring pairs, made of maraging steel [28] for vertical isolation. The suspension wires are also made of maraging steel, with the exception of the lowest wires, which are made of fused silica bonded to the mirror, to reduce thermal noise [29]. Any mechanical noise occurring within the cantilevers or in the wires will propagate to the test mass at some level. In particular, the lowest set of cantilever springs, which are installed in the second mass from the top (the upper intermediate stage, or UIM), will couple most strongly to vertical displacement of the test mass, since there is less vertical isolation between them and the test mass than for those cantilevers that are higher up the chain. In turn, vertical motion of the test mass will couple to its horizontal displacement, which is the degree of freedom which is measured to detect gravitational waves, due to mechanical imbalances in the suspension system and, ultimately, to the Earth’s curvature [26]. Thus, even if the impulsive strain

events at the test mass are small, their combined influence can introduce background noise which could limit the interferometer sensitivity.

Metals can also exhibit creep noise [30]. Although the underlying micro-mechanics of mechanical up-conversion and creep may be related, creep has a event rate that decreases quickly after the initial stress, and experimental investigations have shown that the creep can be reduced with the use of maraging steel [31–34]. Our experiment however focuses on mechanical events that are continuously triggered by a time varying external perturbation, such as the Advanced LIGO suspension cantilevers are subjected to by the local micro-seismic activity of the ground. In addition, since it is virtually impossible to distinguish between events happening in the cantilevers from those happening in the suspension wires or in the clamps, our system mimics as close as possible the Advanced LIGO configuration for cantilevers, wires and clamps.

It is known that crackling noise occurs when metals are stressed in the plastic regime. In the Advanced LIGO suspensions, however, the cantilever and wires loads are solidly within the macroscopically elastic regime, specifically $\approx 50\%$ of the yield stress [28]. To the best of our knowledge, there has been no in-depth investigation for potential discrete, stochastic deviation from linear mechanical behavior in crystalline materials this far below the engineering yield stress. Nevertheless, we can borrow insights from the existing experiments and theories which have studied the problem in the plastic regime. First of all, micro-pillar compression tests have demonstrated the dependence of event size on the driving mode: under load-controlled mode large bursts are seen, while displacement-controlled mode leads to slipping events of smaller sizes [35]. It has also been shown that the distribution of the size of crackling events depends on the stress and stress rate [36], being skewed toward smaller sizes for lower external stress and stress rate. These predictions have only been experimentally validated in the plastic regime, where burst sizes are large enough to exceed instrumental noises. Thus, the question of the existence of non-linear mechanical noise in the elastic regime remains open. Furthermore, the non-linear mechanical noise we are trying to characterize in the elastic regime — which we will hereafter refer to as *up-conversion noise* — can have intrinsically different physical origins from the crackling noise studied in the plastic regime.

Due to the novelty of this investigation and lack of a micro-mechanical model which predicts the exact form of the expected signals, our experiment follows a different approach with respect to what has been done previously. Instead of trying to detect individual slip events, we focus on the stochastic noise that would arise as a sum of a large number of small events. Such noise might have a non stationary nature, with power depend-




Figure 2.1: The aLIGO test mass suspension system consists of a quadruple pendulum incorporating 3 stages of maraging steel cantilever springs. Drawing adapted from ASTON et al. [26].

ing on the external perturbation. In particular, given the performance of the Advanced LIGO seismic isolation system, we expect that residual low frequency motion of the suspension cantilevers could excite broadband mechanical noise, resulting in non-linear up-conversion and a broadband power spectrum of displacement noise, time-correlated with the driving force or force rate. Thus, we may expect an increased rate of larger events when the stress or stress rate of the cantilever is increased with respect to the equilibrium position.

2.2 MEASUREMENT METHOD

A direct measurement of the horizontal displacement noise introduced by up-conversion events in the Advanced LIGO suspension cantilevers would be impossible except with an apparatus which has the same displacement sensitivity as the Advanced LIGO interferometers [4]. However, any up-conversion noise at the level of the UIM cantilever springs will be attenuated by the additional vertical isolation provided by the lower suspension stages and by the relatively small coupling of vertical to horizontal test mass motion. For this reason, the sensitivity of our apparatus does not have to reach the Advanced LIGO level if we measure the vertical displacement of the cantilevers directly. A rough estimate of the sensitivity needed in our setup goes as follows. At 10 Hz the Advanced LIGO design displacement noise is of the order of $4 \times 10^{-19} \text{ m}/\sqrt{\text{Hz}}$ [4]. Assuming a coupling of vertical to horizontal of the order of 10^{-4} due mainly to earth's curvature, this corresponds to a vertical displacement noise, at the test mass level, of $4 \times 10^{-15} \text{ m}/\sqrt{\text{Hz}}$, without assuming any additional isolation between the test mass and the maraging cantilevers. This estimate has been confirmed using a model of the suspension system. Therefore we set a target sensitivity for our system of $10^{-15} \text{ m}/\sqrt{\text{Hz}}$ at 10–20 Hz, which will be sufficient to probe up-conversion noise amplitudes relevant for Advanced LIGO.

However, as the magnitude of up-conversion noise is unknown and likely small, background noise sources will be a strong limiting factor in any measurement attempt. In view of this, an important component of our measurement strategy is to make a differential measurement of the motion of two cantilever springs that are arranged to make their response to background noise sources, such as seismic activity in the lab, as equal as possible. Since up-conversion noise occurs incoherently in each cantilever, a measurement of the cantilevers' differential displacement will be sensitive to up-conversion noise while rejecting any noise they have in common.

We choose to make this kind of differential measurement with a Michelson interferometer [37], wherein a laser beam incident on a partially transmissive mirror is split into two beams which are retro-reflected by end mirrors mounted on masses suspended by the cantilever springs being tested for up-conversion noise, which we will refer to as the *test cantilevers* (see figure 2.2). When these beams recombine at the beam splitter, they will interfere constructively or destructively depending on the differential path length the beams traversed on their way to the end mirrors, providing a means of transduction of motion to optical power, which is then measured by a photodiode. Further detail describing the optical signals present in the Michelson interferometer can be found in Section 2.2.1.




Figure 2.2: Simplified schematic of the Michelson Interferometer layout employed. x_1 and x_2 represent the motion of mirrors 1 and 2, which are suspended from test cantilevers 1 and 2 (not shown). “SY” and “AS” refer to the “symmetric” and “anti-symmetric” ports, respectively.

Furthermore, rather than trying to measure the up-conversion events due to ambient seismic motion, we can apply a controlled driving force, equal for both test cantilevers (*common mode*) to excite more up-conversion events. As will be explained in more detail in Section 2.2.3, this also allows us to enhance the apparatus’ sensitivity by incorporating our knowledge of the drive, and can provide insight into the micro-mechanical nature of the up-conversion events.

2.2.1 UTILITY OF THE MICHELSON INTERFEROMETER

A representative diagram of the interferometer design is shown in Figure 2.2. The quantities x_1 and x_2 represent the vertical displacement of the end mirrors, M1 and M2, from the equilibrium position of the cantilever springs they are suspended from. For the sake of convenience, we will assume that the optical path lengths between the beam splitter and each end mirror are equal. As a first step, we consider the optical signal present at the photodiode at the anti-symmetric port of the interferometer, labeled “AS” in Figure 2.2.

We will consider the laser light's field amplitude incident on the beam splitter to be of the form

$$E_{\text{in}} = E_0 e^{i\omega t} \quad (2.1)$$

where ω denotes the frequency of the laser light source, related to the wavelength $\lambda = \frac{2\pi c}{\omega}$ and to the wave number $k = \frac{2\pi}{\lambda}$. Then, the field exiting the beam splitter at the AS port will be the superposition of the fields which independently traversed the two arms of the interferometer:

$$E_{\text{AS}} = E_1 - E_2, \quad \text{where} \quad (2.2)$$

$$E_i = \frac{1}{2} E_{\text{in}} e^{2ik(L+x_i)}, \quad (2.3)$$

where L is the distance from the beam splitter to the equilibrium point of each end mirror and, in the second equation, $i = 1, 2$ refer to the field propagating in the two interferometer arms. The minus sign in equation 2.2 is due to the fact that the field returning from mirror 2 reflects off of the back surface of the beam splitter, and thus experiences a π phase shift relative to the light which reflected off of mirror 1 and the front surface of the beam splitter.

Thus, the field amplitude and intensity at the AS port are given by

$$E_{\text{AS}} = \frac{1}{2} E_{\text{in}} e^{ik2L} (e^{ikx_1} - e^{ikx_2}) \quad (2.4)$$

$$I_{\text{AS}} = E_{\text{AS}}^* E_{\text{AS}} = \frac{1}{2} E_{\text{in}}^2 [1 - \cos(k(x_1 - x_2))] \quad (2.5)$$

Equation 2.5 shows the optical power measured by a photodiode at the AS port, which is a function of the positions of the two end mirrors. Thus, the Michelson interferometer naturally provides an optical signal that is only sensitive to differential displacements of the two test cantilevers, providing, ideally, an infinite rejection of common mode motion.

However, the linear range of the signal is limited by the wavelength of the light used, as can be seen by the sinusoidal functions of the displacement. So, to ensure linear readout, active feedback is used to keep the interferometer at the proper operating point [38]. Specifically, we employ a feedback loop that stabilizes the differential displacement by applying differential force to the tip of the test cantilevers that is continuously tuned to maintain constant power incident on the photodiode. This does not reduce the information present in the system, as one can reconstruct the linear open-loop behavior of the system by appropriately combining the feedback control and error signals.

A potential flaw with this optical readout scheme is the inability to distinguish fluctuations in the laser source intensity from real displacement fluctuations; the signal described

in equation 2.5 is linearly proportional to the input laser power. The solution to this issue is to read out both interferometer outputs: the symmetric port, as described above, in addition to the anti-symmetric port, “SY”. By injecting the input beam at an angle, one can cleanly separate both output beams.

The signal at the symmetric port can be easily written down by conservation of energy from equation 2.5:

$$I_{SY} = \frac{1}{2}E_{in}^2 (1 + \cos(k(x_1 - x_2))) \quad (2.6)$$

We can now construct a signal that suppresses the linear coupling of intensity to position readout by subtracting the two signals, either with analog electronics or within a digital control and data acquisition system:

$$\begin{aligned} x_e &= I_{SY} - I_{AS} \\ &= E_{in}^2 \cos(k(x_1 - x_2)) \end{aligned} \quad (2.7)$$

With the aid of the feedback control loop, we actuate on the differential mirror positions, which constrains this signal to remain close to zero, which in turn eliminates the direct linear coupling of laser intensity noise to our displacement signal. This interference condition is often called the *half fringe*, meaning that the power at the two detectors is equal: half of the input power.

An ideal Michelson interferometer is insensitive to laser frequency noise. However, any mismatch, ΔL , in the length of the two Michelson arms will result in a coupling of laser frequency noise to the output port powers. Indeed, starting from equation 2.2 and considering that a variation in the laser frequency corresponds to a variation of k , it is easy to show that a change in the laser frequency $\delta\omega$ will introduce a power variation equivalent to a differential displacement of the end mirror δx , given by:

$$\delta x = \frac{\Delta L}{\omega} \delta\omega \quad (2.8)$$

Therefore frequency noise of the laser can be ignored if the length of the two Michelson arms is equalized to within a good accuracy. As discussed below, the safest approach is to implement a way to remotely equalize the length of the two arms.

Finally, although a Michelson interferometer is first order insensitive to translation and rotation of the input beam, any misalignment of the end mirrors translates in a change of the interference of the beams at the output ports, resulting in a reduction of the optical gain. If the vertical motion of the two test cantilevers translates to a differential angular

motion of the end mirrors, an additional up-conversion mechanism will be present that can mimic the one we are looking for, as discussed below. For this reason it is important to install the two blades in an anti-parallel configuration, and decouple efficiently the mirror angular motion from the cantilever. As discussed below, this is done by suspending the mirrors with thin wires.

2.2.2 UP-CONVERSION NOISE MODEL

Absent a detailed micro-mechanical model, we can instead use a simple phenomenological model informed by analogous physical processes, such as Barkhausen noise in magnetized materials [39], to design our analysis method. Specifically, we model the effect of up-conversion events in a cantilever spring as a stochastic displacement noise with unknown spectral properties, but with a magnitude determined by the applied force and/or its derivative. Since this stochastic noise is the result of the sum of a large number of microscopic events, its statistical properties depend on the rate and size distribution of such events. We expect those properties to depend both on material properties, and on the local stress or stress rate in each part of the cantilever. We focus our attention to the case of a cantilever which is subject to a possibly large static load and a time varying external perturbation, typically induced by an external low frequency force. The static load might induce some creep in the metal, but this phenomenon is well known and its magnitude reduces over time [30].

Thus, we consider a cantilever subject to a time dependent force $F(t)$, with a characteristic frequency below the macroscopic resonance of the cantilever. In this case, the local microscopic stress varies over time following the external drive. Thus, we write the up-conversion noise contribution to the displacement as:

$$x_{\text{up-conversion}}(t) = \chi \left[F(t) \right] \delta x_f(t) + \theta \left[\dot{F}(t) \right] \delta x_j(t) \quad (2.9)$$

where δx_f and δx_j are stochastic processes representing the force- and jerk-dependent up-conversion, χ and θ are the functional forms of the noise dependence on the applied force and its derivative. They reflect the intensity of up-conversion noise in the specific cantilever, and they may be a function of drive frequency and amplitude, in addition to the static load, cantilever geometry and material properties.

An important observation can be made at this point. The δx_f and $\delta x_{\text{ext}j}$ terms will in general have nonzero spectral content at the frequencies in the LIGO detection band (10–5000 Hz). On the other hand, the typical force $F(t)$ on the cantilevers is due to the residual coupling of ground motion through the suspension system, which acts as a steep

low pass filter with corner frequency of the order of a few Hz. Therefore, while $F(t)$ has very low amplitude at higher frequencies, the large amplitudes at low frequency can excite up-conversion events, generating noise in the audio band. Thus, it is important to measure the level of non-linear up-conversions from large static strains and low-frequency motions to noise in the audio band, where Advanced LIGO is most sensitive to gravitational waves.

2.2.3 DEMODULATION ANALYSIS

In order to excite up-conversion events we introduce a low frequency, common mode excitation in the two test cantilevers through the application of a force in the form of $F(t) = F_0 \sin(\omega_d t)$, that is much larger than the residual seismic motion.

To mimic the conditions in the Advanced LIGO suspension, this time-varying force is small when compared to the static load applied on the cantilever; in our test setup, for example, the static load is of the order of 20 N, while the time varying common mode drive is of the order of few mN. Therefore we can expand, by a Taylor series, the two functions χ and θ around the point corresponding to the static load. With this assumption, the individual cantilever displacements are given as:

$$x_i(t) = \frac{F_0 \sin(\omega_d t)}{k} + \frac{\alpha}{\sqrt{2}} F_0 \sin(\omega_d t) \delta x_{f,i}(t) + \frac{\beta}{\sqrt{2}} F_0 \omega_d \cos(\omega_d t) \delta x_{j,i}(t) \quad (2.10)$$

Here, k is the cantilever elastic spring constant and the factors $\sqrt{2}$ have been introduced to simplify later equations. In the differential displacement signal, the elastic responses of the cantilevers will cancel out, leaving the incoherent sums of the up-conversion noise terms.

In practice, this will not be the only signal present; there will be many background noise sources. It can occur that the displacement due to up-conversion noise is smaller than the background noise in a given frequency band, and thus will not be directly observed. However, many background noise sources — such as shot noise on the photodetectors, intensity noise of the laser, and seismic noise that couples through mechanical asymmetries — can be assumed to be quasi-stationary, and have no dependence on the common mode drive. Instead, it can be seen that the up-conversion signal written above varies in a known way with the applied drive; using this additional information we can potentially extract this signal from underneath stationary background noise, as will now be explained.




Figure 2.3: A qualitative illustration of the signal described in Equation 2.11, with simulated data. In this case we have assumed that the up-conversion noise is proportional to the derivative of the external drive, therefore the up-conversion noise power is larger when the sinusoidal excitation crosses zero.

We condense the total differential displacement due to up-conversion noise and some stationary background noise $n(t)$, combining the incoherent sums of up-conversion noise in each cantilever into a single term:

$$\Delta x(t) = n(t) + \alpha F_0 \sin(\omega_d t) \delta x_f + \beta F_0 \omega_d \cos(\omega_d t) \delta x_j \quad (2.11)$$

An example of how the δx_j term manifests itself is shown in Figure 2.3.

We now want to take advantage of the periodicity and phase of the envelope of the up-conversion noise processes, and analyze the instantaneous power of the displacement time series, i.e. its square. Simple algebraic computations yield

$$\begin{aligned} \Delta x^2(t) &= 2n(t)F_0 \left[\alpha \sin(\omega_d t) \delta x_f + \beta \omega_d \cos(\omega_d t) \delta x_j \right] \\ &+ n(t)^2 + \frac{F_0}{2} \left[\alpha^2 \delta x_f^2 + \beta^2 \omega_d^2 \delta x_j^2 \right] \\ &+ \cos(2\omega_d t) \frac{F_0}{2} \left[-\alpha^2 \delta x_f^2 + \beta^2 \omega_d^2 \delta x_j^2 \right] \\ &+ \sin(2\omega_d t) F_0 \alpha \beta \delta x_f \delta x_j \end{aligned} \quad (2.12)$$




Figure 2.4: Result of the demodulation described in Equation 2.13 of 30 minutes of simulated data with background and up-conversion noise levels as in Figure 2.3. The distributions are clearly separated, showing a strong up-conversion noise signature.

We can average the above quantity over a period longer than the typical time scale of the random processes, and slower than the external drive sinusoid. Assuming that $n(t)$, $\delta x_f(t)$, and $\delta x_j(t)$ are independent zero-mean random noise processes, the first and last line will have expectation values of zero, while the second line will have some constant expectation value. In contrast, the $\cos(2\omega_d t)$ term provides a time varying component at a known frequency, with a known phase with respect to the driving force. Writing the Fourier transform of the power signal as $\tilde{P}(\omega)$, we can take the expectation value at $2\omega_d$, or *demodulate* the drive-modulated signal, to see the power fluctuations due to up-conversion events:

$$\langle \tilde{P}(2\omega_d) \rangle = \frac{F_0}{4} \left(-\alpha^2 \delta x_f^2 + \beta^2 \omega_d^2 \delta x_j^2 \right) \quad (2.13)$$

In addition, by integrating for many cycles, the determination of the up-conversion noise amplitude of equation 2.13 improves proportionally to the square root of integration time. Thus, it is possible to increase the measurement time to find up-conversion noise power varying with the modeled phase and frequency, even below the background noise.

Furthermore, the common drive can be switched off, which should result in a demodula-

tion result of zero, on average. Thus, we can sample the magnitude of the demodulation amplitude in two different states: with the drive on and up-conversion noise present, and with the drive off and no up-conversion noise present. We expect to observe different means in the underlying distributions, as shown in Figure 2.4.

Thus, the analysis of confidence and uncertainty in our measured results reduces to the standard analysis of whether the two sets of data are unlikely to arise from the same underlying probability distribution. Then, appropriate statistical methods, such as the Student's t-test, can be used to determine if a statistically significant difference in the means of the two sets of results is present or to derive confidence intervals on the upper bound of the difference in the means consistent with our observations. This manner of statistical validation would provide a strong argument for the observation of up-conversion noise.

In practice, the functional dependence of the up-conversion noise on the applied force is unlikely to take the simple linear form we used above. The model can be generalized by writing out more terms in the Taylor expansion of the χ and θ functions, and working out the corresponding periodic fluctuations expected in the displacement power time series. Thus, by examining different harmonics of the drive frequency, we can potentially infer the form of χ and/or θ and the micro-mechanical phenomena they arise from. Without going into the details of those computations, it suffices to say that the analysis will be carried out looking at various frequency components of the up-conversion noise amplitude: at the drive frequency, the second and the fourth harmonics. Additionally, we will allow for modulation both in-phase and in quadrature with respect to the drive.

Finally, while demodulation techniques such as we have described here are useful for discriminating periodic signals from stationary background, there are additional systematic effects from background noise that are also modulated by the common-mode drive and thereby not easily distinguishable from true up-conversion noise. Examples include Barkhausen noise of the magnets used in the electromagnetic actuators driving the test cantilevers, or modulation of the power detected at the photodiode induced by misalignments of the end mirrors due to the common-mode drive. These effects are ideally minimized via careful experimental design and construction, and their contributions to the demodulated signals quantified and accounted for.

2.3 THE INITIAL PROTOTYPE OF THE MEASUREMENT SYSTEM

The initial prototype for this experiment consisted of a Michelson interferometer, with end mirrors attached to the bottom of load masses clamped to the tips of small test cantilevers that were used in Advanced LIGO prototype suspensions. The cantilevers were clamped to a single tall post, and in turn attached to an optical board, where the horizontal Michelson interferometer was mounted. The need to measure vertical motion of the test cantilevers while the interferometer was arranged horizontally introduced additional complexity to the system and reduced its overall rigidity.

The load was rigidly clamped to the tip of the test cantilever, as can be seen in Figure 2.5. This approach had several drawbacks. First of all, any vertical motion of the cantilever tip coupled directly to a tilt of the load mass and of the Michelson end mirror. In turn, this misalignment of the Michelson was a limiting factor for the maximum amplitude of the common mode displacement we could exert. Secondly, this rigid clamp also coupled all of the cantilever transverse and torsional modes to angular motion of the mirror, introducing additional complexity to the actuation and control of the system.

The apparatus was housed in a vacuum chamber to mitigate acoustic noise, and mounted on a stack of two plates standing on rubber springs to reduce seismic noise. Outside of the chamber, a free-running polarized HeNe laser was coupled into a single mode, polarization maintaining, fiber optic cable, which then was fed through to the interior of the chamber.

While the prototype reached a sensitivity on the order of 10^{-14} m/ $\sqrt{\text{Hz}}$ above 400 Hz, the sensitivity at lower frequencies was greatly limited by poor seismic isolation, as can be seen in Figure 2.6. This, in turn, was limited by the available space inside of the available vacuum chamber.

Mitigating these issues became the main consideration when designing the second iteration of the experiment. Specifically, we decided to suspend the cantilevers load masses with steel wires to reduce the coupling of higher order vibrational modes of the cantilevers, and to construct a two stage pendulum seismic isolation system to attempt to reach a sensitivity of 10^{-15} m/ $\sqrt{\text{Hz}}$ at 10 Hz. This figure is motivated by the sensitivity at which a null result would suggest that up-conversion noise would not be a limiting noise source for Advanced LIGO.

Concretely, with the first iteration of the apparatus, we set an upper limit on the possible






Figure 2.5: A loaded maraging steel test blade, showing an electromagnetic actuator. One end mirror of the Michelson interferometer is mounted on the bottom of the loading mass.

up-conversion noise when the cantilevers were subjected to a $1\ \mu\text{m}$ common mode displacement of $3 \times 10^{-15}\ \text{m}/\sqrt{\text{Hz}}$ in the band from 450–500 Hz [40]. However, without a clear observation or verified physical model, the frequency dependence of the noise is unknown — though perhaps reasonably lying between f^{-3} and f^{-1} — making it difficult to extrapolate this upper limit to frequencies relevant for Advanced LIGO. Depending on the noise model used, the extrapolated noise at 10 Hz varies dramatically. From our prototype’s upper limit, we computed an upper limit of the amplitude spectral density of up-conversion noise in the Advanced LIGO upper intermediate mass (UIM, see Figure 2.1) cantilever tips propagating to the gravitational wave strain readout anywhere from 2.4×10^{-21} – $4 \times 10^{-18}\ \text{m}/\sqrt{\text{Hz}}$, depending on the spectral profile of up-conversion noise, where the Advanced LIGO design sensitivity at 10 Hz is approximately $8 \times 10^{-19}\ \text{m}/\sqrt{\text{Hz}}$ [40] (see Figure 2.7). Therefore the results obtained with the prototype were not good enough to rule out up-conversion noise as an important factor in Advanced LIGO’s sensitivity.



2.4 THE IMPROVED MEASUREMENT SYSTEM

The limitations found during the operation of the first measurement system prompted us to design an upgraded, more sensitive measurement system. The scientific goal of the improved system is to reach a displacement sensitivity of the differential motion of the tip of the two test cantilevers of the order of $10^{-15} \text{ m} / \sqrt{\text{Hz}}$, at frequencies of 10–20 Hz and above, thus improving by many orders of magnitude our capability to detect up-conversion noise in the low frequency region.

This section describes the main features of the improved system: passive suspension of the optical board to achieve better isolation from seismic ground motion; use of a near infrared Nd:YAG source to reduce the laser technical noises; and an improved design of the test cantilever load, clamp and displacement readout.

Figure 2.8 shows a rendering of the improved instrument. The optical board that holds the Michelson interferometer hangs vertically inside the support structure. The bread-board is suspended by two stages of vertical and horizontal isolation. Its motion is sensed and controlled using six integrated shadow sensor and electromagnetic coil actuators. More details on the seismic isolation system are given in Section 2.4.3. The entire system is housed inside a vacuum chamber, to reduce contamination of the optics, noise due to air fluctuations, and acoustic disturbances.

2.4.1 OPTICAL SYSTEM

Figure 2.9 shows a schematic of the optical layout. As described above, the readout system is a Michelson interferometer. The two end mirrors of the Michelson interferometer must be horizontal, since they should measure the vertical displacement of the test cantilevers. For this reason, the entire optical system is mounted on a vertical $40 \times 45 \text{ cm}$ board. This allows us to have a much more rigid structure, compared to the first prototype system. The light source is a 1064 nm wavelength Nd:YAG NPRO (Non-Planar Ring Oscillator) laser, which delivers a typical power of 10–20 mW into the interferometer. This power level is enough to reach a shot noise limited sensitivity better than $10^{-15} \text{ m} / \sqrt{\text{Hz}}$ over the entire band of interest. The laser is not actively stabilized in either intensity or frequency at this time. The laser beam enters through a view port (not shown) on the bottom left side of the board and is steered into the beam splitter by two adjustable mirrors. The two arms of the Michelson interferometer are folded in such a way that the beam is almost vertically incident on the two interferometer end mirrors. They are mounted on two 2.2 kg blocks that are suspended with wires from the two test




Figure 2.8: Rendering of the improved measurement apparatus (Section 2.4), showing the suspended optical breadboard (Section 2.4.1), the seismic isolation system (Section 2.4.3) and the support structure.

cantilevers. Lateral motion of the two blocks is also sensed and mitigated with the same kind of integrated sensors-actuators that are used for the main optical board. Both test cantilevers are clamped to the same support visible in the top center of the board. On the two cantilever tips there are two additional displacement sensors and actuators (not shown in the figure) that are used both to maintain the correct half-fringe operating point of the interferometer and to apply the common mode low frequency drive that would excite up-conversion noise. The symmetric and anti-symmetric beams recombining at the beam splitter are picked up by two additional steering mirrors and sent to two photodiodes.

The typical free running frequency noise of an NPRO system like the one we use [41] is




Figure 2.9: Simplified optical scheme of the Michelson interferometer. Only the main beams and optical components are shown: reflections from the secondary surfaces and beam dumps are not drawn for simplicity. Also, actuators and displacement sensors have been removed.

on the order of $100 \text{ Hz} / \sqrt{\text{Hz}}$ at 100 Hz and decreases as f^{-1} . The coupling depends on the Michelson arm length difference, as discussed above. So, to reach our design sensitivity at 10 Hz, the length difference of the arms must be smaller than 0.3 mm. This level of accuracy is not easily obtainable in the installation phase of the optical system. For this reason, one of the two folding mirrors in the interferometer (Figure 2.9) is mounted on a linear motorized translation stage. It is then possible to add an external perturbation on the laser frequency and directly measure the coupling to the Michelson displacement signal. This can be converted into a length difference, via equation 2.8 above, that can

be corrected using the translation stage. This procedure allows us to achieve the needed length balancing. The other folding mirror is mounted on a motorized angular stage, to allow us to fine tune the interferometer alignment in vacuum.

2.4.2 IMPROVED TEST CANTILEVER ASSEMBLY

The two test cantilevers are pre-curved in such a way that when they are loaded at about 50 % of their yield stress (corresponding to 2.2 kg in our case), they are flat. The transverse profile of the cantilever is triangular: in this way the initial curvature is constant along the entire length of the cantilever, and, moreover, the static stress due to the load is constant along the cantilever, except of course close to the clamp, where there is some localized increase of stress.

The load mass is attached to the cantilever through a single steel wire. In this way we obtain a very high decoupling of any torsional and angular motion of the cantilever tip from angular motions of the Michelson mirror, which is rigidly attached to the bottom of the load mass. Indeed the load mass is isolated from lateral motion of the cantilever tip by a pendulum, and from any angular motion by the stiffness of the wire itself, which can be made very small. Moreover, the wire is clamped to the cantilever with two small steel blocks, held together with bolts. This is a scaled down version of the clamp used in the Advanced LIGO system, and it provides a clean solution that avoids friction and additional stress. Additionally, it serves the purpose of making the test system as similar as possible to the system used in the gravitational wave detectors.

2.4.3 SEISMIC ISOLATION

The dominant limitation to the sensitivity of the first version of the measurement system was seismic noise at frequencies below a few hundred Hz. Indeed, the ground motion in a typical urban ground location can be many orders of magnitude larger than our target sensitivity. The measured motion of an optical table in our lab showed a displacement noise of the order $10^{-8} \text{ m} / \sqrt{\text{Hz}}$ at 10 Hz, decreasing with frequency roughly like f^{-3} . The most important degree of freedom in our system is the vertical one, since this corresponds to the direction of the Michelson interferometer measurement. Ideally, if the optical system were infinitely rigid and the two test cantilevers were exactly equal, any vertical motion of the optical breadboard would result in a common mode variation of the interferometer arm lengths. Thus, since a Michelson interferometer has virtually infinite common mode rejection, it should not be affected by seismic motion of the ground.

However, there is a limit to the level the two cantilevers can be made equal: in particular, differences in the material, machining, and clamping can result in a mismatch of the resonant frequency and of the distance from the clamp to the wire suspension point. A trade-off is necessary between the requirements on the cantilever equality and the performance of the seismic isolation system: a worse matching of resonant frequency or distance would require increased performance on the suspension system. It can be shown using a simple elastic model of the two cantilevers that the residual coupling of common vertical motion x_{comm} to differential displacement x_{diff} of the two cantilever tips is given by:

$$\frac{x_{\text{diff}}}{x_{\text{comm}}} \approx \left(\frac{f_0}{f}\right)^2 \left[2\frac{\delta f_0}{f_0} + \frac{\delta L}{L}\right] \quad (2.14)$$

Here, f is the measurement frequency, f_0 is the cantilever mean resonant frequency, δf_0 is the difference between the two resonant frequencies, L is the mean of the cantilever's length from the clamp to the wire attachment point, and δL is the length mismatch. The two expressions above are correct for frequencies larger than f_0 (about 2 Hz) and smaller than the first higher order resonance of the loaded cantilever, about 150 Hz.

A difference in the two resonant frequencies of about 5 mHz, obtained experimentally in the first prototype, and a difference in the two lengths of 0.5 mm, well within machining tolerances, provide us with a common mode rejection factor of about 6000. So, to reach the desired displacement sensitivity at 10 Hz, the suspension system must provide an additional factor of 2000 of vertical isolation at 10 Hz. This is achievable using two cascaded stages with characteristic frequencies close to 2 Hz. Figure 2.10 shows a simplified schematic of the mechanical system. Each stage is composed of maraging steel cantilevers, roughly 30 cm long, 7 cm wide and 2 mm thick. Four cantilevers suspend the intermediate stage from a support structure with steel wires, and two additional cantilevers support the optical breadboard from the intermediate stage, with another two wires attached to the sides of the board, above its center of mass. Each cantilever supports a load of about 10 kg, which corresponds to about 50 % of their yield stress. Both the optical board and the intermediate stage have a mass of about 20 kg. The intermediate stage includes a stack of rubber to provide some passive damping of the suspension resonant modes.

2.4.4 CURRENT SENSITIVITY AND NOISE SOURCES

The design described in the previous sections provides a theoretical sensitivity to displacement noise which is limited by shot noise at all frequencies above about 20 Hz at a level better than $10^{-15} \text{ m}/\sqrt{\text{Hz}}$. Figure 2.11 shows the measured sensitivity of the system in the




Figure 2.10: A simplified schematic of the seismic isolation system, highlighting the key components and the two stages of vertical and horizontal isolation.

present configuration. The same figure shows also the best sensitivity achieved with the first prototype of the experiment, as described in Section 2.3. The vast improvement at low frequencies is very apparent. The additional traces in Figure 2.11 show the projected contribution of various technical noises to the measured sensitivity. The sum of all those noise is capable of explaining almost all of the measured displacement noise. However, the contribution of seismic noise is much larger than what was foreseen in the design of the seismic isolation system, both at frequencies below 40 Hz and at frequencies above




Figure 2.11: Typical sensitivity of the measurement system in the present configuration (solid black trace) compared with the best obtained with the first prototype (dotted red line). The other traces show the contribution of various noise sources to the total displacement noise: actuation noise and scattered light noises are described in Section 2.4.4; electronic noise refers to the sum of photodiode dark noise and analog-to-digital conversion noise; laser intensity noise and shot noise are discussed in Section 2.2.1.

200 Hz.

The main coupling path in the low frequency region has been identified as the following: since the optical board and the intermediate stage of the seismic isolation are suspended by multiple wires and cantilever springs, any difference in the stiffness of the springs causes a direct coupling of vertical motion of the suspension point to angular motion of the suspended body. In particular, the critical angular degree of freedom for our measurement is the roll motion of the optical board (i.e. rotation about an axis perpendicular to the board surface), since any motion in this degree of freedom will create a differential displacement of the two test cantilevers with respect to the Michelson beam splitter. An imbalance of a few percent in the cantilever stiffnesses — well within manufacturing tolerances — is enough to explain this increased coupling of seismic noise at low frequency. We have designed a modification of the seismic isolation system, consisting of an additional stage to be added before the optical board, that will decouple vertical motion from angular motion. At higher frequencies, the increased coupling of seismic

noise is also due to mismatched responses of the suspension cantilevers, this time in their internal resonant modes. This issue will also be mitigated by the addition of the angular decoupling stage as described above.

As discussed in Section 2.2, the direct displacement sensitivity is not the ultimate limit to our measurement system, since the demodulation technique can detect periodic non-stationary noise below the sensing noise, provided that the latter is stationary. Therefore particular care is necessary for all sources of non-stationary noise, especially those that might be modulated by the common mode motion of the two cantilevers. Referring again to Figure 2.11, two of the noise sources listed there are particularly problematic. Scattered light is intrinsically non-stationary, since the amplitude and maximum frequency of this noise source depends on the motion of the scattering element [42]. Scattered light has been mitigated with a careful placement of black glass absorbing baffles and beam dumps. All spurious beams from the anti-reflection coated surfaces are intercepted and dumped. This improvement will be sufficient to reduce scattered light below the target sensitivity. In addition the increased seismic isolation will also help reducing scattered light. Indeed, scattered light up-conversion is highly non linear [42]: residual motion at few Hz will be the dominant contributor to scattered light phase noise, while the slower ≈ 100 mHz motion that we will introduce to periodically stress the cantilevers results in a negligible contribution. The second potentially problematic non-stationary noise source can be traced to the actuation chain used to apply force on the two test cantilevers. In particular, the digital-to-analog converters (DAC) are known to exhibit a significant amount of harmonic distortion. The low frequency common mode drive is up-converted in frequency by the DAC and results in a non stationary noise at the level of the measured sensitivity. This issue is being tackled with an improvement of the control electronics.

Finally, Barkhausen noise in the magnets used for the vertical actuation of the test cantilevers can result in crackling noise like signals. To reduce the impact of Barkhausen noise we are using SmCo magnets which have much lower noise than the more common NdFeB magnets [39]. The estimated contribution of this noise source is well below our present sensitivity.

2.5 DISCUSSION AND OUTLOOK

This chapter presented an instrument designed for the study of the mechanical up-conversion phenomenon in metals. Two key points make the approach presented here different from previous studies. First of all, given the authors' involvement in the gravitational wave observatory Advanced LIGO, this system will study the behavior of metals in the elastic regime, far from the yield stress that would introduce plastic deformations. As already pointed out, to the best of our knowledge there has been no experimental investigations of this kind in this regime. Secondly, since we expect the up-conversion events, if present, to be of very small amplitude, our intended measurement is not the detection of the single events, but rather the statistical properties of the up-conversion noise that arise as the incoherent sum of all the events. In particular, we are interested in the dependence of the noise power on the low frequency external disturbance the metal is subjected to.

At the time of writing, the experimental apparatus has been constructed and commissioned. Although it has not yet reached the design sensitivity, characterization of the background noise was successful. With a measured sensitivity level of $\approx 10^{-14}$ m/ $\sqrt{\text{Hz}}$ at 50–100 Hz, the experiment has already reached a sensitivity level which is significantly better than the first prototype of the instrument. However, several limitations of the present setup have been already identified and are being tackled with small scale, short term modifications of the seismic isolation system and of the control electronics. This will allow meaningful upper limits to be set for the contribution of up-conversion noise to the Advanced LIGO detectors' sensitivities, and may possibly yield a direct detection of up-conversion noise in metals still operating within the elastic regime.

Clearly, one important point that has to be addressed is how to scale the results from the small test cantilevers used in our experiment to the much larger blades installed in the Advanced LIGO suspension. The derivation of this scaling is made more uncertain by the lack of a microscopical model of up-conversion noise in our regime. Since the time of this work, a scaling model for the cantilevers has been published in VAJENTE [43].

CHAPTER 3

THE CALTECH 40M PROTOTYPE INTERFEROMETER

3.1 INTRODUCTION

The Caltech 40m Prototype Interferometer is a controls prototype for the Advanced LIGO observatories — a test bed for novel techniques and technologies where new ideas can be quickly iterated without disrupting an active observatory. To this end, the 40m has been designed and commissioned in a configuration generally analogous to the aLIGO instruments to ensure the applicability of the research done there.

As mentioned in Chapter 1, the optical portion of the aLIGO observatories consists of multiple coupled resonant optical cavities formed by suspended mirrors. As such, all of the appropriate resonance conditions must be met in order for the interferometer to operate in its desired configuration (i.e. the distances between the mirrors forming the arm cavities must be an integer number of wavelengths of the laser light present in the cavities). However, despite the substantial isolation of the mirrors from ground motion from the mirror suspensions and active seismic isolation platforms, the residual motion of the mirrors is too large for the cavities to passively remain resonant. Thus, active feedback control of the various optical cavity lengths plays a critical role in operating the aLIGO interferometers. Furthermore, it is vanishingly unlikely that, upon attempting to bring the instrument under control, all of the resonance conditions will be simultaneously met and one could simply engage global control and be done with it. Instead, we desire to plan, test, and execute a deterministic lock acquisition procedure that will robustly and repeatedly bring the interferometer to its desired operation point from an uncontrolled

state. This controls problem has been the primary focus of the 40m prototype, and has guided the design decisions in its construction and commissioning.

Certain aspects of the aLIGO design are not strictly necessary or practical to this end. For instance, the multiple pendulum design for the aLIGO optics was elided in favor of simple single pendulum suspensions. At the prototype scale, it is also much more practical and safe to use a commercial laser light source instead of a custom high power laser system. Thus, since the available laser power is much less than in the aLIGO system, the masses of the mirrors themselves must be reduced in order to maintain analogous radiation pressure dynamics. A summary and comparison of some important parameters of the aLIGO and 40m interferometers is shown in Table 3.1.

This chapter will describe the commissioning work done to realize the fully resonant Dual Resonant Sideband Extraction (DRSE) scheme used in aLIGO at the 40m in 2015. In particular, it will describe why and how the locking procedure differed from its original conception and the procedure used at the aLIGO observatories.

In addition to the material in this chapter, substantial amount of information regarding the current optical layout, seismic isolation, input laser systems, and optic suspensions at the 40m prototype can be found in WARD [10] and DRIGGERS [44].

Table 3.1: Comparison of aLIGO and 40m Prototype Parameters. [4, 5, 12, 44]

	aLIGO	40m
Arm cavity Length	3994.5 m	37.80 m
Mass of Test Mass	40 kg	250 g
Beam radius on ETM	6.2 cm	5.1 mm
Base modulation Freq.	≈ 9 MHz	≈ 11 MHz
Input Laser Power (Incident on PRM)		
Current Levels	22 W	1 W
Future	125 W	5–10 W
$\lambda = 1064$ nm		
ITM Transmission	1.4 %	1.4 %
ETM Transmission	3.6 ppm	15 ppm
PRM Transmission	3 %	5.6 %
SRM Transmission	32 %	9.9 %
Arm cavity Finesse	≈ 440	≈ 420
Round-trip loss	85–100 ppm	250–500 ppm
$\lambda = 532$ nm		
ITM Transmission	1–2.5 %	1.1 %
ETM Transmission	3–15 %	5.6 %
Arm cavity Finesse	13	108
Observed Power Recycling Gain	≈ 40	≈ 8
Signal Extraction Pole	≈ 380 Hz	167 kHz

3.2 LENGTH SENSING

A representative sketch of the 40m prototype optical layout and placement of optical sensors is shown in Figure 3.1 (from DRIGGERS [44]). While some of the optical sensors are used for monitoring the DC optical power or angular control, the majority are interferometric length sensors that operate via optical heterodyne effects. In the case of the green light ALS signals, differential changes in frequencies of the PSL and auxiliary lasers change the optical beat frequency on a broadband photodiode, which is tracked via a delay line frequency discriminator. The infrared sensors are resonant RF photodiodes (RFPDs) tuned to sense the beat between phase modulation sidebands, which is thereafter demodulated to infer the length fluctuations.

The ALS signals are identified by ALSX/ALSY for the beats between the green light derived from the auxiliary lasers at the X/Y end and the PSL laser. Common and differential arm length fluctuations can then be digitally calculated as the sum and difference of these. In contrast, aLIGO directly beats the X and Y green light for a differential signal, and uses a beat between one auxiliary laser and the PSL for the common signal. Given aLIGO's significant seismic isolation, the laser frequency fluctuations in the single ALS signal generally dominates the cavity length fluctuations at all frequencies.

The IR RFPD signals are identified by the interferometer port their incident field is sampled from and the approximate frequency they are resonant for, which corresponds to the subsequent demodulation frequency and the frequency difference between the field products the RFPD is intended to sense. The demodulation produces signals in two quadratures, referred to as the in-phase (I) and quadrature (Q) signals; these two signals will generally have different responses to the length fluctuations of the interferometer, and can be independently used to control length degrees of freedom.

Section 1.3.3 contains more detail about the length sensing design of the interferometer. Table 3.2 shows the sensors used for control of the interferometer lengths during the lock acquisition process, which will be detailed in Section 3.4. Generally, the signals used in the initial phase (ALS and $3f$) are less precise, but less sensitive to the highly non-stationary coupling between the recycling and arm cavities before the lock is stable. The signals used in the final state have much smaller dynamic range, but much better sensitivity.

It may seem counter intuitive that the REFL_{II} I signal may be used for the simultaneous control of two degrees of freedom. REFL_{II} I has strong contributions from both CARM and PRCL, but the final CARM control bandwidth is around 30 kHz which allows for substantial suppression of the CARM signal in the hundreds of Hz. However,

this control loop will not suppress the residual fluctuations due to the PRCL signal, so the remaining signal may be used to control the SRC length as well.

It may also be noted that a linear combination of signals is used for SRCL control. While REFL₅₅ has greater sensitivity to SRCL fluctuations than REFL₁₁, since only the 55 MHz sideband resonates in the SRC, it still has a strong contribution from the resonance of the 55 MHz sideband in the PRC; it is undesirable to include this PRCL contribution to actuation on the SRC length. So we create a linear combination of REFL₁₁ I and REFL₅₅ Q that is tuned to nullify the presence of a PRCL signal, so that the remaining content can be used to control the SRC length without excess actuation noise.

Table 3.2: Sensors used for Interferometer Length Controls

Length	Initial	Final
CARM	ALSX + ALSY	REFL ₁₁ I
DARM	ALSX - ALSY	AS ₅₅ Q
PRCL	REFL ₃₃ I	REFL ₁₁ I
SRCL	REFL ₁₆₅ I - α REFL ₃₃ I	REFL ₅₅ I - α REFL ₁₁ I
MICH	REFL ₁₆₅ Q	REFL ₅₅ Q




Figure 3.1: Optical layout and position of optical sensors used in the length sensing and control of the 40m prototype interferometer. Not to scale. f_1 and f_2 refer to the phase modulation sidebands imposed onto the main input laser at ≈ 11 MHz and ≈ 55 MHz. Adapted from DRIGGERS [44].

3.3 PERFORMANCE UPGRADES

3.3.1 DELAY-LINE FREQUENCY DISCRIMINATOR FOR ALS

The performance of the ALS subsystem (described in Section 1.3.4) is of particular importance for the lock acquisition of the interferometer. While it is not necessary for ALS to be as sensitive as the IR optical signals derived from the main laser, its sensitivity must be good enough that the error signal used for feedback control can be reliably switched from ALS to IR when the appropriate conditions are reached. Practically speaking, this means that the actuation noise imprinted on the cavity lengths due to the ALS sensing noise must not be so high that the optical plant is changing substantially while the hand-off between the error signals is made.

The 40m prototype uses two delay-line frequency discriminators (DFD) to track the fluctuations of the RF beat frequencies between the green frequency doubled light from the auxiliary laser and the main laser. This differs from the PLL based system described in STALEY et al. [12] due to the fact that the absolute frequency stability requirements at the 40m are less stringent, given the 100× shorter arm lengths.

The essential concept of the DFD is that mixing an RF signal with a delayed version of itself creates a voltage signal that can be used to infer the RF signal's frequency fluctuations. Figure A.1 shows a schematic of the hardware employed. In addition, the 40m employs a digital phase tracker, which extends the linear range of the DFD signals.

The length of the delay line in a DFD is a key determinant of its performance, as the delay itself sets the conversion factor of the RF frequency fluctuations to a measurable audio frequency signal. However, one cannot increase the length without limit as energy loss in the cable will decrease the delayed signal's amplitude and thereby also reduce the conversion. An analysis of these effects will show that the maximum conversion in a DFD is set by choosing the length of delay line that results in 8.7 dB of loss in the specific cabling being used (see Appendix A).

Naturally, it is important to shield the delay line cabling from external EMI and physical vibrations that may get picked up and contaminate the readout. LMR-195 coaxial cable was used for our DFDs, and it was empirically determined that 50 m of cable would result in the best sensitivity. A rack-mounted enclosure was built to house the two bundles of cable that were used as the delay lines for the two ALS DFDs. In addition, the RF amplification of the optical beat signals from the broadband photodiodes was simplified and aLIGO RF demodulator units were used in place of the previous custom built circuits.




Figure 3.2: Comparison of the out-of-loop displacement fluctuations in a single 40m Fabry-Pérot cavity when under active stabilization with the ALS subsystem with the results in IZUMI et al. [23], as a measure of the control noise injected by use of the ALS signal. The dashed line indicates the cumulative RMS of the spectrum as integrated from high to low frequencies. The out-of-loop sensor is the IR PDH signal.

These changes resulted in a significant increase in the sensitivity of the ALS subsystem at the 40m prototype, with a corresponding decrease in the control noise introduced when using the ALS signal as a length feedback signal (see Figure 3.2).

3.3.2 SEISMIC FEED-FORWARD FOR PRC ANGULAR MOTION

In addition, angular stability in the recycling cavities is a particular issue at the 40m prototype, as we do not employ the same degree of passive seismic isolation or active interferometric angular control as aLIGO. Excess angular motion can lead to significant differential motion of the cavity axes of, for instance, the power recycling cavity and the arm cavities, causing constant changes in their coupling and a difficult controls problem.

While angular motion of the resonant field inside of the PRC can be seen on a QPD at the POP port, the observed motion is dominated by angular fluctuations of the arm cavities when the interferometer is fully resonating. However, if one arranges the interferometer

such that light is resonant only in the PRC (by a gross misalignment of the ETMs, for instance), the POP QPD signal is entirely determined by angular motion in the PRC. Since the angular motion is seismically driven, we can simultaneously record the QPD signal along with the output of a nearby seismometer to characterize the relationship between the two, and train a feed-forward stabilization filter. This is done by the Wiener filtering technique that will be discussed in detail in Section 5.2.2, and is analogous to the feed-forward reduction of seismic noise in DEROSA et al. [45]. More details about the training of these specific filters can be found in DRIGGERS [44].

Since the mechanical relationship between seismic motion and the induced angular motion of the PRC cavity axis does not change when the interferometer is fully resonant, this feed-forward subtraction can be continuously run once the PRC is resonant and under stable feedback control. Figure 3.3 shows the effect of the feed-forward filters on the angular degrees of freedom witnessed by the POP QPD and — as an important figure of merit — the reduction of fluctuation in the PRC intra-cavity power, which demonstrates that the cavity axis is truly stabilized and will have more consistent coupling with the arm cavities.




Figure 3.3: Effect of active feed-forward control of PRC angular fluctuations, leading to more stable circulating power. The RMS fluctuations in the angular motion of the PRC mode pick-off at the POP QPD are reduced by a factor of 2–3, leading to a decrease in the RMS relative intensity noise (RIN) of about the same.

3.4 LOCK ACQUISITION PROCEDURE

Although the 40m prototype was designed to be as analogous to the aLIGO detectors as was feasible to serve its function as a controls prototype, the inevitable differences have a significant effect on the operation of the instrument. The most obvious difference is the interferometer baseline and arm cavity length; since the prototype's arms are ≈ 100 times shorter, the optical pole frequencies are ≈ 100 times higher. This is especially important for the common mode optical behavior in the interferometer. At the start of the CARM offset reduction sequence (see Section 1.3.4) the CARM controller need only to account for the single arm cavity pole at 532 nm, while in the final locked state the optical plant instead exhibits an optical pole determined by the coupling of the resonant IR modes of the arm cavity and power recycling cavity.

For aLIGO's arm cavities, the single arm optical pole frequency is ≈ 80 Hz while the coupled cavity creates a pole at ≈ 0.5 Hz[4], which are both solidly within the controllable band of the CDS system (see Section 1.3.3). The prototype's corresponding pole frequencies, however, are 8.8 kHz and 80 Hz; it is much more challenging to design a digital control that is stable over the course of the CARM offset reduction sequence as the optical plant exhibits this change.

The locking scheme employed at the aLIGO observatories uses the ALS error signal to approach IR resonance in the arms until there is enough circulating power to enable the use of the IR power levels transmitted through the arm cavities as an intermediate error signal for the CARM control loop, which is then used until the PDH-style reflection locking signal is available.

Simulated optical responses of the transmitted light signal at the 40m prototype at different points as a function of CARM offset are shown in Figure 3.4, which shows how the dynamics of this signal are affected by the moving IR coupled cavity pole. It can be seen how the coupled cavity pole rapidly transitions through the CARM control unity gain frequency at the CARM offsets where we must cede control from ALS to IR light signals, which led to an inevitable loop instability and loss of lock.

However, it became apparent that the lock acquisition procedure could be simplified greatly since the sensitivity of the 40m ALS subsystem, as compared to the final common optical mode linewidth, was lower than in aLIGO. The aLIGO ETMs unfortunately have errant dichroic coating properties at 532 nm, which led to an effective finesse of about 13 instead of the design value of 100 for the resonant 532 nm light used for the ALS subsystem [12]. The 40m coatings did not suffer from this problem.




Figure 3.4: Simulated response of IR transmitted light signals to CARM fluctuations at the 40m prototype

Specifically, as can be seen in Figure 3.2, the total RMS sensing noise of ALS is ≈ 300 Hz, which is only a few times larger than the common mode linewidth of ≈ 100 Hz. Thus, it could be possible to transition CARM control from ALS directly to the RF REFL signal at zero CARM offset, counting on the fact that the REFL signal is only of appreciable magnitude when the cavities are moving through coincident resonance (this can be referred to as “self-locking”)[46]. Indeed, this scheme proved to be the key to bringing the 40m prototype to its fully operational state.

Specifically, a successful lock acquisition sequence proceeds as follows:

1. With the PRM and SRM misaligned, the arm cavities are individually brought under stable control via the POX/POY signals.
2. The auxiliary green lasers are PDH locked to the arm cavities, and have their zero point set to coincide with the ongoing arm resonances.
3. The control of the arm cavity lengths are transitioned to the ALS signals with a control bandwidth of 120 Hz

4. The cavities are moved away from IR resonance (10 nm) through the introduction of an offset in the CARM error point.
5. The PRM and SRM are realigned, and the DRMI is locked using the $3f$ REFL signals. The DRMI loop triggering is adjusted to include the IR power transmitted through the arms.
6. The CARM offset is ramped to zero. At this point, the interferometer is rapidly passing back and forth through full resonance as the actuation noise due to the ALS sensing noise is larger than the linewidth.
7. A small amount of the REFL_I signal is blended in to the CARM error point, and then an integrator is activated on this portion in order to give control authority to the REFL signal at DC. The same is done for AS₅₅ Q into the DARM error point.
8. The analog AO path of the CARM servo is activated with the REFL_I signal, in order to prepare the blended digital and analog loop shape that will be stable up to 30 kHz of control bandwidth, with a crossover frequency of 100 Hz.
9. The REFL_I contribution to the CARM error signal is increased, and the overall CARM gain is increased to reach a control bandwidth of 1 kHz, at which point the cavity motion is reduced to stay within a linewidth and the arm cavities are resonating, though with significant power fluctuation.
10. DC angular control is activated for all four test masses, through optical levers and transmission QPDs. The alignment may be manually adjusted here to maximize the circulating power.
11. In turn, the ALS contributions to the CARM and DARM error signals are AC coupled, and then ramped to zero entirely. At this point, the ALS subsystem is completely removed from interferometer control and the circulating power is much more stable.
12. The CARM loop gain is increased to a control bandwidth of 10 kHz and low frequency servo boosts activated, which brings the CARM control to its final operating state.

See Figure 3.6 for the open loop gain profile of the CARM feedback control during different stages of this sequence and Figure 3.7 for a time series depiction of the power at




Figure 3.5: Contributions of the ALS and PDH signals to the total OLG just before engaging the AO path. The low frequency authority granted to the PDH signal helps keep the interferometer near resonance before the actuation noise is low enough to remain there.

different ports over time. The contributions to the CARM and DARM error signals from the ALS and PDH sensors during the sequence is represented in Figure 3.8.

After this sequence, the interferometer is fully locked and one can proceed to take steps to improve the overall sensitivity, such as activating the photodiode signal whitening filters and transitioning DRMI control to the REFL $1f$ signals which are now stabilized through the CARM servo.

Step 5 was found to be necessary due to the observation of strong fluctuations in the demodulated intra-cavity signals use to monitor the state of the DRMI (POP₂₂ and POP₁₁₀), that did not seem to correspond to the DRMI moving away from coincident resonance. Including the arm transmissions in the triggering is justified since it is not possible to reach arm powers greater than in the single arm case without a resonating DRMI, and so the triggering threshold was set appropriately.






Figure 3.6: CARM open loop gain at different stages of the lock acquisition process. The initial loop shape corresponds to the use of the ALS error signal with digital actuation on the input mode cleaner length. The high frequency response is then modified by engaging the analog AO path, actuating directly on the laser frequency using a blended ALS and PDH signal. The final high-bandwidth loop profile uses only the PDH error signal. The model of the final loop shape compares favorably with a measurement taken during a lock; this measurement only spans 10–100 kHz as the excitation required to achieve reasonable SNR in regions with substantial loop gain disrupts the lock.



Mar 29 07:22:55 2016 UTC

Figure 3.7: Evolution of the power circulating in the PRC and X arm cavity during lock acquisition of the 40m interferometer. Numbers shown in red correspond to steps in the locking sequence detailed in Section 3.4. The observed powers indicate a power recycling gain factor of ≈ 8 and an interferometer visibility of $\approx 70\%$



Mar 29 07:22:55 2016 UTC

Figure 3.8: Contributions to the CARM and DARM error points from the ALS and PDH derived signals during a lock acquisition of the 40m interferometer. Initially, the ALS signal is roughly calibrated to meters, but the frequency dependent blending does not allow for a direct interpretation of displacement. Numbers shown in red correspond to steps in the locking sequence detailed in Section 3.4. The delay between the complete PDH hand-off for CARM and DARM is due to the fact that the DARM loop bandwidth was being manually adjusted for a time before completely handing off control to the AS55 Q signal.

3.5 ANGULAR DYNAMICS AND CONTROL

In order to robustly and repeatedly lock the interferometer, it is important that the suspended optics are well aligned such that the resonant spatial modes of the various cavities coincide with each other and the input beam, in order to achieve maximal resonant build up and to avoid the confounding effects of higher order spatial mode fields to the interferometric signals used to control the interferometer. Currently, a hierarchical dither alignment scheme is used to define the alignment state before attempting to lock the instrument.

First, the arm cavities are independently locked on IR resonance using the POX and POY PDH signals, with the PRM and SRM misaligned. All four test masses are dithered with independent excitations in the pitch and yaw degrees, and the cavity transmission and length error signals are demodulated at the excitation frequencies. The demodulated transmission signals are used to align the test masses to maximize the resonant power, while the demodulated length signals are used to steer the input beam to center the resonant mode on the mirrors' center of rotation. However, while full actuation of the beam entering the Y arm is available with two active mirrors between the IMC and PRM, the steering into the X arm is only controlled with the beam-splitter. Thus, it was chosen to leave the centering of the beam spot on ITMX untouched by the dither alignment, instead being set by its relative position to the other test masses.

Next, the ETMs are misaligned, the PRM is aligned, and the vertex is locked in the power recycled Michelson (PRMI) configuration. Here, the PRM and BS angles are dithered, and the DC light levels at the REFL and AS ports are demodulated. The PRM alignment is tuned to minimize the light reflected out of the interferometer, maximizing the power circulating in the PRC. The BS alignment is tuned to minimize the power at the AS port, ensuring good overlap of the recombining Michelson fields. Finally, the SRM is aligned, and the vertex is locked in the DRMI configuration. The SRM is dithered and its alignment is tuned to maximize the AS₁₁₀ Q signal at the AS port, indicating maximum buildup of the resonant f_2 field in the SRC.

Step 10 in the lock acquisition process is not ideal, but practically necessary in order to handle angular radiation pressure dynamics. Following Section 1.3.2, $\lambda_a \approx 1.726$ and $\kappa \approx 2.5 \times 10^{-3}$ N m/rad for the 40m prototype ETM suspensions [47]. Then, Equation 1.14 implies that the angular motion of the 40m arm cavities becomes unstable at powers higher than about 3 kW, which is a few times larger than the 1 kW circulating power observed in Figure 3.7. So, dynamical instability is not yet a problem at the 40m.

However, the softening of the total angular spring constant in combination with the lack of an interferometric signal for DC angular control means that any mirror's beam spot miscentering will lead to an exacerbated change in its equilibrium angular position, spoiling the alignment of the arm cavity and its spatial mode matching to the rest of the interferometer. Even if the rest of the dither alignment is perfectly accurate, the current configuration does not afford the centering of the ITMX beam spot, making DC angular radiation pressure torque effects unavoidable.

For this reason, all four test masses use a DC alignment reference that is set during the initial alignment procedure: the ITMs use their optical lever subsystems and the ETMs use a quadrant photodiode sampled from the transmitted IR beam. These sensors are inferior to the interferometric sensors used in aLIGO and would not be robust for long duty cycle locks and their operating points inherent all of the imperfections of the higher noise state present during the initial alignment. However, they are sufficient for hour scale locks and their set points can be adjusted by hand to a certain degree while monitoring the circulating arm powers and symmetry of the beam profile camera images at various ports of the interferometer.

3.6 INTERFEROMETER SENSITIVITY AND CHARACTERIZATION

3.6.1 RECYCLING GAIN

When the interferometer was first locked in the PRFPMI configuration (i.e. with the signal recycling mirror misaligned), it was noted that the recycling gain (the ratio of circulating power in the power recycling cavity to the input power incident on the PRM) observed in the interferometer was much lower than expected [44]. A comparison of the powers shown in Figure 3.7 indicates a gain of ≈ 8 instead of the design value of 40. This is due to excess loss in the interferometer in two places.

During commissioning, the current folding mirrors in the recycling cavities (PR₂, PR₃, SR₂, SR₃) were found to have an excess convexity that brought the cavities too close to geometric instability. The solution at the time was to reinstall them flipped, such that the resonant mode travels through the optic substrate to reflect off an effectively concave surface. The anti-reflective coating on the rear side of these mirrors then introduces a loss of 0.5 % each time the beam traverses the free-space to optic substrate interface. Since the light circulates in both directions at these mirrors, this introduces a 2 % loss into the PRC and SRC. At the time of writing, new dichroic optics are being procured with curvatures specified to ensure stability.

Additionally, around the time when the PRFPMI and DRFPMI locks were achieved, the optical losses of the interferometer arm cavities was measured, and found to be higher than desirable. By comparing the reflectivity of a single bounce from the ITM with the reflectivity of the arm cavity held on resonance, the round trip loss can be inferred [48].

The power recycling gain of the PRC is given by

$$G = \frac{T_{\text{PRM}}}{(1 - r_{\text{PRM}} r_{\text{Arm}}^-)^2} \quad (3.1)$$

where r_{Arm}^- is the mean of the resonant arm cavity reflectivities, each of which is given by

$$r_{\text{Arm}} = r_{\text{ETM}} - \frac{T_{\text{ETM}} r_{\text{ITM}}}{1 - r_{\text{ETM}} r_{\text{ITM}}} \quad (3.2)$$

Losses enter these equations by conservation of energy:

$$R + T + L = r^2 + t^2 + L = 1 \quad (3.3)$$

The measured losses were 240 ppm for the Y arm and 495 ppm for the X arm. Together with the effect of the flipped folding mirrors, these losses imply a recycling gain of 7.8, matching the observations closely.

Since this time, the HR surfaces of all four test masses has been cleaned with the “First Contact” optic cleaning solution, and recent measurements suggest improved round trip losses of ≈ 50 ppm and recycling gains as high as 23 [49, 50]. Together with the new folding mirror optics, this should enable operation at the intended recycling gain.

3.6.2 NOISE PERFORMANCE

Figure 3.9 shows an budgeting of the fundamental noise levels for the 40m, where by “fundamental” we mean fluctuations due to physical processes that are inherent to the physics of the instrument as constructed. The dominant contributions arise from the residual seismic motion that is filtered by the passive seismic isolation stacks and optic suspension, and the shot noise of the light at the AS port used for the strain readout. There is also a contribution from the thermal fluctuations of the optic suspension fibers at their violin mode.

However, as seen in Figure 3.10, the realized displacement sensitivity is many orders of magnitude worse than the fundamental limits. For most frequencies, the sensitivity level is generally explained by excess actuation noise due to distortion in the digital-to-analog converters (DACs) that drive the electromagnetic coil actuators used to control the suspended test masses’ position and angle. One strategy for reducing the influence of this kind of noise is to use digital whitening and analog de-whitening filters to attenuate the distortion that makes its way to the coil drivers. While the 40m has coil de-whitening electronics installed, their use has not yet been commissioned, and they remain bypassed. Any real increase in the sensitivity of the prototype will require that this be addressed.

At higher frequencies, it is suspected that coupling from the vertex length degrees of freedom is responsible for the excess noise, as the usually necessary feed-forward subtraction of these signals has not been employed at the time being.

3.6.3 IN-LOCK CHARACTERIZATION

It is useful to measure the optical response of the RFPDs to the individual length degrees of freedom in the fully resonant state, i.e. the *sensing matrix*. The measurements can be compared with the value predicted via calculation or simulation to make sure that




Figure 3.9: “Fundamental” noise sources that are relevant for the sensitivity of the 40m prototype interferometer. The “quantum noise” trace does not exhibit the same profile as that for aLIGO, as the RSE pole is at a much higher frequency, outside of the band resolvable with our digital control system. The “mirror thermal” trace includes the Brownian noise of the mirror substrate and coating, thermoelastic noise of the substrate, and thermo-optic noise of the coating.

all of the relevant effects are understood, and used to inform changes in the control or electronics design. One way of measuring the sensing matrix is to inject controlled oscillations into the digital control output of the various length servos at an audio frequency outside of the control bandwidth, and demodulating the various signals at the excitation frequency. Notch filters in the controllers are used to make sure that unintended responses are induced via cross-couplings. Through knowledge of the mechanical and electronic responses of the elements in the control loops, one can calculate the optical response to a displacement in a certain length degree of freedom. A measurement of the 40m prototypes DRFPMI sensing matrix is shown in Figure 3.11.

The interpretation of the CARM and REFLI result is not entirely straightforward, due




Figure 3.10: Displacement sensitivity of the 40m prototype interferometer. “Dark noise” refers to the noise present in the output photodiode and associated signal chain when no incident light is present. “DAC” noise is due to distortions from the DACs that drive the test mass position actuators. The trace labeled “fundamental” is the total noise shown in Figure 3.9.

to the analog CARM control. It is not practical to excite CARM or to digitally monitor the RFPD responses above the 30 kHz control bandwidth, so the excitation is subject to loop suppression at the 310 Hz excitation frequency. Furthermore, it is not feasible to notch out the response of the CARM servo at the excitation frequencies of the other lengths, as this would make the loop unstable. Thus, since the analog REFL_{II} I signal is used as the error signal for the high bandwidth CARM loop, the response of REFL_{II} I to any length degree of freedom is suppressed by the CARM loop.

Finally, camera images of the light at the AS, REFL, and POP ports in the locked state can be seen in Figures 3.12, 3.13, and 3.14. The significant asymmetry and side-lobes evident in these pictures indicate significant misalignments in the interferometer. This is ultimately unsurprising given the current lack of interferometric alignment sensors, as are used in




Figure 3.11: Sensing Matrix Measurement of the 40m prototype interferometer during DRFPMI operation. The radial dimension is logarithmic, each grid mark indicating a decade. The I and Q labels refer to the two analog outputs of RF demodulation units, which are thereafter digitized, rotated, and scaled to create the error signals for the length control loops. The response is measured in units of Watts of RF laser intensity fluctuation at the demodulation frequency incident on the photodiode surface per meter of change in the interferometer length coordinate.

aLIGO. Up to this point, these have been elided at the 40m prototype for the sake of controls simplicity.




Figure 3.12: AS port camera during DRFPMI Lock




Figure 3.13: REFL port during DRFPMI Lock




Figure 3.14: POP port during DRFPMI Lock

3.7 FUTURE WORK

3.7.1 LOCK MAINTENANCE

There are several hardware upgrades that could be considered at this point to improve the robustness and reliability of the 40m locking processes.

One of the main limitations on longer lock stretches is the lack of interferometric alignment sensing. As discussed above, the final alignment state of the interferometer includes feedback loops acting on individual optics and referenced to fixed points on in-air optical tables with no seismic isolation, which are limited by the precision of the initial dither alignment servos. Ideally, one would employ resonant wavefront sensors at the AS and/or REFL port to stabilize the differential and common arm angular modes and actuate in a diagonalized basis on the test masses, as in aLIGO. This will be especially necessary for the higher circulating powers afforded by any increase in recycling gain, which should push the radiation pressure dynamics into the unstable regime.

Another potential upgrade that is currently in early testing is the use of fiber coupled light from the auxiliary ALS lasers before the frequency doubling crystal to create an fiber-coupled IR optical beat, in place of the free-space green light beat. This would have the advantage of more available power for the optical beat and eliminate the need for free-space mode matching and frequent manual alignment maintenance.

Finally, while seismic feed-forward is used to reduce angular motion of the PRC and excess length noise in the input mode cleaner using a seismometer located at the interferometer vertex, additional noise reduction can be achieved through the use of the seismometers located at each end station. Reducing the arm length fluctuations at low frequencies will make the ALS to PDH hand-off for CARM control much more reliable.

3.7.2 SENSITIVITY IMPROVEMENT

Now that the aLIGO interferometers are locking robustly and repeatedly, it is no longer as critical to maintain the same configuration at the 40m. The 40m prototype can prove to be a useful test bed for integrating novel technologies into a full suspended interferometer system, such as cryogenic cooling of the test masses.

One aspect of the 40m design that we may consider changing is the signal recycling cavity. In aLIGO, it is desirable to broaden the interferometer's response bandwidth in order to have good sensitivity at higher frequencies. Especially since there are many technical noise

sources at lower frequencies that limit the current performance (see Section 5.1), the low frequency response attenuation due to the RSE technique is not a major drawback.

However, as shown in Table 3.1, the signal extraction pole is 167 kHz, which is far outside what we measure with the data acquisition system. Indeed, even the single 40m cavity pole of 4.4 kHz is near our digital nyquist limit of 8192 kHz. Thus, it is worth examining what benefit could be had at the 40m by employing true signal recycling, where the carrier light resonates in the SRC; amplifying the low frequency response at the expense of response bandwidth.

Figure 3.15 compares the theoretical responses functions and quantum noise levels for signal recycling (SR), resonant sideband extraction (RSE), and the removal of the SRM, keeping all other aspects of the interferometer constant. Although the enhancement of the interferometer response in the SR case does afford a region of increased sensitivity from 200 Hz to 4 kHz, the increased differential mode finesse also worsens the influence of quantum noise. In fact, wholesale removal of the SRM provides a much broader region of quantum noise improvement, from 100 Hz to 30 kHz. This would have the additional benefit of simplifying the length and alignment controls of the interferometer.

There is a much larger configuration space that can be searched, however. For instance, one could use a detuned SRC configuration to achieve sub-SQL performance in a limited frequency band, or increase the mirror masses to reduce the displacement fluctuations due to radiation pressure forces.




Figure 3.15: Comparison of the interferometer response (top) and quantum noise (bottom) for signal recycling (SR), resonant sideband extraction (RSE), and removal of the SRM. The response functions are normalized by the single arm cavity response at DC. The labels on the response curves indicate the theoretical response function pole frequency following IZUMI and SIGG [15]. The quantum noise depicted includes both shot noise on the output photodiode, but also the force noise on the test masses due to radiation-pressure effects. Quantum noise levels are computed with FINESSE [51], incorporating the lowered loss levels discussed in Section 3.6.1 and DC readout. Also shown for comparison is the “standard quantum limit” for the position measurement of a free mass, as described by BRAGINSKY and KHALILI [52].

CHAPTER 4

AUTOMATING INTERFEROMETER CONTROL

As the technology and design behind terrestrial gravitational wave detectors evolves and becomes more complex, the sophistication of the required controls engineering is increased in kind. The adaptability of real-time digital control systems is indispensable for commissioning reliable procedures for controlling the instruments. We can create complex logic and signal flow, execute arbitrary computations on recorded signals, make precise measurements from desktop computers in a control room, and make targeted changes to the control scheme without modifying any hardware.

This chapter will detail the work done on an algorithmic approach to automating aspects of interferometer control, specifically by using a composite feedback control error signal which is made up of multiple available signals weighted in a way that minimizes the noise in the signal. We will refer to this algorithm as CESAR (Combined Error Signal from Automatic Regression).

4.1 INTERFEROMETER LOCKING TECHNIQUES

The initial incarnation of LIGO used a Power-Recycled Fabry-Pérot Michelson configuration (PRFPMI), in which it was necessary to keep four length degrees of freedom under interferometric feedback control using signals derived via optical heterodyne techniques (see Section 1.3). In the PRFPMI configuration, the differential degrees of freedom — the differential arm cavity length (DARM) and short Michelson length (MICH) — are more or less cleanly separable in the signals at the anti-symmetric (AS) port and the power recycling cavity pick-off (POP) port. However, the optical dynamics of the common arm length (CARM) and power recycling cavity length (PRCL) are entangled to a degree where careful design and thought is required to bring the interferometer into its fully operation state.

In REGEHR et al. [53], stable control of a tabletop PRFPMI was achieved using signals derived from a single phase modulation (PM) sideband. Calculating the necessary conditions for stability in the PRFPMI configuration allows one to derive the characteristics of the controllers that are suitable for active feedback control of all four degrees of freedom in the narrow regime around the operating point, using detailed knowledge of the optical system. That is, the response of every optical sensor to every length degree of freedom, a.k.a. the optical sensing matrix, was calculated a priori and used to inform the design of the controllers. Then, the process of lock acquisition is very straightforward: the controller is continuously active, and an unlocked interferometer that coincidentally passes through the operating point is subjected to sufficient corrective actuation to be held there.

However, this process is reliant on a certain confidence that the all four interferometer length degrees of freedom will coincidentally pass through the proper operating point frequently enough, and that the actuators are of sufficient strength to apply the necessary transient force before the system leaves the linear regime. In EVANS et al. [54], the authors devised a multi-step sequence where portions of the interferometer were locked in conditionally stable states, which could be initiated more frequently and repeatably, leading up to a fully resonant interferometer. Auxiliary optical signals, such as the circulating power at various points, were used to dynamically estimate and invert the sensing matrix in real time to continuously control the cavity lengths as it traversed through these different states. This technique had the advantage of widening the phase space where lock acquisition could be initiated, increasing the frequency of lock acquisition attempts and avoiding the need for large transient forces. However, it could be difficult to troubleshoot this system when it was not working, as it relied on accurate and precise knowledge of the

optical parameters of the system that could be challenging to monitor.

As the design process for Advanced LIGO was ongoing, an important aspect was the consideration of the feedback control and lock acquisition with the incorporation of a signal extraction cavity, as part of the DRSE scheme. This additional cavity length necessitated the introduction of new ways to derive information of cavity lengths and procedures for locking the entire interferometer [55]. Some of the successfully demonstrated ideas included: the use of an amplitude modulation induced single sideband resonant solely in the PRC to serve as a stable oscillator for the PM sideband resonant in the SRC [56]; using an auxiliary input laser to provide a subcarrier field at a tunable offset from the carrier to control and dynamically detune the SRC [57]; and using two PM sidebands with different resonant conditions in the PRC and SRC along with double-demodulation to remove the influence of the arm motion induced carrier field fluctuations [58].

The first demonstration of lock acquisition in a suspended DRSE interferometer was performed at the 40m prototype interferometer [10], in a proposed aLIGO configuration using two frontal PM sidebands. Some of the technical challenges observed in this effort led to the development of the Auxiliary Laser Stabilization (ALS) technique, which was also demonstrated at the 40m prototype [23]. ALS involves dichroic mirror coatings and injection of frequency doubled light from auxiliary lasers into the arm cavities to provide independent error signals for the arm cavities that is independent from the highly nonlinear DRMI cavities (see Section 1.3.4).

In addition, the complications of the input optics necessary for the double demodulation technique were omitted in favor of the 3F locking technique [25] in aLIGO, where new DRMI error signals are derived from higher order demodulation products that have a significantly reduced contribution from the carrier field that interacts strongly with the arm cavities, thereby enabling more robust control of the DRMI before the arm lengths are fully stabilized. Together, ALS and 3F signals provide sufficiently separated error signals for the arm and DRMI cavity lengths such that the CARM offset reduction sequence (as summarized in Section 1.3.4) can be reliably performed at the aLIGO observatories [12].

4.1.1 DYNAMIC COMBINATION OF SIGNALS

At their core, the steady-state control loops currently in use are fundamentally similar to the era of analog control. With engineering experience and trial-and-error, a single signal — often drawn from a set of possible choice — is used as an error signal in a manually designed controller and feedback loop. However, as has been discussed in Section 1.3.4,

modern interferometer designs rely on having multiple signals available for providing an estimate of a physical quantity that have partially overlapping regimes of validity or utility. Examples include the CARM length degree of freedom that can be measured via ALS, the DC arm transmission, and demodulated RF photodiode signals and angular motion of the test masses from local suspension sensors, optical levers, or interferometric wavefront sensors.

Using multiple measurements of a quantity will generically have a lower statistical uncertainty associated with the combined estimation as compared to any individual measurement used as an estimator. So, in principle, one can improve the measurement of a interferometer degree of freedom by appropriately combining all of the signals that provide even a weak measurement of that quantity. This would also reduce the influence of sensing noise in a running control loop.

Apart from noise concerns, reducing the active cognitive load on the scientists commissioning and operating gravitational wave detectors is critical for advancing the sophistication of the hardware and technology that can be employed. Constructing automated control systems that require minimal a priori knowledge about the control plant can potentially simplify lock acquisition efforts for many future generations of gravitational wave detectors with novel configurations.

4.2 THE CESAR ALGORITHM

The goal of the CESAR algorithm is to produce a “super-signal” for use in an active feedback stabilization loop that combines multiple candidate error signals in a way that minimizes the uncertainty and noise in the sensing of the feedback error signal without a priori knowledge of the nature of the candidates. This, in turn, reduces the actuation noise present due to the active stabilization. The work presented in this chapter limits itself to linear combinations of the candidate signals; some practical consequences of this limitation will be discussed in Section 4.4. The algorithm is intended to be able to run in real time, adjusting the signal combination as the underlying system parameters evolve or moves through different regimes.

Figure 4.1 sketches the basic topology considered for the application of the algorithm, where x is the target quantity of interest measured by the n candidate signals s_i ; the CESAR error signal e is composed of a weighted sum of the candidates.

We will initially consider the system to be artificially held at some constant state at which $x(t) \approx x_0$, and where the fluctuations in $x(t)$ are small enough that we may treat the x dependent output of each sensor, s_i , as the the first term in the Taylor series, i.e. a purely linear signal:

$$s_i(t) \approx \delta_i(t) + d_i(x(t) - x_0) \quad (4.1)$$

where δ_i represents a random noise process reflecting the sensor noise and

$$d_i := \left. \frac{\partial s_i}{\partial x} \right|_{x=x_0} \quad (4.2)$$

So, we must dynamically evaluate the content of the candidate signals, estimating their response to the desired quantity and the associated sensing noise. First, we must decide what we mean by the “optimal” combination of signals. For the ease of analysis, the metric chosen is simply to minimize the total RMS of the combined CESAR signal, while maintaining a constant response to the target quantity.

The basic procedure of the CESAR algorithm as executed in this work is as follows:

- Continuously excite the target quantity with a constant sinusoidal excitation
- Estimate the response, d_i , of each candidate signal to the excitation, and the uncertainty of this estimation, $\sigma_{d,i}$.




Figure 4.1: Block diagram of the feedback topology considered for the application of the CESAR algorithm. Blocks correspond to Laplace-domain linear time invariant elements. x is the target quantity to be sensed and estimated. There are n candidate signals, s_i , which measure x with some responses d_n and are subject to uncorrelated sensing noises δ_i . The CESAR error signal is the sum of contributions from the candidates, weighted by the c_n coefficients. K is the controller and P is the system plant. Common force and sensing noises, δ_f and δ_s , are included for completeness.

- Calculate the RMS of the independent sensing noises, converted into common units by their estimated responses and separated via frequency domain coherence, $\sigma_{x,i}$.
- Calculate and apply the weights, w_i , that minimize the RMS of the combined signal while maintaining a constant response to the target quantity.
- Periodically repeat the above steps.

The following sections will discuss these steps in more detail.

4.2.1 LOCK-IN RESPONSE SENSING

The method for measuring the response of each sensor signal is fairly straightforward, as long as we adhere to the assumption that the frequency responses in the regime of interest are flat. We can use the common technique of sinusoidally exciting the system with the actuator at some frequency outside of the loop bandwidth. The controller response may feature a notch at the excitation frequency, to ensure that the feedback does not affect the

estimation. Then, each sensor signal can be demodulated at the excitation frequency, to evaluate its response.

Mathematically, we can write the sensor signal in the presence of some forcing term $u = A_D \sin(\omega_D t)$ as:

$$s_i(t) = \delta_i(t) + d_i \left(x(t) + \delta_s(t) + \tilde{P} \star \delta_f(t) + \tilde{P} \star A_D \sin(\omega_D t) \right) \quad (4.3)$$

We then define the output of the complex demodulation as:

$$D_i := 2 \langle |s_i(t) e^{i\omega_D t}| \rangle = d_i P(\omega_D) A_D \quad (4.4)$$

Here, we represent the dynamics of the plant by convolving the applied forces with the impulse response of the plant, \tilde{P} . In practice, we take the expectation value by taking the average value of the demodulated signal, after a suitable low pass to eliminate unwanted terms that arise at $2\omega_D$. We will call this estimate \bar{D}_i .

In order to maintain the same overall response of the CESAR signal throughout the offset stepping process, a reference measurement is made when the algorithm is initiated. Thereafter, the signal responses are determined relative to this initial measurement. Another way of looking at this is that the system attempts to keep the response of the error signal at the excitation frequency constant. So, when the system is first brought under control, this demodulation is performed to determine \bar{D}_0 , and thereafter

$$\frac{\bar{d}_i}{\bar{d}_0} = \frac{\bar{D}_i}{\bar{D}_0} \quad (4.5)$$

If the actuator and sensor calibrations for D_0 are known, these quantities can be determined in physical units. Otherwise, however, it is enough to set $\bar{d}_0 := 1$, since we just need to track changes relative to the starting point.

In order to quantify the error in these estimations, we can associate some standard random error in D_i using the RMS fluctuations of the demodulated signal after the low pass has settled: $\sigma_{D,i}$. Then, the error in \bar{d}_i can be written as:

$$\left(\frac{\sigma_{d,i}}{\bar{d}_i} \right)^2 = \left(\frac{\sigma_{D,i}}{\bar{D}_i} \right)^2 + \left(\frac{\sigma_{D,0}}{\bar{D}_0} \right)^2 \quad (4.6)$$

4.2.2 OPTIMAL WEIGHTING OF THE ERROR SIGNALS

We may create an estimate of $x(t)$ from each candidate sensor

$$x_i(t) := \frac{s_i(t)}{d_i} \quad (4.7)$$

and create the CESAR error signal as a weighted combination of these:

$$e(t) = \sum_i w_i x_i(t) \quad \text{where} \quad \sum_i w_i = 1 \quad (4.8)$$

In this hypothetical state, the content of each estimate is dominated by the δ_i terms. (The reality of in-loop effects and common noise will be considered in Section 4.2.3). Then, the RMS variation of the CESAR error signal can be written as

$$\sigma_e^2 := \sum_i w_i^2 \sigma_{x,i}^2 \quad (4.9)$$

where

$$\sigma_{x,i} := \frac{\sigma_{\delta,i}}{d_i} \quad (4.10)$$

$$= \frac{\sqrt{\langle \delta_i(t)^2 \rangle}}{d_i} \quad (4.11)$$

That is, $\sigma_{x,i}$ is the variance of $s_i(t)$ converted to the units of x . (We will refine the definition below).

We want to find a choice of w_i s that minimizes σ_e^2 . This can be done by taking the gradient with respect to a vector constructed of the w_i s, and using a Lagrange multiplier to enforce the constraint on the sum of the weighting coefficients:

$$L(w_i, \lambda) = \lambda \left(\sum_i w_i - 1 \right) + \sum_i w_i^2 \sigma_{x,i}^2 \quad (4.12)$$

$$\Rightarrow \frac{\partial L}{\partial w_i} = \lambda + 2w_i \sigma_{x,i}^2 = 0 \quad (4.13)$$

$$\Rightarrow w_i \propto \frac{1}{\sigma_{x,i}^2} \quad (4.14)$$

Thus, weighting the signals proportionally to the inverse of the individual converted variances will minimize the variance of the combined error signal. As long as one normalizes the weightings by their sum, initially using $w_i = \sigma_{x,i}^{-2}$ is enough. Given this result, if we want to minimize the RMS of our error signal through combination of various sensors, we must estimate the *independent* variance associated with each sensor.

In practice, we do not have perfect knowledge of the d_i terms, and instead use the estimated value derived from the method in Section 4.2.1.

$$\bar{x}_i(t) = \frac{s_i(t)}{d_i} \quad (4.15)$$

The question is, now, how to write the variance of $e(t)$ in the form of Equation 4.9. We have:

$$e(t) = \sum_i w_i \frac{s_i(t)}{\bar{d}_i} \quad (4.16)$$

$$= \sum_i w_i \frac{d_i x(t) + \delta_i(t)}{\bar{d}_i} \quad (4.17)$$

The method for estimating the individual uncorrelated sensor noise terms, δ_i , will be shown in Section 4.2.3, we will take them as known for now.

In Equation 4.16, there are two mechanisms for the introduction of uncertainty. In addition to the sensor noise itself, the error in the estimation in d_i will have an effect. Intuitively, this makes sense; if the response is not correctly sensed, we arrive at an incorrect estimation of the effective sensor noise level, and may be fooled into placing unfounded confidence in the signal. This is qualitatively different from the direct effect of random sensor noise, and requires that refine our definition of $\sigma_{x,i}$.

Despite this qualitative difference, the simplest practical way to estimate the variance associated with each term in the sum would be to consider the numerator and denominator in Equation 4.16 as quantities sampled with some standard error, thus allowing us to use the rules for the propagation of standard error. While not mathematically rigorous, it allows us to make a quantitative estimation of the total variance of each sensor. That is:

$$e = \sum_i w_i \frac{d_i x \pm \sigma_{\delta,i}}{d_i \pm \sigma_{d,i}} \quad (4.18)$$

$$:= \sum_i w_i (x \pm \sigma_{x,i}) \quad (4.19)$$

Thus,

$$\langle e \rangle = x \quad (4.20)$$

$$\sigma_{x,i}^2 := \left(\frac{\sigma_{\delta,i}}{d_i} \right)^2 + \langle x^2 \rangle \left(\frac{\sigma_{d,i}}{d_i} \right)^2 \quad (4.21)$$

$$\text{and } \sigma_e^2 = \sum_i w_i^2 \sigma_{x,i}^2 \quad (4.22)$$

Here, we must now augment the equations with knowledge gleaned from the closed-loop system. Absent perfect knowledge of $\langle x^2 \rangle$, we may use the variance of the error signal at the previous step, as this will be related to the levels of fluctuations of x while the loop is closed. Ultimately, since it is the relative variances of the sensors that determines the

weightings, very precise determination of the absolute variances is not critical as long as their hierarchy is preserved.

So, we arrive at our expression for our chosen weighting of each sensor's signal, given estimations of the sensor noise, the response of the sensor, and the uncertainty of the determination of the response (before normalization by the sum of the weights):

$$w_i = \left[\left(\frac{\sigma_{\delta,i}}{\bar{d}_i} \right)^2 + \sigma_{e,\text{Prev}}^2 \left(\frac{\sigma_{d,i}}{\bar{d}_i} \right)^2 \right]^{-1} \quad (4.23)$$

4.2.3 INDEPENDENT SENSOR NOISE ESTIMATION

If multiple sensors are combined to create the CESAR error signal, then the sensor noise of each of them will propagate through the loop, and be present in every other sensor's signals. Let us examine how we may arrive at an estimate of each sensor's uncorrelated sensing noise by in-loop measurements. For the remainder of this section, we will assume that $\bar{d}_i \approx d_i$, and neglect effects of miscalibration.

Given equation 4.16, we can close the loop and arrive at an expression for the Laplace domain in-loop error signal:

$$e = \frac{\delta_s + P\delta_f + \sum_i w_i \frac{\delta_i}{d_i}}{1 - G} \quad (4.24)$$

with $G \equiv KP$. Then, each sensor's signal is given by:

$$s_i = \delta_i + d_i \left(\frac{\delta_s + P\delta_f + G \sum_j w_j \frac{\delta_j}{d_j}}{1 - G} \right) \quad (4.25)$$

In the regime where $G \ll 1$, then:

$$s_i \approx \delta_i + d_i(\delta_s + P\delta_f) \quad (4.26)$$

With knowledge of the loop shape, G , one can arrive at a similar, more general expression:

$$s_i - Ged_i = \delta_i + d_i(\delta_s + P\delta_f) \quad (4.27)$$

In both of these expressions, we have come up with a frequency domain expression for each signal composed of measurable quantities that contains the noise common to all

sensors, and the individual sensor's noise. Thus, these signals will have one component that is coherent across all of the sensors, and one that is not. Since all we care about is the noise level of the δ_i terms, and not their time domain representation, we can estimate their power spectral density by performing coherence-weighted frequency domain subtraction. This is implemented with the frequency domain Wiener filter technique detailed in Section 5.2.2, and is applicable for an arbitrary number of candidate signals.

4.3 TESTING AND RESULTS

The test bed that was chosen for the initial testing of the CESAR algorithm was the control of a single Fabry-Pérot arm cavity at the 40m prototype interferometer. This system can be thought of as a simplified version of the control of the common arm length degree of freedom for an aLIGO interferometer, as a similar set of length error signals are available (see Figure 4.2) while not subject to the complications of control of the vertex DRMI, radiation pressure effects, or time-varying dynamics due to a changing coupled cavity pole frequency. This section will detail testing of the CESAR algorithm on a time-domain simulation of a 40m arm cavity, and the subsequent successful application to the real time control of an arm cavity at the 40m prototype.

4.3.1 TIME DOMAIN SIMULATION AND VALIDATION

Before attempting to use the CESAR algorithm on actual interferometer hardware, it was tested with a time-domain interferometer simulation in order to gauge its validity and utility. Specifically, the algorithm was tested on a simulation of a single arm cavity at the 40m prototype interferometer using the “End to End” (E2E) interferometer simulation software [59]. E2E is capable of time domain simulation of the opto-mechanical dynamics of suspended interferometers and the control systems employed to stabilize them. To make this specific simulation an appropriate gauge of the utility of the CESAR algorithm, the following properties of the 40m arm cavity system were replicated as accurately as possible:

- The optical parameters of the arm cavity optics
- The mechanical properties of the optic suspensions
- The ambient seismic noise spectrum, subject to filtering by the seismic isolation stacks.
- The control filters used for single arm locking.
- The sensor noise levels of the PDH, ALS, and transmitted cavity light signals

When executing the simulation, data is written at regular intervals to an output file with the content of a number of signals, including the individual optical sensors and the true physical fluctuations of the cavity length. The software implementation of the control




Figure 4.2: Profiles of the signals used as length error signals and input to the CESAR weighting algorithm. Details of the origins of these signals can be found in Section 1.3.

servo logic also included an input matrix whose elements could be updated by reading an external parameter file at specified times; these matrix elements being intended as the outputs of the CESAR algorithm. This enabled the creation of a program which would coordinate the execution of the simulation, and would make the necessary computations according to the CESAR algorithm to update the input matrix from the output data when the simulation was ready. Looking ahead to its implementation on the 40m interferometer itself, this program was written to be able to easily swap out the commands for retrieving data and communicating its results for those used by the interferometer operation computers, while using the identical code for the algorithm computations.

As an analog to the CARM offset reduction sequence[12], the simulation was arranged to begin with active feedback control of a single arm cavity length with the ALS error signal, with an offset of five cavity linewidths from IR resonance. The control loop would then gradually reduce this offset in small steps, computing and implementing a new error signal input matrix at each step. The control would step continuously until the cavity had passed through resonance and reached an offset of five linewidths on the other side of IR resonance, in order to test that the algorithm would not only smoothly switch to a low-noise signal when available, but also that it would abandon a previously valuable signal when it was no longer in its valid regime. For the simulation, only two signals were included as inputs to the CESAR algorithm: the ALS signal, and the PDH IR optical

heterodyne signal.

Some selected quantities from this simulation are shown in Figure 4.3. Generally, the simulated performance of the CESAR algorithm was successful, though the iterative testing of it with the simulation did motivate some practical modifications.

Specifically, since the algorithm uses information at a previous step to inform the signal weights at the following step, it is in principle possible to assign non-zero weight to a signal whose slope will reach zero at the next step, exhibiting a sensing “singularity” where no length information is actually present in the signal. The preceding step will generally exhibit a small k_i , and thus the problematic signal will be amplified by a large amount, corrupting the overall signal and often causing a lock loss.

To avoid this occurrence, signals with a score of less than 0.1 are reassigned a score of 0, as their inclusion would not improve the overall sensitivity much, and these signals may be near a sensing singularity. There is also a check at each step that the estimation of a signal’s slope, k_i , and the statistical error in that estimation, $\sigma_{k,i}$, did not include zero in $k_i \pm 5\sigma_{k,i}$. Nevertheless, as can be seen around the right side of resonance in Figure 4.3, excess noise can be introduced around sensing singularities even with this check. The most robust response to this factor is simply reducing the step size, such that sensing singularities are approached slowly enough that the converted sensing noise is observed to rise enough to lower the score below the threshold value before the singularity is reached.

4.3.2 40M SINGLE ARM TEST

With a successful test on the E2E time domain simulation complete, the code was deployed on the 40m arm cavity, using three input signals: ALS, PDH, and the DC transmitted cavity power. This is analogous to the signals used during the aLIGO CARM offset reduction sequence. Figure 4.4 shows a “lock acquisition” procedure, where the arm is initially under stable feedback control with an ALS error signal and ends up on IR resonance exclusively using the PDH error signal (due to the PDH signal’s vastly superior noise performance). The available error signals are successfully re-weighted as their equivalent displacement noise evolves, resulting in the overall RMS of the composite CESAR error signal generally decreasing smoothly as IR resonance is approached, and often of lower noise than would be possible when using a single error signal.

In order to achieve this smooth behavior however, the step size had to be quite small — about 50 μm . Thus, the lock acquisition shown took on the order of 20 min to complete. Nevertheless, the CESAR algorithm was able to replicate and improve the general CARM

offset reduction sequence for the single arm case without any a priori knowledge of the functional forms of the available signals or their regimes of validity. The solitary human intervention required was starting the program execution at 5 nm from resonance and defining the three available candidate error signals.




Figure 4.3: Time domain simulation of a 40m Fabry-Pérot arm cavity scan using the CESAR algorithm. Shaded regions, where included, represent 1σ fluctuations of the quantity at that cavity tuning.




Figure 4.4: IR Lock acquisition of a 40m arm cavity using the CESAR algorithm. The starting condition was stable feedback control using solely the ALS signal at 5 nm from resonance. The CESAR program took gradual steps to decrease this offset until a manually specified transmitted power threshold was met, coinciding with IR resonance.

4.4 FUTURE WORK

As detailed here, the performance of the CESAR algorithm is perhaps too slow for convenient usage. One approach to mitigating this factor would be to introduce a persistent memory store that keeps track of the slope and mean value estimates at each operating point, thereby reconstructing the functional form of the signal. Then, when locking, the algorithm could use previously measured features to quickly adjust the blending the signals at each step, without having to wait long enough to remeasure with reasonable statistics. This would additionally discover sensing singularities on the first pass, allowing the algorithm to take larger steps when it is in a regime devoid of singularities.

However, a more pressing limitation for use specifically for CARM offset reduction is that the algorithm as described here does not include considerations for error signals whose frequency response to the physical quantity of interest is non-trivial, as is the case when optical pole effects are present, for example. In the 40m arm cavity case, the single arm cavity pole is on the order of 10 kHz, which is well above the digital control UGF of 100 Hz and thereby easily neglected. However, during the aLIGO lock acquisition sequence, the coupled cavity pole is not only low enough to have an appreciably different frequency response than the ALS signal, this pole moves from ≈ 80 Hz to ≈ 0.5 Hz as the CARM offset is reduced.

One could imagine extending the CESAR algorithm to include multiple sensing excitations at a few frequencies, and interpolating a filter for each signal such that it roughly inverts the changing frequency response. However, this possibility has not been investigated at this time.

Nevertheless, there are situations where offset steps and changing sensing dynamics are not as critical where the current state of the CESAR algorithm could potential be beneficial. For example, one could implement an automated transition from the 3F to 1F DRMI error signals that is robust to changes in sensing magnitudes, or an automated transition from ALS to RF CARM control in 40m-style locking (see Section 3.4) with a static optical pole compensator present for the IR PDH signal.

Finally, work is ongoing to investigate the suitability of using neural networks to reconstruct linear length degree of freedom signals from the ensemble of optical signals, including both DC power and RF heterodyne measurements. If trained on accurate interferometer simulations, the network can provide signals with wider linear regimes that can be used as input to conservative feedback controllers to slow the system down to the point where the unmodified optical signals can be used with their usual controllers.

CHAPTER 5

OFFLINE NOISE SUBTRACTION

5.1 INTRODUCTION

During the first observational run of the Advanced LIGO detectors (O1), from September 12, 2015 to January 19, 2016, two binary black hole merger events — GW150914 [6] and GW151226 [60] — were confidently identified, while a third signal — LVT151012 — was deemed to have an 87% probability of being of astrophysical origin [61]. These observations were a direct consequence of the greatly enhanced sensitivity of the aLIGO detectors over previous generations of terrestrial gravitational wave detectors.

The ultimate sensitivity of the aLIGO detectors is dictated by the physics inherent to their design, such as shot noise of the laser light or thermal fluctuations of the mirror coatings and optic suspensions [4]. However, the performance during the first observing run was also influenced by technical noises which arise from factors such as the instrumentation or control of the interferometer [5]. The confidence in the significance of any given signal and our ability to extract astrophysical information from it is directly impacted by the noise and sensitivity of the detectors at the time, and so there is a strong need to improve their performance by any available means.

In general, the performance of a single detector is characterized by separately considering the plethora of mechanisms by which non-astrophysical signals couple into the strain output of the instrument, such as the shot noise of the light incident on the output photodiodes or thermal motion of the arm cavity mirror surfaces. Once categorized into causally distinct groups, we can predict the instrument's performance from the incoherent sum of these noise mechanisms, and compare it to the observed steady state sensitivity. This is a crucial analysis when working to understand and improve the performance, as a diag-




Figure 5.1: Low frequency noise budget for the L1 aLIGO interferometer during O1, from MARTYNOV et al. [5].

nosis of what aspects or subsystems of the detector are the limiting factors. It also shows us where the noise exceeds the sum of the budgeted noise sources, and thereby where our understanding of the noise is incomplete.

Figure 5.1 shows a noise budget of the lower frequency regime assembled as a representative state of the aLIGO detectors during O1. (The performance at high frequencies is generally well understood, and corresponds to the shot noise at the output photodiodes) It can be seen that at low frequencies there remains some amount of noise that remains unexplained. In addition, there is a significant contribution from technical noise arising from the control of the suspended optics. This is due, in part, the control actuation necessary to keep the instrument well aligned over long periods of time and the gradual shifts in the beam spot positions. Higher mass binary systems, such as the $\approx 30 M_{\odot}$ and $35 M_{\odot}$ black holes that produced GW150914, merge at lower frequencies, making the low frequency performance of the detectors especially important for this kind of event. The recorded signal from GW150914 only spent about 200 ms in the sensitive band of the instrument, resolvable from about 35–250 Hz [6]

The aLIGO detectors employ numerous subsystems that control different aspects of the instrument and monitor its state. These are coordinated and operated in large part by a distributed digital control system which measures and records a large numbers of sig-

nals related to these subsystems, in addition to the main output which measures space-time strain. Thus, numerous signals are synchronously recorded along with the interferometer output, such as those from environmental sensors, mirror suspension actuation, and photodetectors. These auxiliary signals have the potential to witness coupling of unwanted noises into the interferometer, and are used in the commissioning of the detector to diagnose and mitigate such couplings.

Many of these couplings were expected and accounted for in the commissioning and operation of the LIGO interferometers, and could be measured and characterized in a fairly straightforward manner. For instance, excess motion of the interferometer's beam-splitter causes some differential phase shift in the recombining beams exiting the arm cavities, which directly contributes to the output strain signal. The excess beam-splitter motion above 10 Hz itself is primarily due to the Michelson sensing noise's influence on the length actuation signals [62]. Since this coupling arises within a system under our control, we can measure the influence of the beam-splitter actuation on the strain signal, and induce correctional feed-forward actuation to eliminate its contribution to the recorded strain signal once this coupling is well characterized; this has been done in the real-time control system for the LIGO interferometers [13, 16].

Seismometer signals have also been used to train feed-forward subtraction filters that are run in real-time to reduce the physical motion of interferometer elements, see [45] and Section 3.3. This manner of online subtraction has the strong benefit of reducing the gain or dynamic range requirements of the length and/or angular feedback systems.

However, gradual changes in instrument state, such as alignment or thermal state, can cause changes in the expected couplings during an observing run, when it is preferred to make as few configuration changes to the instrument as possible. This may lead to unwanted noise making its way to the recorded strain data, despite the necessary information required to subtract it being available. Furthermore, there is the possibility of unconsidered noise couplings being present that could in principle be predicted from other recorded signals. At this point, the only recourse is to revisit previously recorded data and attempt to regress the unwanted noise out.

One technique for reducing the noise in the strain signal post-facto using auxiliary information is Wiener filtering[63], a multiple-input single-output (MISO) algorithm which optimizes the mean squared difference between the subtraction target and the predicted coupled noise from multiple witnesses, taking the correlations between the witnesses themselves into account. Time domain Wiener filtering has been used successfully in terrestrial gravitational-wave detectors to enhance the performance of the vibration iso-

lation system[64] and reduce the influence of local gravitational field fluctuations[65]. We also find the representation of the Wiener filter in the frequency domain[66] useful; in the single witness case, this is identical to the usual concept of frequency-domain coherence, and so it is also referred to as MISO coherence.

This chapter will describe the manner in which auxiliary signals were used post-facto to estimate noise couplings that existed during the gravitational wave events in O_1 , and thereafter regress and subtract that noise from the recorded strain data using Wiener filters. It will describe the algorithms used for calculating and characterizing noise couplings from recorded aLIGO data, and how these were validated to not corrupt or bias the resultant estimates of the astrophysical source parameters. It will detail the sensitivity improvement that results from this subtraction and the consequent improvement in the confidence in the estimates of the source parameters of GW₁₅₀₉₁₄. It will also describe ongoing efforts to extend these regression techniques to more complicated nonlinear couplings.

5.2 FOUNDATIONS

5.2.1 ADDITIVE SIGNAL REGRESSION

Any procedure that alters the gravitational wave strain data must be carefully considered for biases or distortions that may be introduced, as these may corrupt the validity of the astrophysical information inferred from the data. In this section, we will examine the conditions that allow us to perform regression of unwanted noises while maintaining fidelity to the astrophysical signals. This consideration is similar to that found in SCHÖLKOPF et al. [67].

We will consider independent physical processes and quantities x_i which have some causal additive relationship to a set of observed variables y_j . The influence on the variables may be through nonlinear, frequency dependent, and time-varying functions of the underlying quantities, and may also contain independent additive noise. I.e.

$$y_j := \sum_i f_{i,j}(x_i) + N_j \quad (5.1)$$

where the $f_{n,i}$ are the arbitrary coupling functions, and the N_j are independent noise processes on the y_j variables themselves.

Let us now consider a simple case which will provides the simplest analog to our applications: with two underlying processes, $x_0 = h$ and $x_1 = g$, which influence y_0 and y_1 (see Figure 5.2). We are interested in obtaining the best possible estimate of h , whose influence is only recorded in y_0 . However, a process g also influences y_0 , yet is independently witnessed in y_1 .

Without loss of generality, let us assume $f_{h,0} = 1$, as this is in practice accomplished by the calibration of the instrument. Then, we would ideally reconstruct h in the absence of noise via

$$\begin{aligned} h &= y_0 - f_{g,0}(g) \\ &= y_0 - f_{g,0}\left(f_{g,1}^{-1}(y_1)\right) \end{aligned} \quad (5.2)$$

So, we seek to somehow create a function $W_{1,0}(y_1)$ that approximates $f_{g,0}\left(f_{g,1}^{-1}(y_1)\right)$, and thereby provides an improved estimate of h :

$$\tilde{h} := y_0 - W_{1,0}(y_1) \quad (5.3)$$




Figure 5.2: Simple additive signal model

In the presence of the noise terms, then

$$\tilde{h} = h + f_{g,0}(g) + N_0 - W_{1,0}(f_{g,1}(g) + N_1) \quad (5.4)$$

Since y_1 is completely independent of h , there is no danger of \tilde{h} having different dependence on h than y_0 for any choice of W .

Since h and g are completely independent processes, the only way to reliably reduce a metric of power or amplitude in \tilde{h} by the subtraction of $W_{1,0}(y_1)$ from y_0 is by coherently negating the contribution of $f_{g,0}(g)$, even though there is no immediate causal relationship between y_1 and y_0 . This affords us the freedom to attempt any and all regression techniques to attempt to predict values of y_0 from y_1 , subject to tests of stability and over-fitting.

This simple case translates well to the case of interferometric gravitational wave detectors where the interferometer output signal is unique in that it can conceivably be influenced by space-time strain fluctuations, and is also subject to other influences which may be observed by other auxiliary sensors and signals in the instrument and facility.

Thus, we may construct W functions to predict the values of an instrument's output, leaving \tilde{h} as the residual of this regression, provided the following conditions on the auxiliary signals are met:

- The signal is truly independent on the gravitational wave strain. This would invalidate the starting assumption in Equation 5.1. In practice, this means we must exclude “unsafe” signals that are involved in the feedback control of the interferometer's differential arm degree of freedom, or otherwise in the recording of the instrument's output.
- The functions to be approximated exist and are part of the set of functions that may be regressed by the chosen technique (or can be well approximated by members of this set).

- The noises in the auxiliary sensors must be low enough, such that even applying a perfectly estimated W still leads in reduction of noise in the output signal. I.e., if $W_{1,0}$ in Equation 5.4 is perfectly estimated, it follows that:

$$\tilde{h} = h + N_0 + f_{g,0} \left(f_{g,1}^{-1} (N_1) \right) \quad (5.5)$$

This eradicates any benefit of the regression if N_1 is of sufficient magnitude.

5.2.2 WIENER FILTERING

In this section, we will describe the origin of the time and frequency domain representations of the Wiener Filter, a method for optimal estimation of signal couplings. More detailed and rigorous derivations can be found in WIENER [63] and ALLEN et al. [66].

When diagnosing the noise performance of our instruments, we often examine the signal content in the frequency domain, as the overall signal power present in the interferometer output is dominated by low frequency seismic motion outside of the primary sensitive band. Hence, we will begin by describing how we may subtracted correlated signals in the frequency domain.

We will represent the subtraction “target” as y , a signal that contains some underlying signal of interest (i.e. the gravitational wave strain) in addition to various noise sources. Our “witnesses”, x_i are some set of other signals that may contain information about the target’s noise terms, but *not* the desired signal (as represented in Figure 5.2). This is a critical convenience afforded to us by our confidence that true gravitational wave strain signals are uniquely present in the main readout signal of the interferometer; any attempted noise subtraction from a combination of witness signals may only increase or reduce the influence of noise terms, and cannot fundamentally alter any present strain signals.

The frequency domain representations of the signal content of the target and witness signals are given by their discrete Fourier transforms (DFT), \tilde{y}^k and \tilde{x}_i^k , where the k index refers to a specific DFT frequency. The power spectral density of our signals is proportional to the squared amplitude of the complex DFT values. We frame subtraction in the frequency domain as finding a set of complex coupling constants, \tilde{w}_i^k that predicts values of \tilde{y}^k from the \tilde{x}_i^k such that the squared amplitude of the difference between the prediction, \tilde{z}^k and target is minimized.

That is,

$$\tilde{z}^k := \sum_i \tilde{w}_i^k \tilde{x}_i^k \quad (5.6)$$

where the \tilde{w}_i^k are found by solving

$$\frac{\partial}{\partial \tilde{w}_i^k} \langle \tilde{e}^{*k} \tilde{e}^k \rangle = 0 \quad (5.7)$$

where is the subtraction residual $\tilde{e}^k := \tilde{y}^k - \tilde{z}^k$ and brackets indicate averaging the quantity over subsequent DFT of windowed segments of the time domain signals, as in Welch's method of PSD estimation. (This is necessary because, for a single set of DFTs, one could arrange any linear combination of arbitrary witness DFT values to reconstruct the \tilde{y}^k through a suitable choice of \tilde{w}_i^k even if the witnesses were truly uncorrelated.)

For uncorrelated signals, one would expect that $\langle \tilde{x}^{*k} \tilde{y}^k \rangle$ would be zero, as this is related to the cross spectrum of the signals. Heuristically, we can imagine that Equation 5.7 will be satisfied when there is zero correlation between any of the witnesses and the subtraction residual, as any non-zero correlation would imply that the magnitude of the subtraction residual could be reduced by an additional contribution from the witness.

We will write this condition for each \tilde{x}_i^k as:

$$0 = \langle \tilde{x}_i^{*k} \tilde{e}^k \rangle \quad (5.8)$$

$$= \langle \tilde{x}_i^{*k} \tilde{y}^k \rangle - \langle \tilde{x}_i^{*k} \sum_j \tilde{w}_j^k \tilde{x}_j^k \rangle \quad (5.9)$$

$$\Rightarrow \langle \tilde{x}_i^{*k} \tilde{y}^k \rangle = \sum_j \langle \tilde{x}_i^{*k} \tilde{x}_j^k \rangle \tilde{w}_j^k \quad (5.10)$$

We may use vector notation to summarize and solve this set of equations at the frequency bin specified by k . We define:

$$\mathbf{w}_i := \tilde{w}_i^k \quad (5.11)$$

$$\mathbf{p}_i := \langle \tilde{x}_i^{*k} \tilde{e}^k \rangle \quad (5.12)$$

$$\mathbf{R}_{ij} := \langle \tilde{x}_i^{*k} \tilde{x}_j^k \rangle \quad (5.13)$$

Here, \mathbf{R} contains terms proportional to the power and cross spectra of the witnesses, while \mathbf{p} contains the cross spectra of the witnesses with the target. This allows us to write Equation 5.10 as:

$$\mathbf{p} = \mathbf{R}\mathbf{w} \quad (5.14)$$

$$\Rightarrow \mathbf{w} = \mathbf{R}^{-1}\mathbf{p} \quad (5.15)$$

This is the frequency domain representation of the Wiener filter. Qualitatively, Equation 5.15 ensures that we account for the correlations between the witnesses themselves

when predicting the overall correlation of the set of witnesses with the target, and thereby do not attempt to over-count shared correlations. However, this inversion can become numerically unstable if the witnesses have very high correlations. Computationally, this system can be solved very quickly for smaller numbers of signals, as it entails the calculation and inversion of one small matrix independently at each frequency bin of concern.

There is a dual time domain representation of the Wiener filter that can be reached from the frequency domain result using the fact that the complex cross products like $\tilde{x}_i^{*k} \tilde{y}^k$ are related to the cross-correlations of the time series via the DFT. In the time domain, the Wiener filters are truly FIR filter kernels, applied via convolution:

$$z[n] := \sum_i (w_i * x_i)[n] = \sum_{i,j} w_i[j] x_i[n-j] \quad (5.16)$$

Again, we can posit that that we may solve for the Wiener filters by setting the cross-correlation of the witnesses with the residual to zero at every time lag.

$$0 = (x_i \star e)[n] = \sum_m x_i[n+m] e[m] \quad (5.17)$$

$$= x_i \star y - x_i \star \left(\sum_j w_j * x_j \right) \quad (5.18)$$

One can construct time domain forms of \mathbf{w} , \mathbf{p} , \mathbf{R} such that Equation 5.15 is also the solution, where instead of complex products, one uses cross-correlations. In this case, \mathbf{R} is a block Toeplitz matrix, which makes efficient inversion via the Levinson-Durbin method possible[68], though this is also numerically sensitive to nearly degenerate signals. Another option is equivalently to take the inverse DFT of the Fourier series given by the \tilde{w}_i^k to find their impulse response, which can be used as a FIR filter kernel.

5.2.3 PARAMETER ESTIMATION

The objective of observing astrophysical gravitational wave events is to characterize the systems that gave rise to them. We use Bayesian inference to make quantitative statements of our knowledge of the binary system parameters that gave rise to an observed signal, and the uncertainty of those estimates due to noise in the instrument[69]. Black hole binary (BBH) parameters of interest include the initial and final masses of the systems, the magnitude and orientation of their spins, and their distance from the Earth. Neutron star systems will exhibit behavior dependent on their equation of state, which is a very promising avenue for new discoveries from GWs[70].

Here, we will briefly describe how the parameter estimation computation is carried out, and how the regression of excess noise can improve our confidence in the parameter values. Much greater detail can be found in ABBOTT et al. [71] and AASI et al. [72] and the references therein.

Using Bayes' theorem, we may use our observation of a GW signal in the interferometer strain data, d , to compute the posterior probability density function (PDF) of the system parameters, $\mathbf{\Lambda}$, via:

$$p(\mathbf{\Lambda}|d) = \frac{\mathcal{P}(\mathbf{\Lambda})\mathcal{L}(d|\mathbf{\Lambda})}{e(d)} \quad (5.19)$$

$\mathcal{P}(\mathbf{\Lambda})$ is the *prior* probability on the model parameters, $\mathcal{L}(d|\mathbf{\Lambda})$ is the *likelihood* of the data, and $e(d)$ is known as the Bayesian “evidence” and describes the probability of the data given the model. The evidence is typically used for model selection and enters only as an overall scaling in parameter estimation.

Assuming the discrete detector data d during on observation time T contains the GW signal and detector noise,

$$d = h(\mathbf{\Lambda}_{\text{True}}) + n$$

the log-likelihood function can be computed as

$$\log \mathcal{L}(d|\mathbf{\Lambda}) = -\frac{1}{2} (d - h(\mathbf{\Lambda}), d - h(\mathbf{\Lambda})) \quad (5.20)$$

where (x, y) is the overlap integral defined by:

$$(x, y) = 4\Re \Delta f \sum \frac{\tilde{x}^{*k} \tilde{y}^k(\mathbf{\Lambda})}{S_{n,k}} \quad (5.21)$$

Here, the \tilde{x}^k terms are the discrete Fourier transforms (DFT) at frequencies f_k , and $S_{n,k}$ is the detector's noise power spectral density (PSD) at f_k , which is proportional to $\tilde{n}^{*k} \tilde{n}^k$. Weighting this overlap integral by the PSD has the effect of lending greater importance to regions of lower detector noise relative to the strain signal.

These calculations demand knowledge of the true GW signal that would arise from a set of parameters, $h(\mathbf{\Lambda}_{\text{True}})$, which is most accurately produced by numerical simulations of general relativity [73]. However, the stochastic sampling techniques must generate a very large number of samples to reach a good approximation of the PDFs, which is not feasible to perform with dedicated simulations for each sample. To this end, analytical waveform models are created that are functions of the source parameters, and are compared to numerical simulation results to ensure their utility [74]. These models can quickly be evaluated to produce any desired $h(\mathbf{\Lambda}_{\text{True}})$ within their regime of validity.

Even with the waveform models, the complexity of the parameter space and likelihood function makes exact calculation of the posterior distributions impractical. This is resolved through the use of stochastic sampling techniques, such as Markov Chain Monte Carlo or nested sampling, which are able to approximate the likelihood function and the true PDFs [75, 76].

Reducing the background PSD via off-line noise subtraction directly improves parameter estimation by increasing SNR of the observed strain signal and increasing the likelihood of waveforms similar to the true strain signal through the likelihood function, allowing us to place tighter constraints on the properties of the astrophysical system.

In the following sections, we will primarily examine the affect of noise subtraction on the estimation of the masses of the black holes that give rise to a GW strain signal. The gravitational wave signal emitted by two merging black holes is completely determined by the masses and spins of the black holes themselves, and is typically divided into three phases: inspiral, merger, and ring-down. The low frequency sensitivity of the detectors primarily improves the fidelity of the inspiral phase of the signal, where the black hole orbital velocities and gravitational wave frequency are increasing as they approach coalescence.

In the leading order of post-Newtonian theory, the phase evolution of the strain signal during the inspiral is determined by a quantity called the *chirp mass*[71]:

$$\mathcal{M} := \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \approx \frac{c^3}{G} \left[\frac{5}{96} \pi^{-8/3} f^{-11/3} \dot{f} \right]^{3/5} \quad (5.22)$$

where f is the instantaneous GW frequency, \dot{f} is its time derivative, and m_1, m_2 are the component masses. The ratio of the black hole masses enters at the next order:

$$q := \frac{m_2}{m_1} \quad \text{where } m_2 \leq m_1 \quad (5.23)$$

The spins of the component black holes also affect the inspiral phase evolution, but also at higher order.

Thus, the estimation of the chirp mass and, to a lesser degree, the mass ratio provide a good measure of the fidelity in the observation of the inspiral phase of the GW signal.

5.3 METHODS

The general strategy for efficiently regressing unwanted noise couplings from one interferometer’s output during a GW event proceeded as follows:

- First, we search for a set of auxiliary signals which have non-negligible and non-redundant couplings to the strain signal using a “greedy ranking” algorithm, using the frequency domain Wiener filtering approach detailed in Section 5.2.2.
- Time-domain Wiener filters using the chosen signals are trained on a stretch of data preceding the event, and their performance is validated on a stretch of data following the event.
- If the regression is stable, the trained filters are applied to data during the time of the event, resulting in a GW signal with reduced noise which can then be further analyzed.

This procedure was validated through comparison of the parameter recovery of hardware injections, where a gravitational wave signal with known modeled source parameters was physically introduced into the instrument. This section describes the details of how the above steps were accomplished.

5.3.1 GREEDY RANKING

Due to the complexity of the instrument, and the large number of sensors and signals, it is impractical to examine all available signals without some amount of automation. In addition, multiple sensors may incidentally witness the same noise coupling, such as seismometers located near each other. Finally, the computational cost of training time-domain Wiener filters rises quickly as the number of signals is increased, which greatly expands the size of the matrices being inverted, and when the signals are strongly correlated, as this makes the cross correlation matrix nearly singular. These aspects motivated the creation of an algorithm to efficiently identify a small set of mostly uncorrelated witnesses with the greatest potential of efficacious noise regression.

A list of auxiliary signals is specified which is intended to reasonably cover the signals of potential interest across the interferometer subsystems, such as the various optic suspensions, photodiodes, environmental sensors, feedback control signals, etc., and typically has approximately 800 entries. Then, each signal is independently evaluated for potential coherence with the strain signal in the frequency domain over some frequency band

of interest. The signals are sorted by the integrated coherence over the band, and the top ranked signal is added to list of subtraction candidates. Next, each signal is evaluated again for coherence with the strain signal, but in combination with the first candidate, so that the resultant ranking reflects the signal whose inclusion provides the greatest increase in integrated MISO coherence over the band of interest — i.e. “greedy ranking”. Another possible metric would be to decreased the total signal power in the frequency range of interest, but this was found to bias the selection to channels that had incidental correlation to power line harmonics or other instrumental lines.

With this scheme, the inclusion of redundant signals is not favored, as their inclusion would add negligible marginal increase to the total MISO coherence. In this manner, the list of subtraction candidates grows, until no appreciable improvement is seen from the inclusion of additional signals — usually no more than five witnesses in the periods of time examined.

5.3.2 WIENER FILTER TRAINING AND EVALUATION

This reduction of the witness space from ≈ 800 to ≈ 5 makes the application of the time-domain Wiener filter feasible. Up to one hour each of training and validation data, subject to the interferometer status, is used to ensure that the filters are not over-fitting and reliably result in a decrease in noise power in the strain signal.

The high frequency noise performance in the detectors is generally well understood [5], consisting mostly of shot noise and harmonic lines that are not well witnessed by environmental sensors. Then, since most of the observed noise couplings occur at lower frequencies, the training data can safely be down-sampled to a manageable frequency to save computation time. The resultant Wiener filters will only be applicable for signals at this lowered sampling frequency, but the resultant prediction of the witness’ noise contribution can be up-sampled and subtracted from the full rate strain signal, with appropriate anti-imaging filtering.

As shown in Section 5.2.2, the time domain Wiener filter is optimal in the sense of reducing the RMS difference between the target and prediction. However, the unavoidable seismic noise below 10 Hz makes total RMS of the strain signal a poor metric. Thus, when training time domain Wiener filtering, we employ a frequency selective filter to the witness and target signals, to make a certain target frequency band the dominant contributor to the filtered signal RMS before training the Wiener filters. Since this weighting pre-filter and the trained filters are linear, their effects commute; this means the trained

filters are equally valid on the unweighted signals. In this manner, we are able to exert some frequency selectivity to the training of the Wiener filters. In the results that follow, when a target frequency band is chosen, the weighting filter implemented is: an eighth order Chebyshev type I high-pass filter with a characteristic frequency of the lower bound to mitigate the large amount of low frequency signal power, and a second order Butterworth low-pass filter at the upper frequency bound.

The performance of the trained filters on the validation data can be compared to the subtraction prediction from the MISO coherence of the validation data, which can give an indication of whether the coupling has changed over the timescale of the data sets and whether the filter response necessary to achieve the predicted subtraction is realizable as a causal FIR filter. An example of this comparison can be seen in Figure 5.6.

If the performance of the trained filters is satisfactory, they are then applied to the strain data during the time of interest when a GW event or event candidate was recorded to produce a new strain time series which should have increased fidelity to the gravitational wave strain signal incident on the instrument.

The parameter estimation analysis also makes a further method of subtraction validation possible for specific waveforms and instrument states. Excitations are periodically intentionally injected into the interferometers in a way that mimics an astrophysical gravitational wave. This is used to characterize and calibrate the instrument, as well as to test parts of the data acquisition and analysis pipelines [77]. These injection waveforms are calculated from some set of known source parameters, which is used to test the parameter estimation analyses. Thus, we can test whether this noise subtraction scheme is legitimately reducing unwanted technical noise without distorting the measured GW signals by performing the same subtraction, and checking that the resultant posterior parameter estimation distributions are consistent with those from the pre-subtraction strain signal are not significantly biased away from the known injected parameters. Tests for waveforms resembling GW₁₅₀₉₁₄ will be shown in Section 5.4.2.

5.4 NOISE SUBTRACTION APPLIED TO GW₁₅₀₉₁₄

The procedure outline above was carried out for both the Hanford and Livingston observatories for the black hole binary merger GW₁₅₀₉₁₄, the first confirmed observation of gravitational waves [6]. This section will detail the results of the individual steps taken to produce cleaned strain time series for both detectors and the resultant improvement in the astrophysical parameter estimation of the binary black hole system that generated the signal.

5.4.1 OBSERVED NOISE COUPLINGS

The time intervals chosen for analysis around GW₁₅₀₉₁₄ are shown in Table 5.1.

Table 5.1: Time intervals used for GW₁₅₀₉₁₄ noise subtraction analysis

	Start Time (UTC)	Duration
Training	Sep 14 9:06:23 2015	2600 s
Validation	Sep 14 9:50:53 2015	3600 s
Event	Sep 14 9:50:45 2015	—

HANFORD

The primary noise coupling that was identified by the greedy ranking was a known problem of excess vibration on the input optics periscope where the PSL laser is sent into the vacuum system. While this coupling was later manually mitigated in the hardware [78], the interferometric alignment sensors accurately witnessed the same vibration, leading to coherence with the strain channel. In addition, there remained some residual coupling of the short Michelson signal due to small errors in the online feed-forward subtraction. Exploring multiple frequency bands, no significant couplings were found above 400Hz, leading to the final choice of 20–400 Hz as the target region for the subtraction. The specific channel names are shown in Table 5.2, and the noise contributions are visualized in Figure 5.3.

LIVINGSTON

The prediction for noise subtraction from the Livingston data is more focused at lower frequencies, which is more useful for resolving the initial inspiral phase of a binary black

hole merger signal. At the time of GW150914, the background noise of the Livingston detector included a reasonable contribution from excess angular controls noise at low frequencies, due to imperfect centering of the resonant cavity mode on the test masses leading to a direct coupling of mirror angular motion to a phase shift on the carrier light. Thus, the angular control signals were identified by the greedy ranking as having significant contributions. Smaller contributions were identified from auxiliary cavity length and seismic isolation signals. A ground tilt sensor was also identified as a useful witness, but for a harmonic comb due to a electronics rack power supply rather than ground tilt. No significant noise couplings were observed over ≈ 70 Hz, leading to a target frequency band of 20–70 Hz. The specific channel names are shown in Table 5.2, and the noise contributions are visualized in Figure 5.4.

Table 5.2: Channels identified by the greedy ranking algorithm.

Channel	Description
Hanford	
H1 : IMC-WFS_B_I_YAW_OUT_DQ	Input Mode Cleaner Alignment
H1 : LSC-POP_A_RF45_Q_ERR_DQ	MICH length
Livingston	
L1 : SUS-ETMX_L3_LOCK_Y_IN1_DQ	ETMX yaw angular control
L1 : SUS-ITMX_L3_LOCK_P_IN1_DQ	ITMX pitch angular control
L1 : LSC-POP_A_RF45_I_ERR_DQ	PRC length
L1 : PEM-EY_TILT_VEA_FLOOR_T_DQ	Y-End Ground Tilt-meter
L1 : ISI-ITMY_ST1_BLND_Y_T240_CUR_IN1_DQ	ITMY Chamber Seismic Isolation Seismometer

5.4.2 VALIDATION

Once the greedy ranking had produced a list of subtraction candidates, time domain Wiener filters were trained, and their performance compared favorably to the MISO coherence subtraction prediction during the validation data set; see Figures 5.5 and 5.6. The mismatch in the low frequency subtraction in the Livingston data is likely due to the change in the angle to length coupling that arises when the beam spot position changes; since the beam spots were not under active stabilization at the time, they would slowly drift. Nevertheless, the trained filters are still able to significantly reduce the background




Figure 5.3: Results of Greedy Ranking Search for the Hanford aLIGO Observatory. The bottom subplot is a stacked visualization of the contributions of each chosen signal to the overall coherence of the set to the strain signal. I.e., the portion of the stack corresponding to a signal is the increase in coherence at that frequency that accompanies its inclusion in the set of subtraction witnesses.




Figure 5.4: Results of Greedy Ranking Search for the Livingston aLIGO Observatory. The bottom subplot is a stacked visualization of the contributions of each chosen signal to the overall coherence of the set to the strain signal. I.e., the portion of the stack corresponding to a signal is the increase in coherence at that frequency that accompanies its inclusion in the set of subtraction witnesses.




Figure 5.5: PSDs of H1 Wiener Filter Validation around GW₁₅₀₉₁₄. The cleaned data exhibits an increase from 281 Mpc to 285 Mpc in the angle-averaged inspiral range for a 30 M_{\odot} and 35 M_{\odot} BBH merger.

noise at low frequencies without injecting additional noise, which suggests that the filters have accurately captured real physical couplings without over-fitting.

As mentioned previously, hardware injections were used to verify that the noise subtraction did not introduce any distortion of the gravitational wave signal or bias in the parameter estimation. Conveniently for our purposes, a number of hardware injections were coincidentally performed at both aLIGO observatories not long after GW₁₅₀₉₁₄ with waveforms generated with source parameters based on initial estimates of the maximum likelihood parameters of GW₁₅₀₉₁₄ itself. The detectors were in a similar enough configuration as when GW₁₅₀₉₁₄ occurred that the channels listed above still captured the same couplings. Thus, to test the noise subtraction procedure, Wiener filters were calculated and applied to each hardware injection, whereafter comparisons in the results of the parameter estimation could be made. Figure 5.7 shows one result of these comparisons, where it can be seen that the posterior probability estimate of the system's chirp mass derived from the cleaned data is more centered on the injected value.




Figure 5.6: PSDs of L1 Wiener Filter Validation around GW₁₅₀₉₁₄. The cleaned data exhibits an increase from 216 Mpc to 244 Mpc in the angle-averaged inspiral range for a 30 M_{\odot} and 35 M_{\odot} BBH merger.

5.4.3 EFFECTS ON GW₁₅₀₉₁₄ PARAMETER ESTIMATION

Confident in the results of the validation tests, the trained Wiener filters were applied to the strain data during GW₁₅₀₉₁₄. Figure 5.8 shows a comparison of the cleaned and original time series, in which the low frequency noise in the Livingston data can be seen to differ. Figure 5.9 compares the resultant posterior probability densities for the black hole binary chirp mass and mass ratio, while Figure 5.10 shows a similar plot for the final black hole mass and spin. Table 5.3 summarizes the parameter estimation results of the cleaned GW₁₅₀₉₁₄ data, and compares it with the results from the original recorded data.

Broadly speaking, there is a decrease in the span of the parameter confidence intervals commensurate with the increase of signal SNR in the cleaned data, and the median values are consistent with the original data. Given that the majority of the SNR improvement came from the low frequency noise of the L1 interferometer, it is not surprising that the chirp mass estimate is particularly improved, as its value is the main influence in the lower frequency inspiral phase of the GW signal. The source redshift and distance are gener-




Figure 5.7: Comparison of parameter estimation posterior distributions for a GW₁₅₀₉₁₄-like hardware injection with known source parameters. The 2D colored outlines indicate the contours of the 90 % credible regions of the two PDF estimates.

ally functions of timing and triangulation, and not strongly SNR limited with only two running detectors, so it is not surprising that those parameters are largely unaffected by the subtraction.

Table 5.3: Comparison of GW150914 Parameter Estimation Results, using the IMRPhe-nomPV2 waveform models [74]. Values reported are the medians of the 1D posterior probability distributions as well as the symmetric 90 % credible interval. Δ_{CI} represents the relative change in the span of the 90 % confidence interval; where a positive value indicates a narrower interval. The stated figures for the original data differs slightly from the values given in ABBOTT et al. [71], as they do not include the uncertainty model of the interferometer calibration used at that time, which is now known to be a significant overestimate. The inclusion of this calibration model obscures the reduced uncertainty from the noise subtraction.

Quantity	Original	Cleaned	Δ_{CI}
LI Optimal SNR	13.9	14.5	
HI Optimal SNR	20.9	22.3	
Optimal Network SNR	25.1	26.6	
Detector-frame Total Mass M/M_{\odot}	$71.2^{+3.2}_{-3.2}$	$71.9^{+3.0}_{-2.6}$	12.2%
Detector-frame Chirp Mass \mathcal{M}/M_{\odot}	$30.8^{+1.4}_{-1.5}$	$31.1^{+1.3}_{-1.2}$	14.7%
Detector-frame Primary Mass m_1/M_{\odot}	$38.4^{+5.2}_{-3.2}$	$38.4^{+4.7}_{-2.8}$	10.5%
Detector-frame Secondary Mass m_2/M_{\odot}	$32.7^{+3.1}_{-4.8}$	$33.5^{+2.6}_{-4.2}$	14.0%
Detector-frame Final Mass M_f/M_{\odot}	$67.9^{+2.9}_{-2.9}$	$68.5^{+2.7}_{-2.4}$	11.5%
Source-frame Total Mass $M^{\text{source}}/M_{\odot}$	$65.3^{+3.3}_{-3.0}$	$65.4^{+3.3}_{-2.6}$	7.1%
Source-frame Chirp Mass $\mathcal{M}^{\text{source}}/M_{\odot}$	$28.2^{+1.5}_{-1.4}$	$28.3^{+1.4}_{-1.1}$	9.1%
Source-frame Primary Mass $m_1^{\text{source}}/M_{\odot}$	$35.4^{+4.8}_{-3.0}$	$35.0^{+4.4}_{-2.6}$	11.1%
Source-frame Secondary Mass $m_2^{\text{source}}/M_{\odot}$	$29.9^{+3.0}_{-4.4}$	$30.4^{+2.6}_{-3.7}$	13.8%
Source-frame Final Mass $M_f^{\text{source}}/M_{\odot}$	$62.2^{+3.1}_{-2.7}$	$62.3^{+3.1}_{-2.4}$	7.0%
Mass Ratio q	$0.85^{+0.13}_{-0.20}$	$0.88^{+0.11}_{-0.19}$	11.9%
Effective Inspiral Spin Parameter χ_{eff}	$-0.043^{+0.101}_{-0.123}$	$-0.022^{+0.096}_{-0.096}$	14.1%
Primary Spin Magnitude a_1	$0.288^{+0.539}_{-0.263}$	$0.318^{+0.502}_{-0.291}$	1.1%
Secondary Spin Magnitude a_2	$0.335^{+0.554}_{-0.304}$	$0.295^{+0.580}_{-0.271}$	0.9%
Final Spin a_f	$0.669^{+0.035}_{-0.048}$	$0.676^{+0.033}_{-0.037}$	16.4%
Luminosity Distance D_L/Mpc	433^{+129}_{-160}	483^{+104}_{-186}	-0.3%
Source Redshift z	$0.092^{+0.025}_{-0.033}$	$0.102^{+0.020}_{-0.038}$	0.3%




Figure 5.8: Comparison of Filtered 150914 time series. The signal has been band-passed at 25–400 Hz and notched at multiple instrumental line frequencies.




Figure 5.9: Improvement of GW₁₅₀₉₁₄ detector-frame component mass posterior PDFs. The mass ratio q is defined as $\frac{m_1}{m_2}$ where $m_1 < m_2$ by convention; this leads to the sharp cut-off at $q = 1$. The 2D colored outlines indicate the contours of the 90% credible regions of the two PDF estimates.




Figure 5.10: Improvement of GW_{150914} final black hole parameter posterior PDFs. The 2D colored outlines indicate the contours of the 90 % credible regions of the two PDF estimates.

5.5 NONLINEAR NOISE AND REGRESSION

Naturally, linear couplings of external disturbances into the gravitational wave strain readout are a subset of the full dynamics of the detectors. There are many known nonlinear coupling mechanisms, and surely more exists which have not been fully accounted for. The functional forms can vary greatly, and even modest uncertainty in the parameters involved can make it impractical to reconstruct and regress the unwanted noise.

Recent popular successes in the use of Artificial Neural Networks (ANN) for classification and regression problems has prompted work into developing techniques for performing nonlinear regression in the interferometer data that does not require precise a priori knowledge of all of the system parameters.

In this section, we will detail some particular nonlinear noise sources that are relevant for the current performance of the aLIGO detectors, and the ongoing efforts to develop automated regression techniques that can be put to use for off-line noise subtraction, in a manner similar to the previous section.

5.5.1 KNOWN NONLINEAR NOISE SOURCES

SCATTERED LIGHT

It has been known for some time that the existence of scattered light within the vacuum envelope of a gravitational wave interferometer can degrade the sensitivity [79]. Optical elements of the instrument will invariably have surface defects, which scatter some small amount of light out of the main resonant mode of the interferometer. This light may interact with elements which are moving much more than the carefully isolated test masses, such as the vacuum chamber walls or the active seismic isolation tables, and its phase and amplitude thereby modulated by this motion. If, in turn, the moving element reflects some of this scattered light back into the resonant light in the primary interferometer mode, they may interfere in such a way that the modulation due to the moving elements is imprinted on the strain readout. If the amplitude of the scattering source's motion is larger than the laser wavelength, the modulation will be strongly nonlinear, wrapping over 2π . This has been a major concern in the design in construction of the LIGO and VIRGO instruments [80, 81].

The precise nature of the coupling of the scattered light into the gravitational wave strain readout depends on where in the optical system the scattering is taking place, and the instrument's response to amplitude and phase modulation to the light at that point. In [81],

the authors show that the phase and amplitude noises can generically be written as

$$n_\phi(t) = \sqrt{\frac{P_s}{P_{in}}} \sin(\phi_0 + \phi_r(t)) \quad (5.24)$$

$$\frac{\delta P}{P}(t) = \sqrt{\frac{P_s}{P_{in}}} \cos(\phi_0 + \phi_r(t)) \quad (5.25)$$

where P_{in} is the power in the interferometer field, P_s is the power in the field that has scattered out and back into the main field, ϕ_0 is the static (or average) propagation phase traversed between the interferometer and the scattering object, and $\phi_r(t)$ encodes the phase of the field due to motion of the scattering object via

$$\phi_r(t) = \frac{4\pi}{\lambda} z_r(t) \quad (5.26)$$

(This assumes that $\frac{P_s}{P_{in}} \ll 1$.)

In the case of aLIGO, the functional form of the scattering around the Output Mode Cleaner (OMC) is considered in MARTYNOV [13]. Here, light that scatters off of the HAM6 chamber walls or other nearby elements will interfere with the static field present on the DC readout photodiodes in a phase sensitive manner. ϕ_r in this case may contain low frequency motion on the order of multiple wavelengths corresponding to the motion of the OMC suspension and high frequency motion of the chamber wall smaller than the wavelength. Thus, the noise may take on the general form of

$$\phi_r(t) = \frac{4\pi}{\lambda} (L_{lf}(t) + L_{hf}(t)) := \phi_{lf}(t) + \phi_{hf}(t) \quad (5.27)$$

$$\rightarrow n(t) \propto \cos [e^{\phi_{lf}(t)} (1 + \phi_{hf}(t))] \quad (5.28)$$

The proportionality of this functional form furthermore depends on P_s , which may, itself, vary with the low frequency motion of suspended elements. Through measurement of the motion of various interferometer elements, via seismic isolation and suspension sensors and accelerometers on the vacuum chambers, it may in principle be possible to reconstruct this signal. However, quantities such as ϕ_0 and P_s are difficult to know, and there may be some frequency dependent relationship between motion sensors and the actual motion causing the light modulation. This complicates regression significantly.

BILINEAR ANGULAR NOISE

Precise alignment sensing and control (ASC) of the interferometer mirrors is critical for sustained operation of aLIGO observatories; seismic and tidal forces are continuously

changing, and the active seismic isolation platforms move relative to each other. While interferometric wavefront sensors are used to control and stabilize important alignment degrees of freedom [13, 19, 82], there will always be some degree of residual angular motion. Phase shifts of the laser light within the interferometer arm cavities are caused by motion of the test mass mirror surface that is perpendicular to the incident beam, and are apparent in the gravitational wave strain readout. Thus, any angular motion of a test mass will impart a phase shift on the light as long as the beam spot is not exactly positioned on the center of rotation; the offset acts as a lever arm causing a so-called “bilinear” coupling of the spot position and mirror angle to perceived length fluctuation.

$$\delta x = x_{\text{spot}} \delta \theta \quad (5.29)$$

This means that if the beam spot has a stable offset from the center of rotation there is a direct linear coupling of angular motion to the strain readout.

The relative strength of the test mass actuators is adjusted in an attempt to co-locate the center of the angular control torque with the beam spots to cancel the linear coupling of ASC control signals to the strain readout. However, this procedure is of finite precision, and uncontrolled degrees of freedom may cause the average beam spot position to acquire a constant offset from the center of actuation; this is the cause of the linear couplings seen in previous sections. However, even if the beam spot is centered on average, Equation 5.29 will be non-zero due to the residual motion of the beam spot and the necessary control torque.

Using bilinear combinations of signals has previously proven successful in predicting a subset of transient glitches [83], but a time series regression has not been attempted until recently. While the functional form in Equation 5.29 is simple, there are some unknowns factors between the true physical values of the relevant quantities and the signals at our disposal. While many sensors monitor beam spot motion, the true zero point corresponding to the center of rotation is not inherently known. In addition, the known angular control signals, which are applied to intermediate stages of the mirror suspensions, traverse an opto-mechanical transfer function that depends on the mechanical properties of the suspension and the power dependent angular radiation pressure dynamics [84]. Finally, there will be a bilinear noise contribution from each test mass, but the beam spot motion on all four is correlated through the radiation pressure effects, so any successful regression routine will have to be robust against these degeneracies.

Mock data generation codes were written in order to facilitate repeatable testing of non-linear regression systems, with background noise levels corresponding to recent aLIGO




Figure 5.11: Amplitude spectral density of mock scatter and bilinear noises. A background corresponding to the quantum noise levels is included in all cases, whose noise power cannot be reduced through regression.

sensitivity. The unknown parameters involved were chosen in such a way to make the resultant total noise spectrum consistent with observations. Representative noise spectra can be seen in Figure 5.11.

5.5.2 DEEP NEURAL NETWORKS FOR REGRESSION

Artificial neural networks (ANNs), are computational structures inspired by the interconnectivity of neurons in biological brains. The multitude of interconnected “neurons” represent weighted sums of inputs applied to nonlinear “activation functions” (see Figure 5.12 for a schematic representation). Each neuron has many free parameters associated, making the network a high-dimensional function of its inputs. For some time, it has been known that these complex structures are capable of nonlinear function approximation, if there are a suitable number of neurons [85]. Thus, it is natural to imagine using numerical optimization techniques to try and fit, or “learn,” the free parameters such that a desired functional output given some known inputs can be reproduced; a popular technique for ANNs is the back-propagation method [86], in which the multi-dimensional gradient of




Figure 5.12: Schematic representation of the neural network used for nonlinear regression. The inputs consist of a window of past samples of (potentially multiple) witness channels, with a target output of single strain channel sample y_n . Each circle represents a “neuron”, in which the sum of the neuron inputs i_n multiplied by weighting coefficients w_n with a bias b is used as the input to a nonlinear activation function f to calculate the neuron’s output. The weights and biases are varied and optimized to minimize the RMS difference between the network outputs and the recorded strain signal.

the loss function (or estimation error) with respect to all of the free network parameters is efficiently calculated. Using an ANN for function approximation or regression in this manner falls under the category of “supervised learning,” as the desired output for a set of inputs is known ahead of time.

Recently, the use of ANNs for supervised learning tasks has seen an explosion of success and popularity; so-called “deep” learning uses many intermediate “hidden” neuron layers that require computing power that was not readily available when ANNs were initially considered. However, the vast freedom in choices of network layer number, neuron number, activation function, and a multitude of “hyper-parameters” can make it challenging to construct an ANN that converges in finite time, if at all.

It is hoped that deep ANNs are suitable for the regression of problematic nonlinear noise sources in LIGO data, as described in Section 5.5.1. This work is ongoing at the time of writing, but preliminary results with mocked strain data is encouraging.

The simplest model one could think of would be to use a network where the inputs are individual samples from each candidate signals, and the output is individual strain samples. However, since there are nontrivial frequency-domain dynamics at play in these couplings, this would not be enough information to predict real noise couplings. Equivalently, one can imagine that single input samples are not enough to estimate derivatives or integrals that may represent relevant conjugate dynamical variables that are necessary for accurate regression. Thus, similar to the FIR filter kernel nature of linear time-domain Wiener filters, we have used ANN structures where the input vector is a concatenated series of finite length time windows of the past values of the various candidate signals.

This also facilitates the mixture of signals with different characteristic time scales, as there is no strict requirement that all of the input signals must be at the same sampling frequency. This is useful for the bilinear noise case, in which the relevant physical beam spot motion is only sensed from 0.1–0.3 Hz, and the relevant angular motion is in the audio band.

As with any regression routine, separate data sets are used for training and validation of the predictions. Similar to the time-domain Wiener filter, the quantity being optimized is the mean squared error (MSE) between the target (strain data) and the prediction (ANN output). Thus, it is important to define appropriate frequency selective pre-filters to ensure that the majority of the MSE comes from the frequency band of interest. Typical activation functions and initializers for starting ANN parameters also generally expect the input and target data to be roughly zero mean, unity variance quantities. However, it is not particularly challenging to perform the normalization and pre-filtering in an easily invertible manner, such that the trained network prediction can be converted back into real physical units.

Figure 5.13 shows the loss evolution for the training of an ANN to regress the bilinear noise from four independent beam spot and angular motion noise contributions, using witness signals that contain independent sensing noise terms. It can be seen that the loss on the training and validation sets is monotonically decreasing, though they do not reach the minimum level that would be achieved by perfect regression without noise injection.

The network receives 0.25 s of each witness signal to predict each individual strain sample; the angular control and strain signals are sampled at 256 Hz while the beam spot motion signals are down-sampled to 32 Hz. Seven fully connected neuron layers were used, the input layer having an equal number of neurons as the input channel, and subsequent layers according to $n_i = 2^{7-i}$, such that the final output layer consists of a single neuron. The final layer used a linear activation function, whereas all others used the “exponential




Figure 5.13: Mean squared error in Neural Network predictions for the bilinear noise coupling of four uncorrelated beam spots as a function of training epoch. The regression performance on test data that has not been used for training consistently decreases as training progresses, indicating that the network is not over-fitting to to the training data.

linear unit” as the activation function, defined as:

$$f(x) = \begin{cases} x & \text{if } x \geq 0 \\ \exp(x) - 1 & \text{otherwise} \end{cases} \quad (5.30)$$

Figure 5.14 shows the spectral density of the achieved noise subtraction, which reaches factors of 10–20 in the noisiest regime, though some injection is present above 20 Hz. This injection varies considerably with the specific cost pre-filter being used.

These preliminary results are certainly encouraging, and further work into optimization of ANN network hyper-parameters is ongoing, in addition to the application of ANN regression networks to real interferometer data. Additionally, alternate network units are being explored, such as recurrent LSTM cells [87], which use previous time steps’ outputs as inputs thereby having a form of memory that may make the time-window scheme described above unnecessary.




Figure 5.14: (Top: Amplitude spectral densities showing the successful noise subtraction of bilinear angular noise from the mocked strain signal by a trained ANN on unseen validation data. Bottom: Coherence of the nonlinear ANN noise prediction compared with the MISO linear coherence of the target with the entire set of witnesses. This indicates that the ANN is adding new regression capabilities.

5.6 FUTURE WORK

Going forward, it is evident that noise regression efforts are worth pursuing further. Given the great cost associated with the design, construction, commissioning, and analysis of the LIGO interferometers, being able to reliably improve the data quality through semi-automated processes will ensure a greater science return on the investment of the scientific community and the public.

Future avenues of application could be to perform training of linear subtraction filters in a low-latency manner, such that a cleaned strain time series could be consistently available not long after the raw data is recorded. Running regressions in a constant online manner would also facilitate the use of cleaned data in the LIGO search pipelines, which require the use of the entire run's data to properly estimate the statistical significance of events over the background.

The prospects of re-running searches on previous data would be especially promising if successful nonlinear regression routines are developed, as scattering noise was known to be a significant hindrance to the sensitivity of the aLIGO detectors during the first observing run. In this case, it may be possible for event candidates like LVT₁₅₁₀₁₂ could be promoted to fully confident detections.

APPENDIX A

DELAY-LINE FREQUENCY DISCRIMINATOR ANALYSIS

Here, we will analyze the response and sensitivity of the Delay-line Frequency Discriminator (DFD) scheme used in the 40m ALS system in a manner analogous to the treatment in *Phase Noise Characterization Of Microwave Oscillators* [88], with modifications reflecting the dual quadrature demodulation and digital phase tracker employed at the 40m.

We begin with the RF signal with nominal beat frequency Ω_B and some time-varying phase fluctuations $\phi(t)$:

$$V_B(t) = V_0 \cos(\Omega_B t + \phi(t)) \quad (\text{A.1})$$

We are ultimately concerned with the fluctuations of the beat frequency:

$$\Delta F_B(t) := \frac{1}{2\pi} \frac{d\phi}{dt} \quad (\text{A.2})$$

Here, the capital F is used to distinguish the RF beat frequency from the later use of f for the Fourier frequency of its variations.

$V_B(t)$ is sent into an RF power splitter. One of the split components is then delayed by a time τ before mixed with the non-delayed signal in two mixers 90° out of phase with each other, labelled I and Q. So, we can write the L and R inputs to the mixers as:

$$V_L(t) = V_L \cos(\Omega_B t + \phi(t)) \quad (\text{A.3})$$

$$V_{R,I}(t) = A V_L \cos(\Omega_B(t - \tau) + \phi(t - \tau)) \quad (\text{A.4})$$

$$V_{R,Q}(t) = A V_L \sin(\Omega_B(t - \tau) + \phi(t - \tau)) \quad (\text{A.5})$$




Figure A.1: Schematic of Delay-line Frequency Discriminator used to track optical beat frequency fluctuations in the 40m prototype ALS system.

Here, A refers to the power lost by transmission through the coaxial delay line cabling.

The signals are then mixed in an RF mixer with conversion efficiency ϵ . After low-passing the mixer IF outputs to eliminate the higher harmonics, we obtain the resultant demodulated signals $V_I(t)$ and $V_Q(t)$:

$$V_I(t) = A\epsilon V_L \cos [\Omega_B \tau + \phi(t) - \phi(t - \tau)] \quad (\text{A.6})$$

$$V_Q(t) = -A\epsilon V_L \sin [\Omega_B \tau + \phi(t) - \phi(t - \tau)] \quad (\text{A.7})$$

We may write this in terms of ΔF_B using

$$\phi(t) - \phi(t - \tau) = \int_{t-\tau}^t dt \, 2\pi \Delta F_B(t) \quad (\text{A.8})$$

If we assume that $\Delta F_B(t)$ is band-limited such that the maximum frequency of interest is much lower than τ^{-1} , we may treat $\Delta F_B(t)$ as roughly constant over timescales of τ .¹ This implies

$$\phi(t) - \phi(t - \tau) = 2\pi\tau \Delta F_B(t) \quad (\text{A.9})$$

$$V_I(t) = A\epsilon V_L \cos [\Omega_B \tau + 2\pi\tau \Delta F_B(t)] \quad (\text{A.10})$$

$$V_Q(t) = -A\epsilon V_L \sin [\Omega_B \tau + 2\pi\tau \Delta F_B(t)] \quad (\text{A.11})$$

These signals are whitened and digitized via ADC into the digital controls system.

In *Phase Noise Characterization Of Microwave Oscillators* [88], the authors discuss tuning the delay line length such that $\Omega_B \tau = (n + \frac{1}{2})\pi$, which results in a linear transduction of ΔF_B to V_I when $\tau|\Delta F_B| \ll 1$. However, when the interferometer arm cavities are uncontrolled, the signal may be wrapping through multiple fringes. So, we employ a *phase*

¹If this assumption is not valid, analysis in the frequency domain is still viable, as described in *Phase Noise Characterization Of Microwave Oscillators* [88]

tracker servo in which we digitally rotate V_I and V_Q by a time varying angle $\Theta(t)$ in an attempt to minimize the amplitude of the new I signal:

$$V_I'(t) = \sin [\Theta(t)] V_I(t) + \cos [\Theta(t)] V_Q(t) \quad (\text{A.12})$$

Imagining that through some averaging process, we have determined Ω_B well enough to set $\Theta = \Omega_B \tau + \theta(t)$, we have:

$$V_I'(t) = A\epsilon V_L \sin [\theta(t) + 2\pi\tau\Delta F_B(t)] \quad (\text{A.13})$$

We can imagine that due to the sin function, V_I' will occasionally exhibit a zero crossing, around which we can linearise it as

$$V_I'(t) \approx A\epsilon V_L [\theta(t) + 2\pi\tau\Delta F_B(t)] \quad (\text{A.14})$$

At this time, we can close a purely digital feedback loop around the rotation, using V_I' as the error signal; this is the phase-tracker servo. For convenience, let us switch to the Laplace domain, resulting in

$$\tilde{\theta}(s) = K(s)\tilde{V}_I'(s) \quad (\text{A.15})$$

$$\Rightarrow (K^{-1} - A\epsilon V_L)\tilde{\theta} = A\epsilon V_L 2\pi\tau\tilde{\Delta F}_B \quad (\text{A.16})$$

$$\Rightarrow \tilde{\theta} = \frac{G}{1 - G} 2\pi\tau\tilde{\Delta F}_B \quad (\text{A.17})$$

$$\text{where } G(s) := A\epsilon V_L K(s) \quad (\text{A.18})$$

Equation [A.17](#) is the usual definition of a linear feedback control signal in the Laplace domain, with $G(s)$ as the open loop gain transfer function. Practically, a digital integrator ($K(s) = \frac{K_0}{s}$) is sufficient, and the control bandwidth is limited only by the time step of the digital control system. A unity gain frequency of 2 kHz is typical at the 40m prototype, defined by

$$K_0 = \frac{2\pi f_{UGF}}{A\epsilon V_L} \quad (\text{A.19})$$

While we considered a hypothetical zero-crossing, the phase tracker is in practice easily activated at any time. As long as the control bandwidth is sufficient, the phase tracker will suppress fluctuations of $V_I'(t)$, keeping the linearity approximation legitimate.

At a few hundred Hz, where we seek to use the DFD output as a length control signal, $1 \ll G$, so $V_I' \approx 0$ and

$$\theta(t) = -2\pi\tau\Delta F_B(t) \quad (\text{A.20})$$

Thus, the phase tracker control signal provides us with an easily calibrated linear signal for the beat frequency fluctuations, even when these fluctuations are very large; the only constraint is that the dynamics of $\Delta F_B(t)$ are at frequencies below $f = \tau^{-1}$, and the digital control system runs fast enough for stable feedback control.

Equation A.20 seems to imply that the sensitivity of the phase tracker depends only on the length of the delay, however we have not yet considered the influence of noise. For instance, V_I' will inherit any voltage noise present in V_I and V_Q , following Equation A.12, which will manifest in the phase tracker control signal in the Laplace domain via:

$$\tilde{\theta} = \frac{G}{1-G} \left[2\pi\tau\tilde{\Delta F}_B + \frac{\delta\tilde{V}}{A\epsilon V_L} \right] \quad (\text{A.21})$$

Thus, the signal to noise ratio of the beat note frequency fluctuations to voltage noise out of the mixer will follow

$$\text{SNR} = 2\pi\tau A\epsilon V_L \quad (\text{A.22})$$

There are two terms with delay-line length dependence here: τ and A . $\tau = L/v$, where v is usually about $\frac{2}{3}c$ in coaxial cable. Commercial coaxial cables usually specify their attenuation in dB of power lost per unit length Z , following $A = 10^{-LZ/20}$. So, we see that there are two competing terms that determine the SNR of the DFD; we can find the maximum as follows:

$$\text{SNR} \propto L10^{-\frac{LZ}{20}} \quad (\text{A.23})$$

$$\Rightarrow \frac{\partial}{\partial L} \text{SNR} \propto 10^{-LZ} - \frac{\ln(10)}{20} LZ10^{-LZ} \quad (\text{A.24})$$

$$\Rightarrow (LZ)_{\max} = \frac{20}{\ln(10)} \approx 8.7 \text{ dB} \quad (\text{A.25})$$

Thus result means that, all else being equal, the maximum SNR is governed principally by the loss of the delay line cable and not its length, and achieved at a total delay line attenuation of 8.7 dB

BIBLIOGRAPHY

- ¹A. EINSTEIN, “Die Grundlage der allgemeinen Relativitätstheorie”, *Annalen der Physik* 354, 769–822 (1916) [10.1002/andp.19163540702](#).
- ²C. MISNER, K. THORNE, and J. WHEELER, *Gravitation* (W. H. Freeman, 1973).
- ³A. EINSTEIN, “Näherungsweise Integration der Feldgleichungen der Gravitation”, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, 688–696 (1916).
- ⁴J. AASI et al., “Advanced LIGO”, *Class. Quantum Grav.* 32, 074001 (2015) [10.1088/0264-9381/32/7/074001](#).
- ⁵D. V. MARTYNOV et al., “Sensitivity of the advanced ligo detectors at the beginning of gravitational wave astronomy”, *Phys. Rev. D* 93, 112004 (2016) [10.1103/PhysRevD.93.112004](#).
- ⁶B. P. ABBOTT et al., “Observation of gravitational waves from a binary black hole merger”, *Phys. Rev. Lett.* 116, 061102 (2016) [10.1103/PhysRevLett.116.061102](#).
- ⁷P. R. SAULSON, *Fundamentals of interferometric gravitational wave detectors* (World scientific, 2017).
- ⁸J. H. TAYLOR, L. A. FOWLER, and P. M. McCULLOCH, “Measurements of general relativistic effects in the binary pulsar psr1913 + 16”, *Nature* 277, 437–440 (1979) [10.1038/277437a0](#).
- ⁹J. MIZUNO et al., “Frequency response of michelson- and sagnac-based interferometers”, *Optics Communications* 138, 383–393 (1997) [10.1016/S0030-4018\(97\)00056-4](#).
- ¹⁰R. L. WARD, “Length sensing and control of a prototype advanced interferometric gravitational wave detector”, PhD thesis (California Institute of Technology, 2010).
- ¹¹R. X. ADHIKARI, “Gravitational radiation detection with laser interferometry”, *Rev. Mod. Phys.* 86, 121–151 (2014) [10.1103/RevModPhys.86.121](#).
- ¹²A. STALEY et al., “Achieving resonance in the advanced ligo gravitational-wave interferometer”, *Classical and Quantum Gravity* 31, 245010 (2014) [10.1088/0264-9381/31/24/245010](#).

- ¹³D. MARTYNOV, “Lock acquisition and sensitivity analysis of advanced ligo interferometers”, PhD thesis (California Institute of Technology, 2015), [10.7907/Z9Q81B1F](#).
- ¹⁴E. HALL, “Long-baseline laser interferometry for the detection of binary black-hole mergers”, PhD thesis (California Institute of Technology, 2017), [10.7907/Z9PG1PQ9](#).
- ¹⁵K. IZUMI and D. SIGG, “Advanced ligo: length sensing and control in a dual recycled interferometric gravitational wave antenna”, *Classical and Quantum Gravity* **34**, 015001 (2017) [10.1088/0264-9381/34/1/015001](#).
- ¹⁶B. P. ABBOTT et al., “LIGO: the laser interferometer gravitational-wave observatory”, *Reports on Progress in Physics* **72**, 076901 (2009) [10.1088/0034-4885/72/7/076901](#).
- ¹⁷K. IZUMI and D. SIGG, “Frequency response of the aLIGO interferometer: part 1”, *LIGO Document T1500325* (2016).
- ¹⁸J. MIZUNO, “Comparison of optical configurations for laser-interferometric gravitational-wave detectors”, PhD thesis (Universität Hannover, 1995), [11858/00-001M-0000-002B-0AD7-7](#).
- ¹⁹K. L. DOOLEY, “Design and performance of high laser power interferometers for gravitational-wave detection”, PhD thesis (University of Florida, 2011).
- ²⁰R. W. P. DREVER et al., “Laser phase and frequency stabilization using an optical resonator”, *Applied Physics B* **31**, 97–105 (1983) [10.1007/BF00702605](#).
- ²¹M. W. REGEHR, “Signal extraction and control for an interferometric gravitational wave detector”, PhD thesis (California Institute of Technology, 1995).
- ²²A. J. MULLAVEY et al., “Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers”, *Opt. Express* **20**, 81–89 (2012) [10.1364/OE.20.000081](#).
- ²³K. IZUMI et al., “Multicolor cavity metrology”, *J. Opt. Soc. Am. A* **29**, 2092–2103 (2012) [10.1364/JOSAA.29.002092](#).
- ²⁴K. IZUMI, “Multi-color interferometry for lock acquisition of laser interferometric gravitational-wave detectors”, PhD thesis (University of Tokyo, 2012).
- ²⁵K. ARAI, “Robust extraction of control signals for power-recycled interferometric gravitational-wave detectors”, PhD thesis (University of Tokyo, 2001).
- ²⁶S. M. ASTON et al., “Update on quadruple suspension design for advanced ligo”, *Class. Quantum Grav.* **29**, 235004 (2009) [10.1088/0264-9381/29/23/235004](#).
- ²⁷F. MATICHARD et al., “Seismic isolation of advanced LIGO: review of strategy, instrumentation and performance”, *Class. Quantum Grav.* **32**, 185003 (2015) [10.1088/0264-9381/32/18/185003](#).

- ²⁸S. BRACCINI et al., “The maraging-steel blades of the virgo super attenuator”, *Meas. Sci. Technol.* **11**, 467 (2000) [10.1088/0957-0233/11/5/304](https://doi.org/10.1088/0957-0233/11/5/304).
- ²⁹G. D. HAMMOND et al., “Reducing the suspension thermal noise of advanced gravitational wave detectors”, *Class. Quantum Grav.* **29**, 124009 (2012) [10.1088/0264-9381/29/12/124009](https://doi.org/10.1088/0264-9381/29/12/124009).
- ³⁰W. BLUM, P. EISENLOHR, and F. BREUTINGER, “Understanding creep — a review”, *Metallurgical and Materials Transactions A* **33**, 291 (2002) [10.1007/s11661-002-0090-9](https://doi.org/10.1007/s11661-002-0090-9).
- ³¹M. BECCARIA et al., “The creep problem in the virgo suspensions - a possible solution using maraging steel”, *Nucl. Instrum. Meth. A* **404**, 455 (1998) [10.1016/S0168-9002\(97\)01123-6](https://doi.org/10.1016/S0168-9002(97)01123-6).
- ³²G. CAGNOLI et al., “Mechanical shot noise induced by creep in suspension devices”, *Phys. Lett. A* **237**, 21 (1997) [10.1016/S0375-9601\(97\)00710-X](https://doi.org/10.1016/S0375-9601(97)00710-X).
- ³³A. AGEEV et al., “Excess noise in the steel suspension wires for the laser gravitational wave detector”, *Phys. Lett. A* **246**, 479 (1997) [10.1016/S0375-9601\(98\)00526-X](https://doi.org/10.1016/S0375-9601(98)00526-X).
- ³⁴S. BRACCINI et al., “Monitoring the acoustic emission of the blades of the mirror suspension for a gravitational wave interferometer”, *Physics Letters A* **301**, 389–397 (2002) [https://doi.org/10.1016/S0375-9601\(02\)00991-X](https://doi.org/10.1016/S0375-9601(02)00991-X).
- ³⁵R. MAASS et al., “Slip statistics of dislocation avalanches under different loading modes”, *Phys. Rev. E* **91**, 042403 (2015) [10.1103/PhysRevE.91.042403](https://doi.org/10.1103/PhysRevE.91.042403).
- ³⁶N. FRIEDMAN et al., “Statistics of dislocation slip avalanches in nanosized single crystals show tuned critical behavior predicted by a simple mean field model”, *PRL* **109**, 095507 (2012) [10.1103/PhysRevLett.109.095507](https://doi.org/10.1103/PhysRevLett.109.095507).
- ³⁷A. MICHELSON and E. MORLEY, “On the relative motion of the earth and the luminiferous ether”, *American Journal of Science* **34**, 333 (1887) [10.2475/ajs.s3-22.128.120](https://doi.org/10.2475/ajs.s3-22.128.120).
- ³⁸J. BECHHOEFER, “Feedback for physicists: a tutorial essay on control”, *Rev. Mod. Phys.* **77**, 783 (2005) [10.1103/RevModPhys.77.783](https://doi.org/10.1103/RevModPhys.77.783).
- ³⁹R. WEISS and D. KELLEY, “Collection of reports on barkhausen noise”, *LIGO Document T0900061* (2009).
- ⁴⁰G. VAJENTE, E. A. QUINTERO, and X. NI, “Upper limits on crackling noise from the crackler apparatus”, *LIGO document T1400752* (2014).
- ⁴¹F. BONDU et al., “Ultrahigh-spectral-purity laser for the virgo experiment”, *Opt. Lett.* **21**, 582 (1996) [10.1364/OL.21.000582](https://doi.org/10.1364/OL.21.000582).
- ⁴²T. ACCADIA et al., “Noise from scattered light in virgo’s second science run data”, *Class. Quantum Grav.* **27**, 194011 (2010) [10.1088/0264-9381/27/19/194011](https://doi.org/10.1088/0264-9381/27/19/194011).
- ⁴³G. VAJENTE, “Crackling noise in advanced gravitational wave detectors: a model of the steel cantilevers used in the test mass suspensions”, *Phys. Rev. D* **96**, 022003 (2017) [10.1103/PhysRevD.96.022003](https://doi.org/10.1103/PhysRevD.96.022003).

- ⁴⁴J. C. DRIGGERS, “Noise cancellation for gravitational wave detectors”, PhD thesis (California Institute of Technology, 2015), [10.7907/Z94F1NNP](https://doi.org/10.7907/Z94F1NNP).
- ⁴⁵R. DEROSA et al., “Global feed-forward vibration isolation in a km scale interferometer”, *Classical and Quantum Gravity* **29**, 215008 (2012) [10.1088/0264-9381/29/21/215008](https://doi.org/10.1088/0264-9381/29/21/215008).
- ⁴⁶K. IZUMI, D. SIGG, and L. BARSOTTI, “Self-amplified lock of an ultra-narrow linewidth optical cavity”, *Opt. Lett.* **39**, 5285–5288 (2014) [10.1364/OL.39.005285](https://doi.org/10.1364/OL.39.005285).
- ⁴⁷M. RAKHMANOV et al., “Dynamical properties of LIGO single loop suspended mirrors”, *LIGO Document T000134* (2000).
- ⁴⁸T. ISOGAI et al., “Loss in long-storage-time optical cavities”, *Opt. Express* **21**, 30114–30125 (2013) [10.1364/OE.21.030114](https://doi.org/10.1364/OE.21.030114).
- ⁴⁹J. EICHHOLZ, *40m ELOG 13299*, (Sept. 2017) <https://nodus.ligo.caltech.edu:8081/40m/13299>.
- ⁵⁰G. VENUGOPALAN, *40m ELOG 12586*, (Oct. 2016) <https://nodus.ligo.caltech.edu:8081/40m/12586>.
- ⁵¹D. D. BROWN and A. FREISE, *Finesse*, (May 2014) <http://www.gwoptics.org/finesse>.
- ⁵²V. BRAGINSKY and F. KHALILI, *Quantum measurement* (Cambridge University Press, 1995).
- ⁵³M. W. REGEHR, F. J. RAAB, and S. E. WHITCOMB, “Demonstration of a power-recycled michelson interferometer with fabry–perot arms by frontal modulation”, *Opt. Lett.* **20**, 1507–1509 (1995) [10.1364/OL.20.001507](https://doi.org/10.1364/OL.20.001507).
- ⁵⁴M. EVANS et al., “Lock acquisition of a gravitational-wave interferometer”, *Opt. Lett.* **27**, 598–600 (2002) [10.1364/OL.27.000598](https://doi.org/10.1364/OL.27.000598).
- ⁵⁵K. A. STRAIN et al., “Sensing and control in dual-recycling laser interferometer gravitational-wave detectors”, *Appl. Opt.* **42**, 1244–1256 (2003) [10.1364/AO.42.001244](https://doi.org/10.1364/AO.42.001244).
- ⁵⁶J. E. MASON and P. A. WILLEMS, “Signal extraction and optical design for an advanced gravitational-wave interferometer”, *Appl. Opt.* **42**, 1269–1282 (2003) [10.1364/AO.42.001269](https://doi.org/10.1364/AO.42.001269).
- ⁵⁷D. A. SHADDOCK et al., “Power-recycled michelson interferometer with resonant sideband extraction”, *Appl. Opt.* **42**, 1283–1295 (2003) [10.1364/AO.42.001283](https://doi.org/10.1364/AO.42.001283).
- ⁵⁸G. MÜLLER et al., “Dual-recycled cavity-enhanced michelson interferometer for gravitational-wave detection”, *Appl. Opt.* **42**, 1257–1268 (2003) [10.1364/AO.42.001257](https://doi.org/10.1364/AO.42.001257).
- ⁵⁹H. YAMAMOTO, M. EVANS, and E. MAROS, *End-to-end model*, (2015) <https://labcit.ligo.caltech.edu/~e2e>.

- ⁶⁰B. P. ABBOTT et al., “Gw151226: observation of gravitational waves from a 22-solar-mass binary black hole coalescence”, *Phys. Rev. Lett.* **116**, 241103 (2016) [10.1103/PhysRevLett.116.241103](#).
- ⁶¹B. P. ABBOTT et al., “Binary black hole mergers in the first advanced ligo observing run”, *Phys. Rev. X* **6**, 041015 (2016) [10.1103/PhysRevX.6.041015](#).
- ⁶²R. X. ADHIKARI, “Sensitivity and noise analysis of 4 km laser interferometric gravitational wave antennae”, PhD thesis (MIT, 2004), [1721.1/28646](#).
- ⁶³N. WIENER, *Extrapolation, interpolation, and smoothing of stationary time series* (The MIT Press, 1964).
- ⁶⁴J. C. DRIGGERS et al., “Active noise cancellation in a suspended interferometer”, *Review of Scientific Instruments* **83**, 024501 (2012) [10.1063/1.3675891](#).
- ⁶⁵J. C. DRIGGERS, J. HARMS, and R. X. ADHIKARI, “Subtraction of newtonian noise using optimized sensor arrays”, *Phys. Rev. D* **86**, 102001 (2012) [10.1103/PhysRevD.86.102001](#).
- ⁶⁶B. ALLEN, W. HUA, and A. OTTEWILL, “Automatic cross-talk removal from multi-channel data”, (1999), [arXiv:gr-qc/9909083](#).
- ⁶⁷B. SCHÖLKOPF et al., “Removing systematic errors for exoplanet search via latent causes”, in *Proceedings of the 32nd international conference on machine learning*, Vol. 37 (2015), pp. 2218–2226.
- ⁶⁸J. DURBIN, “The fitting of time-series models”, *Review of the International Statistical Institute* **28**, 233–244 (1960) [10.2307/1401322](#).
- ⁶⁹M. BAYES and M. PRICE, “An essay towards solving a problem in the doctrine of chances”, *Philosophical Transactions* **53**, 370–418 (1763) [10.1098/rstl.1763.0053](#).
- ⁷⁰J. S. READ et al., “Measuring the neutron star equation of state with gravitational wave observations”, *Phys. Rev. D* **79**, 124033 (2009) [10.1103/PhysRevD.79.124033](#).
- ⁷¹B. P. ABBOTT et al., “Properties of the binary black hole merger gw150914”, *Phys. Rev. Lett.* **116**, 241102 (2016) [10.1103/PhysRevLett.116.241102](#).
- ⁷²J. AASI et al., “Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network”, *Phys. Rev. D* **88**, 062001 (2013) [10.1103/PhysRevD.88.062001](#).
- ⁷³J. G. BAKER et al., “Gravitational-wave extraction from an inspiraling configuration of merging black holes”, *Phys. Rev. Lett.* **96**, 111102 (2006) [10.1103/PhysRevLett.96.111102](#).
- ⁷⁴M. HANNAM et al., “Simple model of complete precessing black-hole-binary gravitational waveforms”, *Phys. Rev. Lett.* **113**, 151101 (2014) [10.1103/PhysRevLett.113.151101](#).

- ⁷⁵M. van der SLUYS et al., “Parameter estimation of spinning binary inspirals using markov chain monte carlo”, *Classical and Quantum Gravity* 25, 184011 (2008) [10.1088/0264-9381/25/18/184011](https://doi.org/10.1088/0264-9381/25/18/184011).
- ⁷⁶J. VEITCH and A. VECCHIO, “Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network”, *Phys. Rev. D* 81, 062003 (2010) [10.1103/PhysRevD.81.062003](https://doi.org/10.1103/PhysRevD.81.062003).
- ⁷⁷C. BRWER et al., “Validating gravitational-wave detections: the advanced LIGO hardware injection system”, *Phys. Rev. D* 95, 062002 (2017) [10.1103/PhysRevD.95.062002](https://doi.org/10.1103/PhysRevD.95.062002).
- ⁷⁸R. SCHOFIELD, *Lho logbook entry 22497*, (2015) <https://alog.ligo-wa.caltech.edu/aLOG/iframeSrc.php?callRep=22497>.
- ⁷⁹J.-Y. VINET, V. BRISSON, and S. BRACCINI, “Scattered light noise in gravitational wave interferometric detectors: coherent effects”, *Phys. Rev. D* 54, 1276–1286 (1996) [10.1103/PhysRevD.54.1276](https://doi.org/10.1103/PhysRevD.54.1276).
- ⁸⁰D. J. OTTAWAY, P. FRITSCHER, and S. J. WALDMAN, “Impact of upconverted scattered light on advanced interferometric gravitational wave detectors”, *Opt. Express* 20, 8329–8336 (2012) [10.1364/OE.20.008329](https://doi.org/10.1364/OE.20.008329).
- ⁸¹B. CANUEL et al., “Displacement noise from back scattering and specular reflection of input optics in advanced gravitational wave detectors”, *Opt. Express* 21, 10546–10562 (2013) [10.1364/OE.21.010546](https://doi.org/10.1364/OE.21.010546).
- ⁸²N. MAVALVALA, “Alignment issues in laser interferometric gravitational-wave detectors”, PhD thesis (MIT, 1997), [1721.1/10769](https://doi.org/10.1721.1/10769).
- ⁸³P. AJITH et al., “Instrumental vetoes for transient gravitational-wave triggers using noise-coupling models: the bilinear-coupling veto”, *Phys. Rev. D* 89, 122001 (2014) [10.1103/PhysRevD.89.122001](https://doi.org/10.1103/PhysRevD.89.122001).
- ⁸⁴J. A. SIDLES and D. SIGG, “Optical torques in suspended fabry–perot interferometers”, *Physics Letters A* 354, 167–172 (2006) [10.1016/j.physleta.2006.01.051](https://doi.org/10.1016/j.physleta.2006.01.051).
- ⁸⁵K. HORNIK, “Approximation capabilities of multilayer feedforward networks”, *Neural Netw.* 4, 251–257 (1991) [10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T).
- ⁸⁶D. E. RUMELHART, G. E. HINTON, and R. J. WILLIAMS, “Learning representations by back-propagating errors”, *Nature* 323, 533–536 (1986) [10.1038/323533a0](https://doi.org/10.1038/323533a0).
- ⁸⁷S. HOCHREITER and J. SCHMIDHUBER, “Long short-term memory”, *Neural Computation* 9, 1735–1780 (1997) [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- ⁸⁸*Phase noise characterization of microwave oscillators*, Hewlett-Packard (1985).