

Syracuse University

SURFACE at Syracuse University

Dissertations - ALL

SURFACE at Syracuse University

5-12-2024

Overcoming Controls and Noise Challenges for High Power Interferometric Gravitational-Wave Detectors

Elenna Capote
Syracuse University

Follow this and additional works at: <https://surface.syr.edu/etd>



Part of the [Physics Commons](#)

Recommended Citation

Capote, Elenna, "Overcoming Controls and Noise Challenges for High Power Interferometric Gravitational-Wave Detectors" (2024). *Dissertations - ALL*. 1973.
<https://surface.syr.edu/etd/1973>

This Dissertation is brought to you for free and open access by the SURFACE at Syracuse University at SURFACE at Syracuse University. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE at Syracuse University. For more information, please contact surface@syr.edu.

ABSTRACT

In August 2017, the Advanced LIGO and Advanced Virgo detectors made the first coincident detection of gravitational waves from a binary neutron star merger, GW170817. This multi-messenger detection, measured with both gravitational wave detectors and electromagnetic telescopes, emphasized the ability of ground-based gravitational-wave detectors to make exciting astrophysical discoveries that revolutionize understanding of the universe. In this new era of gravitational-wave astronomy, high detector sensitivity and reliable operation are required to enable more of these exciting detections.

Ground-based gravitational-wave observatories are generally limited in sensitivity at low frequency by technical noise, at mid frequency by coating thermal noise, and at high frequency by quantum shot noise. This dissertation covers efforts to study and improve the controls scheme of Advanced LIGO to reduce low frequency technical noise and increase the operating power of the detector to improve the shot-noise-limited sensitivity. Reliable modeling of the alignment controls is presented, which enabled an increase to the highest detector operating power ever achieved with Advanced LIGO. Improvements in detector noise and operational stability summarized in this work also enabled the highest detector sensitivity ever achieved at the LIGO Hanford Observatory, more than doubling both the detector's sensitive volume and gravitational-wave detection rate.

This dissertation also summarizes work to study lower thermal noise optical coatings, making use of an improved analysis method for multimodal measurements. Preliminary results of the coating thermal noise of a novel amorphous coating mixture are presented, with a proposal for next steps. Finally, this work analyzes current detector limitations, and considers design requirements for future proposed detectors that will enable exciting astrophysics in the new age of gravitational-wave astronomy.

**OVERCOMING CONTROLS AND NOISE
CHALLENGES FOR HIGH POWER
INTERFEROMETRIC GRAVITATIONAL-WAVE
DETECTORS**

By

Elenna Capote

M.S., California State University Fullerton, 2020

S.B., Massachusetts Institute of Technology, 2019

DISSERTATION

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHYSICS

Syracuse University

June 2024

Copyright © Elenna Capote 2024
All rights reserved.

ACKNOWLEDGEMENTS

I begin my acknowledgements with the recognition of all my advisors who have guided and supported me through my academic career. First, I thank Stefan Ballmer, my PhD advisor, who saw my potential and has fostered it ever since. Thank you Stefan for having my back when I faced challenges, and thank you for making sure I was in the right place at the right time so I could achieve so much during these four short years. Next, I thank Josh Smith, my master's advisor, who is the one who introduced me to LIGO, taught me how to work in a lab and on a team, and helped me write my first paper. To this day, I know I can seek you out for advice whenever I need it. And third, I thank Deepto Chakrabarty, my undergraduate advisor. Without Deepto, I can easily say I would not have a PhD, let alone an undergraduate degree. I am so grateful that in my most difficult moments I had an advisor who still believed in me and championed me to the end.

To Ionel Tifrea, you may not have been my advisor, but you supported me all the same. Thank you for letting me walk into your statistical mechanics class, pushing me to finish my degree, and encouraging me into a PhD.

I cannot possibly list every wonderful person I worked with at LIGO Hanford, but I will try. Thank you to Jeff for teaching me about suspensions, Jim for answering all of my seismic questions, Erik and Jonathan for always lending a hand with the data access, Rick S. for making it possible for me to be a fellow, and Jason O. for letting me tag along in the PSL. Thank you to all the great operators: TJ, Tony, Corey, Ryan S, Ryan C, Ibrahim, Austin, and Oli. Working with all of you makes every day better. Thank you Camilla for being a great friend and collaborator.

Thank you Daniel for the wine and whiskey, and stopping by in the control room to drop off fantastic advice during commissioning struggles. Thank you Keita for reading my alogs, giving me hallway advice, and correcting my mistakes when I make them. Thank you Sheila for all your patience and guidance as I slowly learned how to be a commissioner. Also, thank you for asking me to look into the wavefront sensors. That moment is where it all began! Above all, thank you to Jenne for your guidance, mentorship, and wonderful faith in me. You let me make mistakes and learn from them, and always challenged me further.

To Dan B, Evan H, Kevin and Cao, thank you for working late into the night with me. I can't think of a better group of people to work with. Thank you Dan for teaching me how to use finesse and introducing me to the superior snack of mandarin oranges. Thank you to Kevin for always helping me with code and showing me geomspace. Thank you Evan for teaching me so much about controls. I have learned so much from shadowing you. Thank you Cao for always answering my questions and taking the time to laugh with me at silly things.

Gabriele, you stand out as a mentor to me in my graduate career. I am incredibly grateful for every chance I've had to work with you.

Thank you to Varun and Danny V. for teaching me how to use an SR785, take a transfer function, build electronics, and measure noise. Thank you Steve for opportunity to work on the GeNS and teaching me everything I know about coatings. Thank you Nick for being a great lab buddy and friend.

Craig, I will always be grateful for your git repositories and the time you took to teach me how to make good measurements. Georgia, thank you for your advice, guidance, and patience first as a commissioner and then as a professor.

Thank you to my committee members: Stefan, Georgia, Josh Russell and Mitch Soderberg. Thank you for reading this whole thing (really!), giving me great feedback, and supporting me through my defense.

To my San Diego yoga family— you lifted me up when I was in a dark place and helped me realize the person I could be. Thank you, all.

Ellie, every day of my life with you in it has been better than anything I could ever imagine. I love you.

And finally, to my mother, Janet Fox. You have known me and loved me from the very first moment. You are my number one cheerleader. Nothing in my life is possible without you.

Contents

List of Tables	xii
List of Figures	xv
Preface	xvi
1 Introduction to Gravitational Wave Detectors	1
1.1 Basics of a Gravitational Wave	1
1.2 Basics of a Gravitational Wave Detector	3
1.2.1 Fabry-Pérot Cavity	7
1.2.2 Power Recycling	9
1.2.3 Signal Recycling	9
1.2.4 Mode Cleaners	11
1.3 How to Lock an Advanced Gravitational Wave Detector	11
1.3.1 Lengths and Sidebands	12
1.3.2 Advanced LIGO Locking Sequence	14
2 Fundamental and Technical Noises of Gravitational-Wave Detectors	18
2.1 Fundamental Noise Sources	18
2.1.1 Quantum Noise	19
2.1.2 Thermal Noise	24
2.1.3 Seismic and Newtonian Noise	27
2.1.4 Other Fundamental Noises	28
2.2 Technical Noise Sources	28
2.2.1 Auxiliary Length Control Noise	31
2.2.2 Alignment Control Noise	34

2.2.3	Laser Noises	35
2.2.4	Other Technical Noises	36
2.3	Sensitivity Figures of Merit	37
3	Coating Characterization	43
3.1	Coating Thermal Noise in Gravitational Wave Detectors	44
3.2	Quantifying the Mechanical Loss of Coating Materials	46
3.2.1	Vibrational Mode Ringdowns	46
3.2.2	Gentle Nodal Suspension	47
3.2.3	Multimodal Measurements	55
3.3	Extracting Coating Loss From Coated Sample Loss	55
3.4	A Digital Lock-In Amplifier for Multimodal Analysis	59
3.5	Zirconia Titania Measurements	61
3.5.1	Measurement Procedure	62
3.5.2	Coating Loss of Zirconia-Titania Mixtures	62
3.5.3	Constraining Cation Concentration	66
3.5.4	Determining the Young's Modulus	66
3.5.5	Bulk and Shear Loss: fitting frequency dependence	68
3.5.6	Building a High Reflective Stack	70
3.6	Conclusions	71
4	Overview of Alignment Sensing and Control	72
4.1	Basics of Optical Cavities	72
4.1.1	Cavity Eigenmodes	74
4.1.2	Higher Order Mode Spacing	75
4.2	Alignment Sensing Concepts	78
4.2.1	Alignment Sensors	78
4.2.2	Alignment Degrees of Freedom	80
4.3	Alignment Sensing in Advanced LIGO	83
4.3.1	Interferometric DOFs	86
4.3.2	Pointing DOFs	90
4.3.3	Arm Transmission Sensing	91
4.4	Open Loop Transfer Functions	92
4.4.1	Measurement Principles	94

4.4.2	Alignment Control Actuation	97
4.4.3	Open Loop Gain Results	99
4.5	The Benefit of POP Wavefront Sensing	104
4.5.1	Current Alignment Control Configuration	105
4.5.2	Finesse Model	106
4.5.3	Benefits of POP Alignment Control	109
4.6	Conclusions	111
5	Noise Characterization of Alignment Sensing and Control	112
5.1	Noise Characterization of Alignment Sensors	114
5.1.1	Shot Noise	114
5.1.2	Dark Noise	115
5.1.3	ADC Noise	115
5.2	Contribution from Optic Motion	117
5.2.1	Improving Noise from Local Controls	119
5.2.2	Noise Improvements	122
5.3	Contribution from REFL Port Table Motion	127
5.3.1	Applying HAM1 Feedforward	129
5.3.2	Noise Improvements	131
5.4	Contribution from Centering Loops	134
5.5	Alignment Coupling from Spot Positions	135
5.6	Residual Noise in REFL Control Loops	140
5.7	Conclusions	142
6	Alignment Controls at High Power	146
6.1	Radiation Pressure in Fabry-Pérot Cavities	147
6.1.1	Beam Size and Radiation Pressure	152
6.2	High Power in Advanced LIGO	153
6.3	Controls Perspective of Radiation Pressure Torque	154
6.4	Compensating for Radiation Pressure	156
6.5	Thermal Effects in Advanced LIGO	163
6.6	Modeling Alignment Controls with High Power	164
6.6.1	Modeling the HARD Loops	164
6.6.2	Modeling the SOFT Loops	169

6.7	Conclusions	170
7	The Performance of the Laser Interferometer Gravitational-Wave Observatory During the Fourth Observing Run	172
7.1	Introduction	172
7.2	O4 Overview	172
7.2.1	Astrophysical Range	173
7.2.2	Duty cycle	174
7.2.3	Lock Acquisition	174
7.3	Instrument Noise	177
7.3.1	Quantum noise	178
7.3.2	Thermal noise	180
7.3.3	Laser noise	183
7.3.4	Controls noise	187
7.3.5	Photodetector dark noise	194
7.3.6	Scattered light noise	194
7.4	Instrument improvements	195
7.4.1	Pre-Stabilized Laser upgrade	195
7.4.2	Test mass replacements	196
7.4.3	Squeezer upgrades	197
7.4.4	Auxiliary controls upgrades	199
7.4.5	Online noise subtraction	201
7.4.6	Scattered light mitigation	202
7.4.7	Parametric instability mitigation	203
7.4.8	Output photodetector electronics upgrades	204
7.4.9	Grounding improvements	204
7.5	Summary of High Power Operation	205
7.5.1	Arm power characterization	205
7.5.2	High power challenges	207
7.6	Future work	211
8	Technical Noise and Data Quality Considerations for Future Detectors	213
8.1	Introduction	213
8.2	Broadband Technical Noises	215

8.2.1	Auxiliary Length Control Noise	215
8.2.2	Alignment Sensing and Control Noise	219
8.2.3	Science impacts of low-frequency noise	223
8.3	Narrow Noise Features	227
8.4	Discussion	232
9	Future Work	235
9.1	Coating Thermal Noise	235
9.2	Alignment Sensing	236
9.3	Alignment Control	236
9.4	Future Detectors	237
	Bibliography	238

List of Tables

1	Advanced LIGO main cavity lengths and sensing design.	13
2	Deposited sample parameters measured via ellipsometry	64
3	Rutherford Backscattering Spectroscopy results: zirconia concentration, density, index, extinction coefficient	66
4	Rutherford Backscattering Spectroscopy results: elemental breakdown, atomic density	66
5	Coating material, HR layer thickness, total layers	71
6	Reflection port input matrix	87
7	Antisymmetric port input matrix	91
8	Arm transmission input matrix	92
9	Input power, power-recycling gain, and arm cavity power to LHO during O4 commissioning	154
10	Summary of the operating power parameters of the LIGO interferometers during O4	206
11	Detection parameters for different detector configurations	226
12	Summary of future XG science goals with technical noise and data quality considerations	233

List of Figures

1	Test Points Experiencing a Gravitational Wave	3
2	A basic Michelson interferometer	4
3	Schematic of a Fabry-Pérot cavity	8
4	A cartoon of the optical layout of Advanced LIGO, including squeezer. Diagram made by G. Mansell.	10
5	Advanced LIGO O4 Optical Layout	17
6	Advanced LIGO Design Sensitivity	20
7	Sensitivity Improvement from Frequency Dependent Squeezing	23
8	LIGO Hanford Detector Noise Budget, O4	29
9	MICH to DARM Transfer function with and without feedforward	33
10	Range integrands from O3 to O4	40
11	Comoving volume of the LIGO Hanford Detector	41
12	The Advanced LIGO+ Design Curve	44
13	Diagram of a Gentle Nodal Suspension	47
14	GeNS Optical Lever Schematic	48
15	Modal basis of a 76 mm diameter, 1 mm thick fused silica disk	49
16	Optical Lever Transimpedance Amplifier Schematic	54
17	Modeled Frequency Response QPD Transimpedance Amplifier	54
18	A spectrum of excited modes on an uncoated 1 mm thick optical sample	56
19	A comparison of dilution factors determined analytically and with an FEA	58
20	Semi-degenerate modes measured in the GeNS	60
21	Fitting a damped sinusoid ringdown	61
22	ZrTi Loss versus Annealing Temperature	63
23	Measured Frequency Dependence of zirconia titania loss	65
24	Frequency Shifts of Coated ZrTi samples	67

25	Sample Loss Frequency Dependence	69
26	A single-bounce mode scan of the aLIGO input beam	77
27	Basic diagram of a wavefront sensor	78
28	A diagram of the HARD and SOFT mode degrees of freedom in a Fabry-Pérot cavity	80
29	A diagram of the Michelson alignment degree of freedom	81
30	A diagram of the recycling cavity alignment degree of freedom	82
31	Arm cavity basis alignment in yaw	84
32	Arm cavity basis alignment in pitch	85
33	REFL sensing matrix in pitch	88
34	REFL sensing matrix in yaw	89
35	Advanced LIGO Alignment Sensor Layout, O4	93
36	A basic diagram of a control loop	95
37	An open loop gain measurement of CHARD pitch and yaw	100
38	An open loop gain measurement of DHARD pitch and yaw	101
39	An open loop gain measurement of MICH pitch and yaw	102
40	Open loop gain measurements of the input and recycling cavity alignment control	103
41	Sensing Matrix Measurements at POPX	106
42	Finesse3 simulation of a POP WFS in pitch	107
43	Finesse3 simulation of a POP WFS in yaw	108
44	Finesse3 simulations of a 9 MHz POP WFS	109
45	LHO Noise Budget from O3	113
46	A diagram of the RF signal chain	116
47	A REFL WFS noise budget	118
48	Optic motion in REFL WFS	120
49	Model of IM pitch-to-pitch transfer function	121
50	Measured PR3 pitch-to-pitch transfer function	123
51	REFL WFS noise contributions	124
52	Improvement to differential arm length, IM damping	125
53	Reduction of DARM Error signal RMS, PR3 damping	126
54	Pitch/Yaw Cross Coupling in the Transmission QPDs	128
55	NonSENS Offline Subtraction	130

56	A model of the CHARD pitch closed loop response	131
57	HAM1 Feedforward applied to ASC	132
58	Differential arm length improvement with HAM1 feedforward	133
59	DC centering loop coupling to CHARD	135
60	DARM Improvement from DC centering loop bandwidth reduction	136
61	A diagram of spot-miscentering leading to angle-to-length coupling.	137
62	Measuring angle-to-length coupling using DHARD Yaw injections	139
63	Calibrated CHARD in-loop error signal noise	141
64	Calibrated INP1 in-loop error signal noise	142
65	Calibrated PRC2 in-loop error signal noise	143
66	A comparison of O3 and O4 alignment noise	144
67	A diagram of the hard and soft mode	149
68	Sidles Sigg Modes in the Advanced LIGO QUAD	150
69	A diagram of radiation pressure torque as a feedback loop	155
70	Modeled pitch-to-pitch suspension transfer functions for the HARD and SOFT modes at different cavity powers	157
71	Control Diagram with Radiation Pressure	158
72	Control Diagram using Radiation Pressure Compensation	159
73	HARD loop controllers for Radiation Pressure Compensation	161
74	HARD Open Loop Gain with Radiation Pressure Compensation	162
75	A DHARD P modeled with RoC changes	165
76	Model and Measurement of CHARD P	166
77	Model and Measurement of DHARD P	167
78	The effect of drive mismatch on HARD yaw suspension plant	168
79	Model and Measurement of CHARD Y	168
80	Model and Measurement of DHARD Y	169
81	Model and Measurement of CSOFT P	170
82	Modeled drive mismatch on CSOFT P	171
83	Cumulative events up to January 16, 2024, including the entirety of O4a . . .	173
84	Range trends for the LIGO Hanford and LIGO Livingston observatories during the first half of the fourth observing run	175
85	Noise budgets for the LIGO Hanford detector and LIGO Livingston detector during observing run four	179

86	Thermal noise sub-budget for the LIGO Livingston detector	182
87	Laser noise subbudget for LIGO Hanford and Livingston	184
88	Laser frequency noise for LIGO Hanford and Livingston	185
89	Beam jitter noise compared at both detectors	187
90	A comparison of the differential arm length (DARM) and Michelson length (MICH) coupling at LHO	189
91	Limiting alignment control loops measured at the LIGO Hanford Observatory	192
92	Local damping noise contributions to gravitational wave strain noise measured at the LIGO Livingston Observatory	193
93	Comparison of detector noise with and without squeezing	205
94	Low frequency noise trends during LHO power up	212
95	Proposed Cosmic Explorer Design Sensitivity with low frequency cutoffs . . .	225
96	A comparison of gravitational-wave signals from pulsars and known instru- mental lines	228

Preface

The work presented in this thesis stems from my participation in the LIGO Scientific Collaboration (LSC). This work does not reflect the scientific opinion of the LSC and it was not reviewed by the collaboration.

For my father, Miguel Alberto Capote
November 7, 1954 - April 22, 2007
He was a scientist, an engineer, an artist, and a jokester.
He taught me how to ask questions and love learning.
May he rest in peace.

Chapter 1

Introduction to Gravitational Wave Detectors

1.1 Basics of a Gravitational Wave

Before discussing gravitational-wave detectors, we must have some baseline understanding of what a gravitational wave is—how it propagates, how it is created. An understanding of gravitational waves begins with some principles of General Relativity.

To start, some spacetime metric $g_{\mu\nu}$ describes spacetime with gravitational waves propagating in it. Using a standard Minkowski metric $\eta_{\mu\nu}$ for flat spacetime and add some small perturbation $h_{\mu\nu}$ where $|h_{\mu\nu}| \ll 1$,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (1.1)$$

The vacuum solution yields a wave equation

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

where $\bar{h}_{\mu\nu} = h_{\mu\nu} - \eta_{\mu\nu} h^\alpha_\alpha$. To achieve $\bar{h}_{\mu\nu} = h_{\mu\nu}$, $h_{\mu\nu}$ must follow the conditions of the “transverse-traceless gauge”, which are

$$A^{\mu\nu} k_\nu = 0 \quad (1.3a)$$

$$A^\mu_\mu = 0 \quad (1.3b)$$

$$A_{\mu\nu} U^\nu = 0 \quad (1.3c)$$

where $A^{\mu\nu}$ and k_ν are the amplitude and wave number of the solution to Equation 1.2, and U^ν is any velocity four-vector.

The wave equation in Equation 1.2 yields a *gravitational wave*, that is, a wave that propagates through spacetime itself and transfers gravitational-wave energy. This result is one of many revolutionary results that Einstein predicted in his theory of General Relativity, and one that has proven to be the most difficult to prove accurate. Nevertheless, with both indirect proof from observations such as the Hulse-Taylor pulsar [1], and the first direct measurement GW150914 made by Advanced LIGO [2], we now know that gravitational waves exist and are measurable.

Following the usual convention, the gravitational wave propagates in the \hat{z} direction, meaning that two elements will make up the amplitude of the wave, h_+ and h_\times in the x- and y-direction,

$$A^{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.4)$$

with these elements representing two different polarizations of the gravitational wave.

Consider a propagating gravitational wave with plus polarization only. In this case, the x-direction will measure a positive amplitude, or an increase in the spacetime distance between two points, while the y-direction will experience a negative amplitude, or the decrease in the spacetime distance between two points. In this way, gravitational waves are often described as “stretching” and “squeezing” spacetime. This behavior is demonstrated in Figure 1, where a ring of test points experiences a gravitational wave passing through the plane of this dissertation.

This propagating behavior of a gravitational wave leads to the concept of gravitational wave *strain*, which comes from the amplitude of each term, $|h_{\mu\nu}|$. To measure strain, we have to understand how much the length between two points will change as a result of a gravitational wave. If we consider two points separated by a distance L , the relative distance change between these two points as a result of a passing gravitational wave is

$$\delta L = \int_0^L |g_{\mu\nu} dx^\mu dx^\nu|^{1/2} \approx \left(1 + \frac{h_{xx}}{2}\right) L. \quad (1.5)$$

Therefore, we can measure gravitational wave strain as some change in length over a given distance $h \approx \delta L/L$.




Figure 1: A ring of test points experiencing a gravitational wave passing through the plane of this dissertation. One direction is compressed while the perpendicular direction is stretched. Both cross (h_{\times}) and plus (h_{+}) polarizations are demonstrated.

Gravitational waves result from sources with a time-varying mass quadrupole moment, in other words, a system with mass moving in an asymmetric manner. In principle, this could be any asymmetric mass, but the biggest challenge in direct measurement of gravitational waves lies in the weakness of the field itself. Therefore, gravitational waves are most likely to be detected from high-mass, high-acceleration systems that produce the most significant gravitational wave amplitude. A binary system of orbiting black holes, for example, has significant enough mass and a varying quadrupole moment to make measurable gravitational waves. However, this still requires a strain measurement on the order of 10^{-21} . Therefore, an instrument that measures gravitational waves will need to make some of the most precise length-change measurements in the world.

1.2 Basics of a Gravitational Wave Detector

This basic background of gravitational waves lends some understanding of how one might detect them. First, the measurement quantity is strain, some length change over a baseline.

Given the very small value of expected strain, a longer baseline improves the sensitivity to gravitational waves. These waves also propagate length changes in a differential manner, extending length in one direction while shortening it in the other. An experiment that exploits this behavior would be useful.

There is an interesting history of gravitational wave detectors and design, but we will jump in to the considerations that led to the creation of Advanced LIGO. Rainer Weiss first proposed in 1972 that a gravitational wave detector make use of Michelson interferometer topology [3]. A basic Michelson inteferometer, shown in Figure 2, takes some input laser beam, splits it down two perpendicular arms where the beams are reflected by end mirrors and recombine again at the beamsplitter. The resulting light signal is read out just beyond the beamsplitter, at a location called the “anti-symmetric port”. A photodetector placed at the anti-symmetric port measures the brightness of the transmitted beam. A gravitational wave passing around this experiment would shorten one arm and lengthen the other, causing destructive interference in the recombined beam, resulting in a power change on the detector.




Figure 2: A basic Michelson interferometer. Detectors placed at the symmetric or antri-symmetric port can measure the reflected and transmitted electric fields respectively. From [4].

Following the derivation in Saulson [5], the electric field of the laser input to the Michelson interferometer is

$$E_{in} = E_0 e^{i(2\pi ft - kx)} \quad (1.6)$$

where E_0 is the amplitude of the electric field, f is the frequency of the wave, and k is the wavenumber, $k = \frac{2\pi}{\lambda}$, where λ represents the laser wavelength. The field encounters a beamsplitter with 50/50 reflection and transmission

$$r = 1/\sqrt{2} \quad (1.7a)$$

$$t = i/\sqrt{2}. \quad (1.7b)$$

The fields propagating in the x and y arms are then

$$E_x = \frac{iE_0}{\sqrt{2}} e^{i(2\pi ft - kx)} \quad (1.8a)$$

$$E_y = \frac{E_0}{\sqrt{2}} e^{i(2\pi ft - ky)} \quad (1.8b)$$

The fields traverse the length of their respective arms (labeled L_x and L_y in this case), reflect off the end mirrors acquiring a -1 sign, and return to the beamsplitter, once again encountering 50/50 reflection/transmission. The output field at the anti-symmetric port is

$$\begin{aligned} E_{out} &= \frac{iE_0}{2} e^{i(2\pi ft - 2kL_x)} + \frac{iE_0}{2} e^{i(2\pi ft - 2kL_y)} \\ &= \frac{iE_0 e^{i(2\pi ft)}}{2} (e^{-i2kL_x} + e^{-i2kL_y}) \\ &= iE_0 e^{i(2\pi ft - k[L_x + L_y])} \cos k(L_x - L_y). \end{aligned} \quad (1.9)$$

However, the photodetector at the output port measures power, not electric field,

$$\begin{aligned} P_{out} &= |E_{out}|^2 \\ &= E_0^2 \cos^2 k(L_x - L_y) \\ &= \frac{P_0}{2} (1 + \cos(2k\Delta L)) \end{aligned} \quad (1.10)$$

where P_0 is the input power and $\Delta L = L_x - L_y$. If the arms are set to be exactly equal, the light measured at the output port experiences constructive interference, as the two fields travel the same distance. If the length of the arms are such that $2k\Delta L = n\pi/2$ where n is an integer, then the fields will destructively interfere and the power measured at the output is zero.

Light acquires phase ϕ as it travels depending on the distance d it travels and its wavelength λ . The phase change can be written in terms of distance or time of travel t ,

$$\Delta\phi = 2\pi \frac{d}{\lambda} = 2\pi \frac{ct}{\lambda}. \quad (1.11)$$

Considering the null geodesic of a photon traveling the length of the arms in the presence of a gravitational wave,

$$\begin{aligned} ds^2 &= (\eta_{\mu\nu} + h_{\mu\nu})dx^\mu dx^\nu \\ 0 &= -c^2 dt^2 + (1 + h(t))dx^2 + (1 - h(t))dy^2 + dz^2. \end{aligned} \quad (1.12)$$

The distance the light travels in the x-arm increases due to gravitational wave strain while the distance decreases in the y-arm.

The light travel time can be found by integrating

$$\tau_x = \int_0^{L_x/c} \sqrt{1 + h(t)} dt \quad (1.13a)$$

$$\tau_y = \int_0^{L_y/c} \sqrt{1 - h(t)} dt. \quad (1.13b)$$

While the gravitational wave amplitude is changing in time, it is reasonable to assume the period is much longer than the round-trip time of the light in the interferometer. Therefore, h is constant in this integration, and also much smaller than the total length of the detector arm length.

$$\tau_x = \frac{L}{c} + \frac{L_x}{2c} h \quad (1.14a)$$

$$\tau_y = \frac{L}{c} - \frac{L_y}{2c} h. \quad (1.14b)$$

These values are then doubled to account for the fact that a photon travels twice the length of each arm when recombining at the beamsplitter. With the assumption that h is very small, $L_x = L_y$ to make this calculation simpler. Therefore, the difference in the light travel time between the two arms in the presence of a gravitational wave is

$$\Delta\tau = \tau_x - \tau_y = 2\frac{L}{c}h. \quad (1.15)$$

This travel time difference results in a phase difference between the two light fields

$$\Delta\phi = \frac{4\pi hL}{\lambda}. \quad (1.16)$$

A more rigorous treatment of the phase would consider the actual form of $h(t)$ instead of estimating it to be constant. This is important to consider, because indefinitely increasing the length of a gravitational wave detector's arms would not necessarily result in an increasingly

better detection of a gravitational wave. The full form of the phase difference, skipping over the rigorous details, applies $h(t) = h_0 e^{i2\pi f_{GW} T}$,

$$\Delta\phi(t) = h(t) \frac{4\pi L}{\lambda} \operatorname{sinc} \left(\frac{2\pi f_{GW} L}{c} \right) e^{2i\pi f_{GW} L/c}. \quad (1.17)$$

This result demonstrates that there is a zero point in the phase response of a gravitational wave detector every time $\frac{2\pi f_{GW} L}{c} = n$ where n is a positive integer. This describes the situation in which the wavelength of the gravitational wave is equal to $2L$. The gravitational wave passes through one wavelength in the time it takes the photon to travel the full length of the arm, nullifying the effect of the contracted and extended length of the arm during the photon's travel.

For a gravitational-wave detector with arm length 4 km, this null point first occurs at $f_{GW} \approx 12$ kHz, much higher than the expected gravitational wave frequencies for known systems. A proposed future ground-based gravitational-wave detector is called Cosmic Explorer [6], and is designed to have 40 km arms. This first null point then occurs near 1.2 kHz, which is reasonably within the detection band. However, this detector proposal is near the full limit of possible arm length for gravitational-wave detectors. For example, if the detector arm length were set to 100 km, this null point would occur near GW frequencies of 400 Hz, compromising the sensitivity in a frequency band where ground-based detectors are most sensitive.

For a ground-based detector this is not much of a problem, as other limitations prevent the indefinite extension of detector arm length. However, space-based detectors like the Laser Interferometer Space Antenna [7] will have to account for this effect within the detection band due to an arm length on the order of a million kilometers.

Despite the ingenuity of the Michelson topology, it alone is insufficient to make a detection of a gravitational wave. Also, because the arm length cannot be extended significantly to increase the sensitivity to gravitational waves, other methods are used to improve detector sensitivity. As will often be described in this dissertation, the best way forward usually involves adding more optical cavities to the design.

1.2.1 Fabry-Pérot Cavity

By adding an additional mirror in the Michelson interferometer arms, light can reflect back and forth between the two mirrors before exiting to the anti-symmetric port, effectively




Figure 3: A basic Fabry-Pérot cavity

increasing the light travel distance and improving the sensitivity. One such way to do this is by adding a Fabry-Pérot cavity to each arm. Each arm of the Michelson then consists of an input mirror and end mirror, and the multiple reflections between the two mirrors increases the distance each photon travels. A Fabry-Pérot cavity also uses interference properties; by precisely controlling the length of the cavity, the light resonates in the cavity and amplifies through constructive interference.

Figure 3 shows a basic schematic of such a cavity. Assuming the two mirrors are lossless, the cavity dynamics can be understood completely from the cavity length L and the reflection or transmission of each mirror, represented by r_i and t_i . Some monochromatic light with wavelength λ is injected in the Fabry-Pérot cavity. The three resulting fields, reflected, transmitted, and circulating are defined as [5]

$$E_{refl} = E_0 \left(-r_1 + \frac{t_1^2 r_2 e^{-i2kl}}{1 - r_1 r_2 e^{-i2kl}} \right) \quad (1.18a)$$

$$E_{trans} = E_0 \left(\frac{t_1 t_2 e^{-ikl}}{1 - r_1 r_2 e^{-i2kl}} \right) \quad (1.18b)$$

$$E_{circ} = E_0 \left(\frac{t_1}{1 - r_1 r_2 e^{-i2kl}} \right) \quad (1.18c)$$

The cavity achieves resonance when the circulating field within the cavity is maximized,

$$\phi_{rt} = 2kL = \frac{4\pi L}{\lambda} = 2\pi N \quad (1.19a)$$

$$L = \frac{N\lambda}{2} \quad (1.19b)$$

where λ is the wavelength of the laser incident on the cavity, and N is some integer. In other words, when the cavity length is well matched to the laser wavelength, resonance occurs. This relationship assumes a lossless cavity. To account for the round-trip loss a parameter

called the *finesse* is defined [8],

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2}. \quad (1.20)$$

that depends on the reflectivities of the two cavity mirrors.

The intensity of the field within the cavity therefore has a periodic behavior. As the length of the cavity is changed (or the frequency of the laser is scanned), sharp maxima in the intensity occur when $\phi_{rt} = 2\pi N$, with the width of the maxima defined by $\Delta\phi = 2\pi/\mathcal{F}$.

The Advanced LIGO Fabry-Pérot cavities have a finesse of 450 [9] and length 4 km, which corresponds to a light storage time of $\tau_s \approx 2$ ms, or an effective arm length of approximately 2800 km.

1.2.2 Power Recycling

The Fabry-Pérot cavity arms amplify the input power according to the finesse, which in turn increases the sensitivity of the detector. However, even with the highly reflective mirrors designed for gravitational-wave interferometer Fabry-Pérot cavities, significant power will transmit the input mirror of each arm and return to the input port. By placing an additional partially transmissive mirror on the input of the interferometer, labeled “power-recycling mirror” in Figure 4, this effectively adds another input mirror, treating the Fabry-Pérot cavities in the arms as the far mirror. This coupled cavity amplifies the power build-up in the interferometer further via the *power-recycling gain*, which for Advanced LIGO is between 40-50 W/W.

1.2.3 Signal Recycling

Similar to the added mirror at the input of the interferometer, signal recycling adds a mirror at the output of the interferometer. A true signal-recycling scheme shifts the resonance of the interferometer away from f_l , defined by the length of the Fabry-Pérot cavities, to $f_l \pm f_{sig}$, where f_{sig} is the frequency of the gravitational-wave signal measured by the detector [5]. However, this method overall narrows the bandwidth of the interferometer, sacrificing broadband sensitivity for a more targeted measurement of certain signals.

Instead, many gravitational-wave detectors such as Advanced LIGO apply *resonant sideband extraction*, which appears similar to signal recycling, except that the mirror at the output is placed to be anti-resonant to the main cavity beam at f_l and resonant to the sidebands, allowing them to leak to the detection port [10]. This effectively broadens the




Figure 4: A cartoon of the optical layout of Advanced LIGO, including squeezer. Diagram made by G. Mansell.

bandwidth of the interferometer while maintaining the benefit of high resonating power in the Fabry-Pérot cavities. Of course, this is confusingly labeled as a “signal-recycling mirror” in Figure 4 for historical reasons. There is a nice symmetry in having both a “power-recycling cavity” and “signal-recycling cavity”, even if it is truly a *signal extraction cavity*. For the majority of this dissertation, this cavity will be referred to as the “signal-recycling cavity” due to Advanced LIGO convention, but the final chapter will correctly name it the “signal-extraction cavity”.

1.2.4 Mode Cleaners

Not satisfied with the amount of cavities required for a gravitational-wave detector? Fear not, there are more. Although these two cavities are both called “mode cleaners”, they have similar, but slightly different purposes and designs. The Input Mode Cleaner (the original MC) is designed to stabilize the input laser frequency and spatial mode as it is injected into the interferometer [11]. The IMC is a high-finesse triangular cavity at the input of the interferometer. Once the laser beam is created, stabilized, modulated, and injected into the interferometer, the IMC provides additional frequency stabilization, angular control, and higher order mode suppression. This cavity is also shown on the diagram in Figure 4.

Conversely, the Output Mode Cleaner (OMC) is a bowtie cavity at the output of the interferometer. The design suppresses the radio frequency sidebands that modulate the main beam to allow for PDH locking [12]. The cavity is designed to be resonant only to the zeroth order mode audio-band carrier beam, rejecting both the sidebands and any higher order mode that transmits the output of the interferometer. The DC photodiodes that readout the gravitational-wave signal are on transmission of this cavity. It is also shown in Figure 4.

1.3 How to Lock an Advanced Gravitational Wave Detector

So far, the optical layout of an advanced gravitational-wave detector has only gotten more complex. Now that there is a baseline understanding of the layout and requirements, this section will discuss how all of these degrees of freedom are controlled in order to bring a gravitational-wave detector to a state where it can detect gravitational waves. This section will mostly focus on Advanced LIGO, but describes other ground-based detectors where possible.

1.3.1 Lengths and Sidebands

Several different required cavities have been suggested for gravitational-wave detectors. This section will list the cavity lengths and discuss the reasoning behind those length choices. A full discussion of these design choices can be found in [13]. First, for a dual-recycled Michelson with Fabry-Pérot cavity arms, there are five main length degrees of freedom:

1. Differential Arm Length, **DARM**: the difference in arm length between the two 4 km Fabry-Pérot cavity arms. Defining the two arm lengths as L_x and L_y (X-arm and Y-arm), this is $L_x - L_y$.
2. Common Arm Length, **CARM**: the common arm length of the two 4 km Fabry-Pérot cavity arms, $L_x + L_y$.
3. Michelson Length, **MICH**: the length of the Michelson cavity, or the distance between the beamsplitter and input test mass of the X and Y arm. Defining the distances from beamsplitter to each input test mass as l_x and l_y respectively, this is $l_x - l_y$.
4. Power-recycling cavity length, **PRCL**: the length of the power-recycling cavity, $l_p + (l_x + l_y)/2$ where l_p is the distance from the power-recycling mirror to the beamsplitter.
5. Signal-recycling cavity length, **SRCL**: the length of the signal-recycling cavity, $l_s + (l_x + l_y)/2$ where l_s is the distance from the beamsplitter to the signal-recycling mirror.

Note: some definitions for DARM and CARM involve the *average* distance value, and therefore include an additional 1/2 factor. In addition to these main five lengths, the input mode cleaner length and the output mode cleaner length must be locked.

As stated previously, Advanced LIGO design includes two 4 km Fabry-Pérot cavity arms. Other ground-based detectors such as Virgo in Italy [14] and KAGRA in Japan [15] use 3 km arms. The goal of this design choice is to achieve the longest possible baseline L over which to measure strain. This length choice sets the macroscopic DARM and CARM lengths.

Next, the macroscopic lengths of the three “DRMI” (dual-recycled Michelson) lengths must be chosen. Advanced LIGO makes use of the Pound-Drever-Hall Locking technique (PDH), which requires the injection of radio frequency (RF) sidebands in addition to the main laser beam [16]. The sidebands chosen for Advanced LIGO are $f_1 = 9.1$ MHz and $f_2 = 5f_1 = 45.5$ MHz. With these choices, the cavity lengths must be set to follow the desired resonance conditions for these sidebands and the carrier beam.

For the power-recycling cavity length, the carrier and both sidebands must be resonant, which is achieved when,

$$l_{PRCL} = \left(n + \frac{1}{2}\right) \frac{c}{2f_1} \quad (1.21)$$

where n is an integer. The arms add a 180° phase flip to the sidebands, creating the additional $1/2$ factor.

Next, the signal-recycling cavity must be antiresonant to both the carrier and f_1 sideband in order to achieve resonant sideband extraction. Therefore,

$$l_{SRCL} = m \frac{c}{2f_2}, \quad l_{SRCL} \neq k \frac{c}{2f_1} \quad (1.22)$$

which allows the 45 MHz sideband to resonate but not the carrier or 9 MHz.

Finally, the Michelson cavity length must be set so that the 45 MHz is preferentially transmitted to the SRC. This length offset is the *Schnupp asymmetry*, $l_{Schnupp} = l_x - l_y$, which allows some light to leak out of the Michelson cavity so that it can be locked using PDH. Table 1 gives a summary of these key lengths for Advanced LIGO. As a note, the arm length is set so that the sidebands are anti-resonant in the arms.

Cavity	Length	Sensor
Arm	3994.5 m	CARM: REFL 9 I, DARM: DCPD
PRC	57.6557 m	POP 9 I
SRC	56.0084 m	POP 9 and 45 I
MICH	$l_{Schnupp} = 0.08$ m	POP 45 Q

Table 1: Advanced LIGO main cavity lengths and sensing design.

Due to the difference in length choices, Virgo and KAGRA make use of different sideband frequencies. KAGRA, at this time, operates only a power-recycled Michelson due to asymmetries in the Fabry-Pérot input mirrors that cause too many losses to operate with an SRC. Virgo followed this operation in observing run 3, but has transitioned to a dual-recycled topology for observing run 4.

Not yet discussed in this section is a third sideband, $f_3 = 13f_1 = 118.3\text{MHz}$. This sideband is injected to allow better alignment sensing of the signal-recycling cavity, and will be referenced later.

Each of these lengths must be sensed using an appropriate signal. The benefit of PDH locking is the beatnote between the carrier and RF sideband creates a linear control response

near the lock point. The resulting RF beat signal is demodulated, and the I-phase or Q-phase provides the error signal for the degree of freedom. Lengths sensed via cavity reflection, such as CARM or PRCL are best sensed in I-phase, while lengths sensed via transmission with an offset are best sensed in Q-phase. The three DRMI lengths are sensed at the pick-off port of the power-recycling cavity (POP), on transmission of power-recycling mirror 2 (see Figure 5). The PRCL signal appears most strongly in the 9 MHz I-phase beat signal, Michelson in 45 MHz Q-phase, and SRCL in 45 MHz I-phase. However, the SRCL and PRCL signals both appear in 45 MHz because the 45 MHz is resonant in both cavities, so the 9 MHz signal is used to subtract the PRCL signal from the SRCL signal, as 9 MHz is anti-resonant in the SRC.

CARM is sensed on reflection of the interferometer, at the “common length readout” labeled in Figure 5. The beat of the carrier with the 9 MHz sideband in I-phase provides the best signal. CARM is sensitive to both the common arm length and the power-recycling cavity length due to the coupled-cavity nature of Advanced LIGO. However, the CARM bandwidth is set much higher (16 kHz) than the PRCL bandwidth (30 Hz), so that the CARM signal is quickly driven to zero within the PRCL bandwidth.

DARM is sensed through the intentional DC offset of the arms, creating approximately 20 mA of photocurrent on the DC photodiodes (DCPDs) on transmission of the OMC. The length offset scales quadratically with the power on the DCPDs, so it is kept intentionally small so that the signal is within the linear regime.

Control gain hierarchy is important in this sensing scheme, ensuring that the lower bandwidth loops have a stable reference to the carrier. CARM is the highest bandwidth loop (16 kHz), and DARM is next, around 50 Hz. The corner hierarchy order is PRCL (30 Hz), SRCL (20 Hz) and then MICH (8 Hz). This avoids cross-coupling in the signals due to some degeneracy in the lengths. The higher bandwidth loop is driven to zero fast relative to the lower bandwidth loop.

1.3.2 Advanced LIGO Locking Sequence

Given the complexity of a detector like Advanced LIGO, moving into a state where gravitational waves can be observed (often referred to as “observing” or “science” mode) is not a trivial matter. All cavities must achieve resonance conditions at the same time. The benefit of Pound-Drever-Hall locking is the linear response of the control signal; however, all optics must be close to good placement and alignment to be in that linear regime. The Advanced

LIGO locking process is a procedure developed due to the effort and cleverness of many people who carefully considered a reliable method to achieve resonance. The process requires auxiliary lasers and sensors and takes time to carefully move through each stage.

Starting from a fully unlocked state, the locking sequence begins with the two Fabry-Pérot arm cavities. Using a green laser created through frequency doubling of the 1064 nm laser, each arm is individually locked using PDH. This green laser is part of the Arm Length Stabilization system (ALS) [17]. The light is sent from the end stations of each arm, and length and alignment sensors on the platform behind each end test mass provide the control signals. Usually the input test masses are held in one place and the end test masses and transmission monitor platform are adjusted in alignment until the resonance conditions are met and the length lock is triggered.

Next, the X arm is locked to the pre-stabilized laser, creating the common degree of freedom (ALS COMM) and the Y arm is locked to the X arm for the differential degree of freedom (ALS DIFF). While still locked on green, a small amount of 1064 nm light is allowed to resonate in each arm, finding the respective resonance condition for the main beam. A voltage-controlled oscillator (VCO) adjusts its offset to find the appropriate condition. Once this condition is found, a deliberate offset is added to ALS COMM to maintain both arms off-resonance to the carrier light. The green laser provides the signal to stabilize the arms during the next process.

Next, the 1064 nm laser is used to lock the dual-recycled Michelson cavities using the PDH technique. The alignments of the power-recycling mirror, beamsplitter, and signal-recycling mirror are all adjusted as necessary to achieve the resonance of these cavities with the carrier. The locking signals are provided via the 9 and 45 MHz RF sidebands. Once these cavities are locked, the “1f” signal sensors are handed off to “3f” signal sensors, 27 and 135 MHz. When the arms are brought to resonance, the “1f” signals will flip sign due to the additional 180° phase shift from the arm cavity, which is anti-resonant to the sidebands. This would destabilize the DRMI lock and cause a lockloss. The “3f” signals are not sensitive to this sign flip and provide better sensors during the locking process [18].

With the corner cavities locked, the arms can now be brought to resonance, making use of the known offset in the ALS VCO. The carrier beam is allowed to resonate a small amount in each arm—the common offset is still engaged. Now, this is referred to as the “CARM offset”, and this whole process as “CARM offset reduction”. As the CARM offset is carefully reduced, the cavity signals must be carefully measured to find out what part of

the fringe each step is located on, otherwise the offset reduction will fail. Furthermore, there is a radiation pressure effect from the light resonating in the cavity due to the intentional offset that must carefully be controlled. During this process, the differential arm alignment loop is engaged to further stabilize the cavity alignment.

First, CARM is sensed on transmission of the arm, and is at low bandwidth. As the CARM offset is reduced, the bandwidth can slowly be raised. When enough signal is present, the CARM sensor can be shifted to the reflection port, and the bandwidth increased. Again, the offset is measured, and the bandwidth increased further. At some point, the bandwidth is too large to be controlled digitally, so CARM is transitioned to analog controls. At this point, the arms can be brought fully to resonance, where CARM has a bandwidth of close to 16 kHz. DARM is also brought to resonance, controlled via anti-symmetric port sensors.

At this time, the DRMI controls can be transitioned back the “1f” signals, whose sensors are the lowest-noise and in-vacuum. A purposeful differential arm length offset is set, and all alignment controls are engaged. With the DARM offset light transmitting to the anti-symmetric port, the output mode cleaner is locked, transitioning the DARM control to “DC readout”.

This whole process is performed at low power, 2 W input from the pre-stabilized laser. Once full DC readout lock is achieved, the power is increased. At full operation, all length and alignment control bandwidths are reduced, actuators are moved to the lowest noise state, and squeezing is injected. Now, the interferometer is ready for observing.

This process can take anywhere from 30 minutes to a full hour, depending on the environmental conditions and starting point of each mirror. It must be performed after every “lock loss”. A lock loss occurs when a disturbance moves the interferometer controls out of the linear regime, where they cannot be returned to their operating point. At that moment, all resonance conditions are lost and the laser power drops out of the cavities. Lock losses are caused by environmental effects like earthquakes, wind and storms. Anthropogenic noise can cause lock losses too, if loud and close to the interferometer. Sometimes controls become unstable during operation. Whenever possible, lock loss causes are studied and mitigated.

Figure 5 shows a full schematic of the Advanced LIGO interferometers during Observing Run Four. This diagram also includes information about the seismic isolation systems for each mirror, suspensions, and some sensor locations.

LIGO Optical Layout, L1 or H1
 with Seismic Isolation and Suspensions
 aLIGO and A+ O4 Configuration
 G1200071-v8
 J. Kissel Jul 05 2023




Figure 5: A full schematic of the Advanced LIGO optical layout, including seismic isolation, suspension type, auxiliary cavities and sensors. From [19].

Chapter 2

Fundamental and Technical Noises of Gravitational-Wave Detectors

Ground-based gravitational-wave detectors face a variety of noise challenges. The limiting noises of these detectors are roughly divided into three categories: fundamental noise, technical noise, and environmental noise. Fundamental noise sources generally include noises that can be derived from first principles, such as quantum noise and thermal noise. Technical noise encompasses noises that result from the operation of the interferometer, such as controls noise and noise from electronics. Finally, environmental noise includes noises due to the surrounding environment, such as anthropogenic noise.

In order to achieve the incredible precision of these detectors, all sources of noise must be addressed, studied, and mitigated. Some noises are inherent to the design of the detector; the seismic noise level depends inherently on the choice of seismic isolation system. Other noises can be mitigated through careful study and clever methods, such as test mass charging noise. This chapter will attempt to summarize the wide array of limiting noise sources, and the related mitigation strategies.

2.1 Fundamental Noise Sources

Fundamental noise sources, as previously stated, are generally noises that can be derived from first principles. When designing a gravitational wave detector, these noises are often inherent to the design itself and once a detector is built they are unlikely to be improved without a significant intervention in the system.

Figure 6 shows the design sensitivity curve of Advanced LIGO, calculated into strain amplitude spectral density units $1/\sqrt{\text{Hz}}$. At low frequency (approximately 5-60 Hz), seismic noise, Newtonian noise, and suspension thermal noise all have the potential to be limiting noise sources. Quantum noise, specifically quantum radiation pressure noise, also can limit at low frequency. At mid-frequency (approximately 60-200 Hz), coating brownian noise limits. At high frequency (200 Hz and up), the limit is again quantum noise, specifically quantum shot noise. For completeness, excess gas noise, coating thermo-optic noise, and substrate brownian noise are included, although these are a factor of ten or more lower than the total noise estimate, and unlikely to be limiting noises in the current detector design.

The reasoning behind the somewhat strange-looking strain units is two-fold: strain, $h(t)$ is a unit-less quantity, and in this plot is depicted in amplitude spectral density units (ASD). An amplitude spectral density is defined as the square root of the power spectral density (PSD), which has units of $1/\text{Hz}$ in strain PSD. A power spectral density is defined as the Fourier transform of the autocorrelation function of a time series $x(t)$,

$$S_x(f) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x \star x(\tau) e^{-i2\pi f\tau} d\tau \quad (2.23)$$

where $x \star x(\tau)$ is the autocorrelation function,

$$x \star x(\tau) = \int_{-\infty}^{\infty} x(t)x(t+\tau)dt. \quad (2.24)$$

In this case, our time series is the strain time series $h(t)$, which is unitless, but any time series of any quantity can be transformed into a power spectral density.

2.1.1 Quantum Noise

A gravitational wave interferometer measures a passing gravitational wave by measuring a change in power at the antisymmetric port due to the phase shift of the light impressed by the wave. Therefore, gravitational wave measurements rely on the ability to measure power, or in other words, count incident photons on a detector. However, counting individual photons is a discrete process, and photons do not arrive on a photodetector continuously. The overall effect of this Poisson process is known as *shot noise*.

The probability distribution of a Poisson process is defined by the mean number of events in some time interval, λ ,




Figure 6: The Advanced LIGO design sensitivity curve. This curve includes known fundamental noise sources, and their estimated contribution based on the design of the detector. This design curve includes updated parameters based on an improved understanding of the coating thermal noise. Figure first presented in [20].

$$P(N|\lambda) = \frac{\lambda^N e^{-\lambda}}{N!}. \quad (2.25)$$

For a laser with power P and frequency ν_0 , there are N events in τ seconds or

$$\frac{N}{\tau} = \frac{P}{h\nu_0} \frac{\text{photons}}{\text{second}} \quad (2.26)$$

Therefore, the instantaneous power P measured will have some average power P_0 and uncertainty based on the uncertainty in the photon counting

$$P = P_0 \pm \sqrt{\frac{h\nu_0}{\tau} P_0}. \quad (2.27)$$

We can transform this measure of instantaneous power into a measure of the power spectral density of shot noise,

$$S_P(f) = 2h\nu_0 P_0 \left[\frac{\text{W}^2}{\text{Hz}} \right]. \quad (2.28)$$

In our gravitational wave detector with Fabry-Pérot arm cavity length L and cavity pole f_c , we care about how much this noise in power contributes to the strain amplitude sensitivity,

$$\sqrt{S_h(f)} = \frac{1}{L} \sqrt{\frac{h\nu_0}{2P_0}} (1 + if/f_c) \left[\frac{1}{\sqrt{\text{Hz}}} \right]. \quad (2.29)$$

Quantum shot noise dominates at high frequency where the detector response rises like f due to the cavity pole, $f_c \approx 430\text{Hz}$ in Advanced LIGO.

At low frequency, quantum radiation pressure noise limits. The discrete nature of photon arrival time on the test masses means that the masses experience a photon flux, which imparts momentum on their surface. Given the high power contained in the arm cavities of the Fabry-Pérot interferometer, this is a significant force. The contribution to strain dominates at low frequency, with a $1/f^2$ behavior,

$$\sqrt{S_h(f)} = \frac{1}{\pi M f^2 L} \sqrt{2h\nu_0 P_0} \left[\frac{1}{\sqrt{\text{Hz}}} \right]. \quad (2.30)$$

Here, the mass of the test masses is represented as M .

Both of these noises are represented within the “quantum” curve in Figure 6. By examining the two contributions, it is evident that by increasing the power in the cavity, the sensitivity at high frequency can be improved with the square root of the power. Meanwhile,

this would worsen the noise at low frequency by the same amount. However, increasing the mass of the mirrors would reduce the low frequency noise coupling.

Another method to reduce the presence of this noise in gravitational wave detectors involves quantum squeezing, which relies on the fact that these noises are correlated. Using the Heisenberg uncertainty principle, quantum shot noise and radiation pressure noise can be represented as phase and amplitude uncertainty,

$$\Delta\phi\Delta A \geq 1. \tag{2.31}$$

Quantum squeezing plays with this quantum nature of light by minimizing the noise in one quadrature, such as shot noise, while maximizing the noise in the other quadrature, radiation pressure noise. This is known as *frequency independent squeezing*. This is beneficial because it reduces high frequency quantum noise, however, it increases the impact of the low frequency noise. This method benefited gravitational wave detectors in a time when the low frequency noise was dominated by other sources. However, with increased demand for more significant quantum squeezing and an improvement of other low frequency noises, a new method is required. This method is referred to as *frequency dependent squeezing*. Noise is minimized or maximized in one quadrature or the other in a frequency dependent way, such that the resulting angle of squeezing rotates through frequency space. This requires an additional cavity in the gravitational wave detector design to perform this rotation. Figure 4 shows this additional filter cavity, which was installed in Advanced LIGO ahead of the fourth observing run.

Frequency dependent squeezing was deployed in the Advanced LIGO detectors during the fourth observing run. The implementation of this technology is discussed in Ganapathy et al [21], and the benefit is best summarized in Figure 7. The LIGO Livingston detector strain sensitivity without any squeezing is shown in the black trace. Frequency independent squeezing is injected in the green trace, showing a large reduction in noise at high frequency due to the reduction of shot noise, but at the sacrifice of the low frequency sensitivity from increased radiation pressure noise. With frequency dependent squeezing injected in the purple trace, the low frequency sensitivity is maintained while reaping the benefits at high frequency. The gray trace represents residual classical noise in the detector that cannot be suppressed with squeezing.




Figure 7: The LIGO Livingston detector improvement from different squeezing methods is shown. The black trace demonstrates the detector strain noise with no squeezing injected. The green trace shows the effect of frequency independent squeezing—improvement at high frequency due to shot noise reduction and a worsening of radiation pressure noise at low frequency. The purple trace measures noise with frequency dependent squeezing, improvement at high frequency is achieved while maintaining the low frequency sensitivity. The gray trace represents residual classical noise that cannot be changed with squeezing. Figure first presented in [21].

2.1.2 Thermal Noise

The test masses and suspensions used in gravitational-wave detectors are made of molecules constantly vibrating and creating some inherent uncertainty in their displacement. Einstein first characterized this phenomenon, called Brownian Motion [22]. Nyquist observed a similar phenomenon in fluctuating voltage across resistors [23]. The overall theory governing this behavior is a result of the Fluctuation Dissipation Theorem, first formalized by Callen et al [24, 25, 26, 27].

For a system at thermodynamic equilibrium, the spectral density of the thermal driving force is [28]

$$F_{th}^2 = 4k_B T \mathcal{R}(Z(f)) \quad (2.32)$$

where k_B is the Boltzmann constant, T is the temperature of the system and $\mathcal{R}(Z(f))$ is the real part of the system's impedance.

The power spectral density of the displacement of such a system is then

$$S_x(f) = \frac{k_B T}{\pi^2 f^2} \mathcal{R}(Y(f)) \quad (2.33)$$

and $Y(f) = Z^{-1}(f)$ is the system's admittance.

Therefore, regardless of a system's microscopic model, the macroscopic behavior can be described simply with a mechanical model that describes the system's impedance. This theorem links such disparate systems with common behavior such as those described by Einstein and Nyquist, mentioned above.

Moving specifically to the thermal noise from the test mass itself, we can consider the situation where a Gaussian beam is incident on a mirror. We follow the derivation as shown in Levin [29]. In our gravitational wave detector, the phase shift of the light contains information about the motion of the test mass surface, defined as

$$x(t) = \int f(\vec{r}) y(\vec{r}, t) d^2 r, \quad (2.34)$$

where \vec{r} is the transverse location of a point on the test mass surface and $y(\vec{r}, t)$ is the displacement of the boundary along the direction of the laser beam. The form factor $f(\vec{r})$ depends on the laser beam profile, and is normalized to one, $\int f(\vec{r}) d^2 r = 1$ (this value $f(\vec{r})$ is distinct from the frequency f).

To calculate the system admittance, $Y(f)$, we begin with a generalized force $F(t)$ that drives only the generalized momentum conjugate x . The Hamiltonian of such a system is

$$H_{int} = -F(t)x(t). \quad (2.35)$$

By defining $F(f)$ and $x(f)$ as the Fourier transform of the force and momentum conjugate, we can calculate the admittance to be

$$Y(f) = 2\pi i f \frac{x(f)}{F(f)}. \quad (2.36)$$

We then insert the definition of $x(t)$ into the Hamiltonian to deduce the nature of the driving force $F(t)$,

$$H_{int} = - \int F(t) f(\vec{r}) y(\vec{r}, t) d^2 r, \quad (2.37)$$

which leads us to understand that there is a pressure term $P(\vec{r}, t) = F(t) f(\vec{r})$ being applied to the test mass surface. We can apply an oscillatory pressure $P(\vec{r}, t) = F_0 \cos(2\pi f t) f(\vec{r})$, which leads us to infer that whatever pressure we apply to the test mass the test mass then dissipates, according to $Y(f)$. We call this dissipated power W_{diss} , and use our previous Equations 2.33 and 2.36 to determine the power spectral density of displacement as

$$S_x(f) = \frac{k_B T}{\pi^2 f^2} \frac{W_{diss}}{F_0^2}. \quad (2.38)$$

Assuming the test mass is structurally homogeneous, the internal friction comes from homogeneously distributed damping. This can be characterized via a material term known as the Young's modulus

$$E = E_0 [1 + i\phi(f)] \quad (2.39)$$

where $\phi(f)$ is referred to as the material's "loss angle". We then relate W_{diss} to the energy of elastic deformation

$$W_{diss} = 2\pi U_{max} \phi(f) \quad (2.40)$$

where U_{max} is the energy of elastic deformation when the test mass is maximally contracted or extended according to the oscillatory pressure being applied in $P(\vec{r}, t)$.

Remembering that we have applied a Gaussian beam profile to our test mass, we can then define the form factor $f(\vec{r})$ according to the Gaussian intensity profile

$$f(\vec{r}) = \frac{1}{\pi r_0^2} e^{-r^2/r_0^2} \quad (2.41)$$

where r_0 is the radius of the laser beam. In the case of our gravitational wave detector, the characteristic test mass size R is much larger than the laser beam size, so we can apply an expansion around r_0/R to define U_{max} [29].

$$U_{max} = \frac{F_0^2}{\pi^2 E_0 r_0} (1 - \sigma^2) I \left[1 + \mathcal{O}\left(\frac{r_0}{R}\right) \right], \quad (2.42)$$

where σ is the Poisson ratio of the material, E_0 the Young's modulus, and $I \simeq 1.87322$.

Applying this form to Equations 2.40 and 2.38, we derive the form of the power spectral density of thermal noise in displacement to be

$$S_x(f) = \frac{4k_B T}{f} \frac{1 - \sigma^2}{\pi^3 E_0 r_0} I \phi \left[1 + \mathcal{O}\left(\frac{r_0}{R}\right) \right]. \quad (2.43)$$

Thermal noise specifically due to the high-reflective mirror coatings is a main limiting noise source for gravitational wave detectors. This is shown in Figure 6, taking the form of strain amplitude spectral density

$$\sqrt{S_x(f)} = \frac{2}{\pi L f^{1/2}} \sqrt{k_B T \frac{1 - \sigma^2}{\pi^3 E_0 r_0} I \phi} \left[\frac{1}{\sqrt{\text{Hz}}} \right] \quad (2.44)$$

specifically emphasizing the $1/f^{1/2}$ dependence of this noise. However, this long derivation has emphasized that this behavior is generally applied to all materials. Figure 6 also shows the appearance of thermal noise due to the suspensions and the substrate itself. Due to the high quality of the suspension and substrate material (fused silica), these quantities are much lower than coating brownian noise, and are not as significant in the noise contribution.

Much more will be discussed about coating thermal noise in Chapter 3. However, the results of this section demonstrate that choosing materials for a gravitational wave detector suspension, mirror substrate, and mirror coatings is not a trivial matter. Careful study of materials with low internal friction must be made in order to understand the behavior of such materials and minimize their effect in noise.

Examining the dependencies of Equation 2.43, there are many factors that contribute to

coating thermal noise in a detector. Several factors can be grouped under “material properties” (σ , E_0 , ϕ), hence the utmost importance of studying low internal friction materials. Another is beam size; the inverse dependence on laser beam radius r_0 indicates that increasing the beam size will correspondingly reduce thermal noise (up to the point that $\mathcal{O}(\frac{r_0}{R})$ terms become significant). However, cavity beam size has other implications in terms of gravitational-wave detector control, which will also be discussed in later chapters.

Finally, as shown in Eq. 2.43, thermal noise also directly depends on temperature, T . The Advanced LIGO detector is operated at room temperature, and current upgrades to the facility maintain this design. However, cryogenic detectors are one option to overcome the thermal noise challenges, and upgrades such as LIGO Voyager and future detectors such as Einstein Telescope and some versions of Cosmic Explorer follow this principle [30, 31]. Currently, the Kamioka Gravitational Wave Detector (KAGRA), located in Japan, operates with cryogenic technology [15]. Nominally, any detector using cryogenic technology will require investigation into new coating materials, substrates, suspensions and possibly laser wavelengths to achieve this goal. KAGRA operates with sapphire substrates and suspensions, with a goal of achieving 20 K operation in its fully sensitive state [32]. This choice allows KAGRA to make use of $1\ \mu\text{m}$ lasers and accompanying technology, similar to room temperature detectors. Detectors like Voyager, ET, and versions of CE currently plan to use cryogenic silicon and focus on operating at 123 K, requiring a shift towards $2\ \mu\text{m}$ technology [33, 31, 30].

2.1.3 Seismic and Newtonian Noise

The Advanced LIGO detectors are ground-based gravitational wave detectors, and therefore experience the impacts of seismic motion and other effects of the earth. At 100 Hz, typical seismic motion is approximately $10^{-11}\ \text{m}/\sqrt{\text{Hz}}$, which is significantly higher than the desired sensitivity of $10^{-20}\ \text{m}/\sqrt{\text{Hz}}$ [34, 35].

Advanced LIGO applies multiple strategies to isolate seismic noise. Both passive and active isolation techniques are applied. Three main active platforms, called the hydraulic pre-isolater, single stage internal isolator, and the two-stage internal isolator provide active suppression of ground motion [36]. These active stages provide seismic noise suppression of multiple orders of magnitude above 1 Hz for both the main cavity and auxiliary cavity optics [36].

To further isolate the optics, the Advanced LIGO test masses are suspended on quadruple

pendulum suspensions which provide passive filtering of seismic motion above the pendulum resonances around a few Hz [37]. There is a $1/f^2$ noise reduction factor for every suspension stage, achieving the $1/f^8$ slope of the seismic noise demonstrated in Figure 6. Auxiliary cavity optics are also suspended, but due to the reduced sensitivity requirements compared to the test masses, these are triple, double, or single suspension isolation systems. Because these suspensions act as low pass filters, seismic noise below the suspension resonances at a few Hz is not isolated. Therefore, active control systems are required to suppress this noise further.

Newtonian noise, also called gravity gradient noise, results from the density fluctuations of matter around the detector. These density fluctuations cause gravitational gradients that act on the test masses and create excess motion. This noise has its strongest effect at low frequency, and Advanced LIGO is not limited by it and currently employs no technology to mitigate it. However, future gravitational wave detectors are likely to be limited by this noise source at low frequency [31]. Mitigation strategies involve measuring the density fluctuations with an array of sensors and subtracting the noise from data [38, 39, 40].

2.1.4 Other Fundamental Noises

Excess gas in the chambers of the gravitational wave detector can cause noise in two main ways. Scattered light of the interferometer beam off particles can cause excess phase noise. Also, gas particles impinging on the suspended mirrors cause ambient motion of the test masses due to the imparted momentum. Advanced LIGO is held at ultra-high vacuum, reducing the effect of this noise, also demonstrated in Figure 6.

Thermo-optic and thermo-refractive noise are linked due to the fact that temperature fluctuations are responsible for both. Evans et al [41] developed a formalism for handling both at once, combined under one heading as thermo-optic noise as in Figure 6. Due to the opposing nature of each of these sources of noise, it is possible to optimize a coating material such that the thermo-optic and thermo-refractive noises cancel, driving down their contributions in gravitational wave detectors.

2.2 Technical Noise Sources

The fundamental noise budget shown in Figure 6 immediately becomes insufficient to describe gravitational wave detector noise once a real detector is built. That is because the




Figure 8: The LIGO Hanford detector noise budget for observing run four. A variety of limiting fundamental and technical noises are demonstrated. Additionally, a significant amount of unknown noise is demonstrated between 15 and 100 Hz, shown in the discrepancy between the red measured noise trace and the black sum of known noises trace. The O4 Livingston detector sensitivity and the O3 Hanford detector sensitivity are shown for comparison. Measured in December 2023 [42, 43].

operation of real detectors requires many additional components that can cause noise. These noises are generally considered “technical noises” due to the fact that they occur due to the operation of the detector and are considered reduceable with continued study and mitigation once a detector is built. To see the full impact of many sources of noise in gravitational wave detectors, a figure like Figure 8 demonstrates these noises, in addition to fundamental noises.

Furthermore, due to the technical nature of these noises, they can often be measured in the real system. Therefore, unlike the derivations in the previous section that are used to make analytical estimations, this section presents measured noise levels in the Advanced LIGO detectors.

The noise contribution of a witness channel to the differential arm length can be measured in two similar, but distinct ways. For both methods, we define some witness channel w and the differential arm length channel d . One method exclusively measures a linear coupling of w to d by measuring a transfer function,

$$\frac{|S_{wd}(f)|^2}{S_w(f)} = \gamma_{wd}^2(f)S_d(f). \quad (2.45)$$

Here $S_{wd}(f)$ is the cross-spectral density of the witness to the length channel, and S_x is the power spectral density of the particular channel. The strength of the relation is governed via γ^2 , or the coherence in power. This method is only useful if the witness channel has a strong coherence with the length channel, and only predicts linear relationships.

A more broad method involves using the relation [44]

$$S_{w \rightarrow d}(f) = G(f)S_{w,0}(f) \quad (2.46)$$

where $S_{w,0}$ is the power spectrum the witness channel under normal operating conditions. The coupling factor $G(f)$ is estimated by

$$G(f) = \frac{S_{d,\text{exc}}(f) - S_{d,0}(f)}{S_{w,\text{exc}}(f) - S_{w,0}(f)} \quad (2.47)$$

where $S_{d,\text{exc}}$ and $S_{w,\text{exc}}$ are the power spectra of the length channel and witness channel during some excitation. The excess noise in each channel during an excitation above ambient is compared with the ambient noise in the channel to derive the overall coupling. This method allows the calculation of both linear and nonlinear couplings to be measured, and can be used to measure channels with very low coherence to the main channel. However, some care is required with this method. Different excitation amplitudes should be applied and compared

to ensure the coupling does not depend on excitation amplitude alone. Both this method and the transfer function method are regularly used to estimate technical noise contributions to gravitational wave strain.

Sources of technical noise are many, but this section will attempt to address the most common and well-understood noises that present in gravitational wave detectors.

2.2.1 Auxiliary Length Control Noise

The differential arm length, or DARM, is the channel that achieves coupling of the gravitational wave signal, and therefore the channel from which strain is derived. This channel is a main figure of merit for understanding how length noises couple to gravitational wave strain. The residual displacement noise of each length degree of freedom can couple into the differential arm length readout.

The coupling of Michelson noise, or MICH, occurs because Michelson motion between the beamsplitter and input test masses creates phase fluctuations in the light at the gravitational-wave readout. The Michelson cavity motion is indistinguishable from differential arm cavity motion at the antisymmetric port *except* for the fact that DARM motion is amplified due to the finesse of the Fabry-Pérot cavities, and MICH motion is not. Therefore, the coupling factor of MICH motion is the inverse of the arm cavity gain [13],

$$G_M = \frac{1}{G_a^2} = \frac{\pi}{2\mathcal{F}}, \quad (2.48)$$

where \mathcal{F} represents the arm cavity finesse. For Advanced LIGO, arm cavity finesse is approximately 440, so $G_M \simeq 1/280$.

The signal-recycling cavity length couples into differential arm length in two ways. First, the presence of a DARM length offset, which propagates DARM light through the signal-recycling cavity, causes a radiation pressure coupling with a $1/f^2$ shape [13, 45],

$$G_S^{(1)} = 0.012\text{m/m} \left(\frac{10\text{Hz}}{f} \right)^2 \left(\frac{P_{\text{arm}}}{750\text{kW}} \right) \left(\frac{\delta L_-}{10\text{pm}} \right) \left(\frac{\mathcal{F}}{450} \right). \quad (2.49)$$

This coupling has been parametrized into typical Advanced LIGO design values. The factor δL_- represents the DARM length offset, and P_{arm} represents the intra-cavity arm power.

The second SRCL coupling results from small detunings of the SRC from resonant side-band extraction, and instead rises like f^2 [13, 45],

$$G_S^{(2)} = 3 \times 10^{-5} \text{m/m} \left(\frac{f}{100\text{Hz}} \right)^2 \left(\frac{\phi_s}{10^\circ} \right) \left(\frac{\delta L_-}{10\text{pm}} \right) \left(\frac{\mathcal{F}}{450} \right), \quad (2.50)$$

where ϕ_s represents the SRC detuning.

The coupling of the power-recycling cavity length and the common arm length do not have explicit forms in the same way the length couplings above do. These are common lengths, and in principle should not couple at all into differential arm length readout. In the case of PRCL, this coupling most likely results in cross-coupling between PRCL and the other two corner lengths, MICH and SRCL. This can result from poor sensor diagonalization, or other opto-mechanical coupling. CARM primarily couples at high frequency due to arm imbalances that result in contrast defect light reaching the anti-symmetric port.

The contributions of the three corner (MICH, PRCL, SRCL) length controls are shown in Figure 8, summed into one length control contribution. These noises limit at low frequency.

Length Feedforward

The coupling of noise from MICH and SRCL creates a significant portion of noise in DARM if left completely unmitigated. The current loop design for the Michelson length control and signal-recycling cavity length control include low pass filtering above the loop bandwidth to suppress reinjected sensing noise into DARM. In Advanced LIGO, the loop bandwidth for MICH is approximately 10 Hz, and SRCL a few tens of Hz, meaning most of the noise coupled back into DARM is sensing noise for the length degrees of freedom. Specifically, this is shot noise from the length readout of each of these length control loops.

Because this length noise that is impressed on DARM is witnessed by other sensors, it is possible to apply feedforward cancellation of this noise. This is done in real-time mechanically by feeding the expected sensing noise forward to the test masses.

The noise coupling $G(f)$ couples residual length noise in the differential arm length as demonstrated in Equation 2.46. That residual length noise can be written as

$$r = \frac{x + Kn}{1 - H} \quad (2.51)$$

where K is the control gain, H is the total open loop gain of the auxiliary length control loop, x is the free-running displacement noise and n is the sensing noise. Some feedforward filter, $F(f)$ can be applied to the control signal k and fed to the differential arm length




Figure 9: Coupling of the Michelson length to differential arm length without and with feedforward applied. Measurement procedure and data described in [46, 47].

control to cancel out the length noise. However, the control signal k and residual noise r are different quantities

$$k = \frac{Px + Kn}{1 - H} \neq r \quad (2.52)$$

due to the fact that the displacement noise also sees the plant P that is being controlled in the loop. Therefore, feedforward control can be chosen to either subtract sensing noise n via applying $F = G$ or displacement noise x by applying $F = G/P$, but not both. Furthermore, applying a cancellation to sensing noise can worsen the coupling of displacement noise, and vice versa.

For Advanced LIGO, the coupling of the auxiliary length is mostly above the control bandwidth of the loops, and therefore it is advantageous to apply feedforward to cancel sensing noise. Some worsening of the displacement noise is witnessed, but this is below the gravitational wave band, and therefore does not effect the noise appreciably. To minimize the injection of excess displacement noise, a high pass filter can be applied to the feedforward around the crossover between the regions where displacement and sensing noise dominate. However, this must be done carefully, as the additional phase loss from a high pass filter can

compromise the efficacy of the feedforward that requires precise fitting such that F matches well to G .

The worsening of displacement noise due to feedforward should be considered more carefully for future gravitational wave detectors that seek to achieve detections at even lower frequencies than Advanced LIGO. This is further considered in Chapter 8.

Figure 9 shows the improvement of MICH to DARM coupling with feedforward applied. The coherence of the measurement is poor below 7 Hz, but the increasing trend of the coupling below the MICH control bandwidth (8 Hz) is evident. Feedforward improves the coupling above 8 Hz, but worsens it below. A factor of one hundred or more sensing noise suppression is possible with feedforward.

2.2.2 Alignment Control Noise

In addition to the length control of the cavities listed above, the alignment must be controlled. This can be thought of as controlling the individual angular position of each cavity mirror. However, similar to length control, this can be thought of in sensible global control bases.

To understand alignment noise coupling, a change in mirror angle is directly a change in length [48]

$$\Delta L = d_{spot} \times \theta_{mirror} \quad (2.53)$$

where d_{spot} represents some beam position relative to the mirror center of rotation and the mirror angle θ_{mirror} is relative to the input axis of the cavity. This relationship becomes more complex with the awareness that both d_{spot} and θ_{mirror} have some frequency dependence. The work in Barsotti et al [48, 49] defines the resulting coupling as a bilinear process, and related through the approximate convolution of the two spectra

$$\Delta L(f) = d_{spot}(f) * \theta_{mirror}(f) \approx d_{spot}^{RMS} \times \theta_{mirror}(f) + \theta_{mirror}^{RMS} \times d_{spot}(f). \quad (2.54)$$

For the arm cavities, this is a direct coupling, as a change in angle of one of the mirrors is a direct change in the differential arm length. For auxiliary cavity alignment coupling, the alignment noise couples to DARM with the same coupling factor as the auxiliary cavity length [48, 49].

The majority of this dissertation will focus on alignment control topology and alignment noise couplings. Therefore, information in this introductory section is brief, and the author

highly encourages the reader to continue forward in this work in order to learn more of the exciting and tumultuous world of alignment sensing and control.

The total contribution of all alignment degrees of freedom to the strain sensitivity is shown in Figure 8. These noises limit at low frequency.

2.2.3 Laser Noises

Laser noise combines noise sources that are coupled through the laser beam itself. These noises are frequency noise, intensity noise, and beam jitter. In Figure 8, laser frequency and intensity noise are combined into one “laser noise” trace, while beam jitter is demonstrated separately.

Frequency Noise

Frequency noise is a common noise, in that it should not transmit to the anti-symmetric port of the interferometer. However, due to both intentional and unintentional asymmetries in the Fabry-Pérot arm cavities, it can contaminate gravitational-wave strain. Following the work in the thesis by E. Hall [45], the coupling of frequency noise can be understood by writing down the equations of an asymmetric Michelson cavity with differential mirror reflectivities r_x and r_y and arm lengths l_x and l_y ,

$$\begin{aligned} E_{out}(\omega) &= \frac{E_{in}(\omega)}{2} [r_x e^{2i\omega l_x/c} + r_y e^{2i\omega l_y/c}] \\ &= E_{in}(\omega) e^{2i\omega l_+} [r_+ \cos(2\omega l_-/c) + i r_- \sin(2\omega l_-/c)] \end{aligned} \quad (2.55)$$

where $r_{\pm} = (r_x \pm r_y)/2$ and $l_{\pm} = (l_x \pm l_y)/2$. The cosine term shows the response to the usual Michelson offset of length l , which occurs through both the Schnupp asymmetry and the differential arm length offset. The sine term results from “contrast defect” light propagating to the anti-symmetric port. Contrast defect occurs due to arm mismatch, such as reflectivity mismatch between the Fabry-Pérot arms, or other mode mismatch that reduces common-mode rejection.

Frequency noise is the most significant contribution to the “laser noise” term shown in Figure 8. It mostly limits at high frequency, but there is some frequency noise present below 100 Hz. The high frequency noise is most easily improved with mode matching adjustments that reduce contrast defect light at the anti-symmetric port. The low frequency noise coupling is currently not well understood, but is easily reduced through improvements in the

auxiliary length feedforward, leading to the conclusion that it couples through the auxiliary length degrees of freedom.

Intensity Noise

There are a few different paths to couple intensity noise in a gravitational wave detector. The DC readout scheme allows input intensity noise to modulate the DARM offset light propagating to the output port. Contrast defect also further propagates intensity noise, as input intensity noise can modulate any contrast defect exiting to the anti-symmetric port. Further arm mismatch can exacerbate this, as common-mode rejection is reduced with further asymmetry. Also, intensity noise of the RF sidebands can propagate to the output port if not sufficiently suppressed by the output mode cleaner. Since the sidebands are anti-resonant in the arms, they do not see the cavity pole, and their intensity noise can propagate directly to the gravitational-wave readout. Intensity noise is included in the “laser noise” projection in Figure 8, but the measured contribution is small.

Beam Jitter

The coupling of vibrational noise due to first order misalignment of the input beam is referred to as beam jitter [50]. In Advanced LIGO, vibrational noise along the input chain from the output of the pre-stabilized laser to the arms causes this noise. The coupling results from defects in the mirrors that cause mode mismatch between the arms.

Figure 8 shows that the largest beam jitter coupling occurs at mid-frequency. The sharp drop-off at 900 Hz is a result of the limitations of the injection channel used to make the jitter projection. One notable source of jitter noise is the pre-stabilized laser periscope, which has a mechanical resonance at 120 Hz. This resonance is driven by surrounding noise sources, such as HVAC fans. The vibration of the periscope is imprinted on the laser beam as it is injected into vacuum, hence the jitter peaks above 100 Hz.

2.2.4 Other Technical Noises

There are many more technical noises that can or do limit gravitational wave detectors. Electronics noise, such as dark noise in photodetector readouts, quantization noise from digital-to-analog converters, and electrostatic coupling due to wiring are several noises that

impact Advanced LIGO. Technically, environmental noise is a separate category, but technical limitations in cavity baffling can cause scattered light due to equipment shaking from environmental noise.

Figure 8 specifically shows the contribution of DAC quantization noise from the penultimate mass actuator of the test masses, which is close to being a limiting noise source at low frequency. Photodetector dark noise is also measured, but unlikely to limit the sensitivity. Suspension damping noise is a result of noisy local sensors on the test mass quadruple suspensions. The noise from the local sensors is reinjected into other control degrees of freedom that couple directly into the gravitational wave readout.

2.3 Sensitivity Figures of Merit

There are many ways to quantify gravitational-wave detector sensitivity beyond examining the strain noise. Ultimately, it is an astrophysical instrument, and there are a multitude of astrophysical signals to be measured. Depending on the desired signal, sensitivity figures of merit can provide motivation to improve ranges or features of the strain sensitivity spectrum. In this section, I will explore a few different ways to quantify gravitational-wave detector sensitivity.

When quantifying detector sensitivity, the assumption is first made that the detector noise is stationary and therefore the strain noise spectral density can completely characterize the sensitivity of the detector. For such a given strain noise spectral density $S_n(f)$, the average signal-to-noise (SNR) ratio of a gravitational-wave signal is [51]

$$\langle \rho \rangle = \sqrt{4 \int_{f_{low}}^{f_{high}} \frac{|\tilde{h}(f)|^2}{S_n(f)} df} \quad (2.56)$$

where $\tilde{h}(f)$ defines some waveform. This equation leads us to distance when some SNR is chosen and Equation 2.56 is solved for distance D , which is part of the implicit definition here. Therefore, detector range is intimately tied with the source parameters. Furthermore, the integration range shown in Equation 2.56 is between two arbitrary frequencies, f_{low} and f_{high} . The integration range may be as broad as to integrate over all frequencies, but in our real detector we may instead choose realistic frequencies that demonstrate the band over which we are sensitive. For Advanced LIGO, f_{low} can be as low as 10 Hz. The strain sensitivities shown so far include the sensitivity up to a few kilohertz, but there is a more

physical limit that we can choose for f_{high} . The innermost stable circular orbit, or ISCO, often defines the upper limit of integration.

$$f_{isco} = \frac{c^3}{6\sqrt{6}\pi GM} \quad (2.57)$$

where M is the total mass of the system, G is the gravitational constant, and c the speed of light.

One of the most common sensitivity figures of merit currently used for gravitational-wave detectors is referred to as the “BNS range”, short for “binary neutron star range”. This describes the distance at which two neutron stars, of $1.4 M_{\odot}$ (1.4 solar masses), can be measured in a detector with a signal-to-noise ratio of 8, averaged over the whole sky. For binary neutron stars, $f_{isco} = 1570$ Hz.

Defining a waveform for a particular system can be a complex process, so for simplicity we consider this approximation [51]

$$\tilde{h}(f) = \frac{1}{D} \left(\frac{5\pi}{24c^3} \right)^{1/2} (G\mathcal{M})^{5/6} (\pi f)^{-7/6} e^{i\Psi(f,M)} \quad (2.58)$$

which assumes an optimally oriented and optimally located equal-mass binary. The chirp mass \mathcal{M} is

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (2.59)$$

and Ψ is some real function of frequency, parameterized by the total mass M . Notice that through this definition of the waveform, we find the explicit dependence on distance. Applying the other parameters we previously mentioned, and solving for D , the range is

$$D = \frac{1}{8} \left(\frac{5\pi}{24c^3} \right)^{1/2} (G\mathcal{M})^{5/6} (\pi)^{-7/6} \sqrt{4 \int_{f_{low}}^{f_{isco}} \frac{f^{-7/3}}{S_n(f)} df}. \quad (2.60)$$

In Advanced LIGO, a more specific definition of BNS range is called “SenseMon range” and has some historical importance in its usage. The significant difference is an additional division of a factor of 2.6; this is related to accounting for sources that may not be optimally oriented or located due to the actual detector orientation on the earth [51]. The quantity demonstrated here is usually referred to as “horizon distance”, whereas if divided by 2.6 it is “range distance” [52]. Further parameterization of Equation 2.60 can give the result in megaparsecs, or Mpc, which is useful astrophysical distance parameter.

Why use binary neutron star systems to quantify this metric? Any system can be used, but BNS systems are the most useful for this characterization due to the integration range. The ISCO frequency is inversely proportional with total mass, so a two $30M_{\odot}$ black hole binary system integrates over a much smaller range of frequencies for the same detector sensitivity. A binary neutron star system gives the fullest accounting of the detector sensitivity range for a known astrophysical system that could be measured in a ground-based gravitational-wave detector.

Figure 10 shows the sensitivity improvement of the LIGO Hanford detector between observing run three to observing run four, compared with the modeled sensitivity based on the O4 target detector parameters [53, 54]. The left plot shows the integrated range for a $3M_{\odot}$ black hole merger and the right the integrated range for a $100M_{\odot}$ black hole merger. The y-axis plot units are defined such that the area under each curve gives the total range to that type of event. Overall, the LIGO Hanford detector is more sensitive to a lighter system than a heavier system. Also, these plots indicate the best region to target to improve detector sensitivity to intermediate-mass black hole mergers is from 10-100 Hz, compared to benefits up to 1 kHz for very light systems. The improvements from O3 to O4 were across the entire detection band, benefitting both types of systems. Compared to the full design target, the most improvement would be achieved by targeting low frequency noise, opening the detection significantly to high-mass systems.

Furthermore, for the current range of the Advanced LIGO detectors, figures such as “range distance” or “horizon distance” are sufficient for quantifying distance to sources. These numbers can then be translated into a spacetime volume per unit detector time, known as the “comoving volume” [52]. However, as detectors become increasingly more sensitive, this type of volume calculation becomes inaccurate—one must account for cosmological effects. This is referred to as the redshifted volume, and involves an additional factor of $(1 + z)$ where z is the redshift [52]. Current Advanced LIGO sensitivities do not require this additional correction, but future detector sensitivities, such as that of Cosmic Explorer [6], likely will.

A comparison of the comoving horizon volume of the LIGO Hanford detector from O3 to O4 versus source-frame mass is shown in Figure 11. Here, the volume ratio is shown to have double everywhere in the upgrade from O3 to O4 [53]. In this way, the visualization of detector sensitivity is recast in terms of the astrophysical sources that could be measured in the detector.




Figure 10: A range integrand comparison of the LIGO Hanford detector O3 range, pre-O4 range, and the full design sensitivity for the A+ LIGO upgrade. The left plot shows the integrand for a $3 M_\odot$ system and the right a $100 M_\odot$ system. The heavier system motivates sensitivity at low frequency, while the lighter system benefits from improvements everywhere in the band. Plots made by E. Hall in [53].




Figure 11: A comparison of the comoving horizon volume versus different mass systems. The improvements of the fourth observing run improved the volume sensitivity everywhere. The A+ design sensitivity volume is shown for reference. Plot made by E. Hall [53].

For completeness, the waveform used to calculate the range in Equation 2.60 is an approximation. Effects that can adjust this metric include the consideration of spin, tidal effects (for matter-based systems only), or other waveform parameters. Although this range metric is limited to 1570 Hz, other tidal interactions in the post-merger signal can be detected at higher frequencies. Clark et al [55] explores the physics of binary neutron star systems in the post-merger signal, which can provide information about the equation of state of matter at high densities. Specifically, peak frequency locations above 1 kHz can provide tight constraints on neutron star radius. This is further discussed in Alford et al and Raithel and Most [56, 57]. While not yet observed in gravitational wave detectors, there is possibly a high frequency dependency in core-collapse supernovae [58]. Overall, despite this constraint on binary neutron star sensitivity, the high frequency portion of the gravitational wave system can yield key physics, most of which is not yet observed.

Overall, understanding the dependence of detector sensitivity on the different systems Advanced LIGO can measure can further motivate where to prioritize range improvements. High frequency range improvements will benefit low mass systems such as neutron stars more than high mass systems such as black holes. While current generation detectors seek to target a variety of gravitational-wave sources, future detector designs can target certain frequency

bands for specific scientific goals, such as NEMO [59] which targets post-merger signals, or the low frequency Einstein Telescope [30], which seeks to achieve the lowest frequency band yet in ground-based interferometers.

The range quantities listed here are not the only quantities for consideration in the astrophysical impact of ground-based gravitational wave detectors. Section 8.2 further explores sensitivity figures of merit in the context of low frequency noise.

Chapter 3

Coating Characterization

As described in Section 2.1.2, noise from thermal motion in mirror substrates, mirror coatings, and suspension fibers all have the potential to be limiting noises in advanced gravitational wave detectors [28]. The most significant of these sources in detectors like Advanced LIGO is coating thermal noise (CTN), resulting from thermal motion in the high reflective optical coatings on the input (ITM) and end (ETM) test masses in the Fabry-Pérot arms. This is a noise source that limits gravitational wave detectors in the most sensitive band, around 100 Hz. The Advanced LIGO+ design noise budget demonstrates the contribution of coating thermal noise, shown in Figure 12.

In order to reduce this noise in the detector, extensive investigations are underway to develop optical coating materials with low mechanical loss. This chapter discusses efforts to investigate the mechanical loss of new amorphous materials for room temperature gravitational wave detectors. By using a room temperature Gentle Nodal Suspension experiment [60], new coatings can be characterized in an efficient and repeatable way. This chapter presents such an experimental set up at Syracuse University, and the work done to develop and characterize the system. Then, this chapter presents investigations of one new possible coating material: mixtures of zirconia and titania. Unfortunately, while this material initially showed favorable results, further investigation demonstrated that contaminants in the material made it impossible to conclusively determine the effectiveness of these mixtures as low loss coatings. These results are presented as a proof-of-principle that this experiment is a useful one for characterizing optical coatings, and the chapter closes with a quick summary of the complete procedure for characterization, assuming the material has favorable properties.




Figure 12: Upgrades to the Advanced LIGO detectors, referred to as A+, involved increasing the injected quantum squeezing, which reduces the noise depicted in the purple curve. To make full use of the improved sensitivity due to quantum squeezing, a factor of two reduction in coating thermal noise is required, reducing the red curve to the level shown. Plot taken from [54].

3.1 Coating Thermal Noise in Gravitational Wave Detectors

Advanced LIGO coatings are Bragg reflectors made out of alternating layers of low- and high-index of refraction materials. The low index is silica, SiO_2 , and the high index material is tantala doped with about 25% titania, $\text{TiO}_2:\text{Ta}_5\text{O}_2$ [61, 62, 63]. These are amorphous coatings deposited using ion beam sputtering to create a high reflective (HR) “stack” that meets reflectivity requirements. For The aLIGO ITMs, this corresponds to a transmission of 1.4%, for ETMs, 5 ppm [64, 65]. The long derivation in Section 2.1.2 shows that thermal noise arising from an incident gaussian beam depends on the loss angle of the material

[29, 25]. To re-emphasize the result of that derivation, the power spectral density of thermal noise, [29] is

$$S_x(f) = \frac{4k_B T}{f} \frac{1 - \sigma^2}{\pi^3 Y r_0} I \phi \quad (3.61)$$

where σ , Y , and ϕ are the Poisson ratio, Young's modulus, and dissipational loss angle of the coating material respectively, and r_0 is the radius of the incident laser beam. k_B and T represent the Boltzmann constant and temperature of the system, and $I \simeq 1.87322$. Note that for simplicity, this form does not show higher order forms of r_0/R . This form of the coating thermal noise is also an approximation, as the materials in the HR stack have different material properties (such as σ , Y and ϕ), and both bulk and shear loss mechanisms within the coating layers contribute different values of mechanical loss ϕ to the overall loss estimate.

Equation 3.61 shows the power spectral density of thermal noise; the *amplitude* spectral density shown in Figure 12 requires applying a square root to Eq. 3.61. Roughly, the expected form of the thermal noise contribution from coatings in the Advanced LIGO strain amplitude is a noise with a slope of about $1/f^{1/2}$. Due to the inverse relationship of this noise with beam size, the Advanced LIGO beam size is set to be as large as possible to minimize the effect of coating thermal noise in the detector. Advanced LIGO uses a beam size of approximately 5 cm on the ITMs and 6 cm on the ETMs [48].

Once design considerations like beam and test mass size are made, the major contributions to thermal noise arise from the material properties of the chosen coating material. Of the two materials which comprise aLIGO coatings, titania-doped tantala has a higher loss angle, and it is the loss from this material that sets the current CTN limit in aLIGO. To quantify the expected CTN level in aLIGO, small witness samples are made during the coating process of the ITMs and ETMs. These smaller samples are measured via direct coating thermal noise measurements. The total estimated CTN for the current Advanced LIGO coatings is $1.1 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}}$ at 100 Hz [66]. These results correspond to a titania-doped tantala loss angle of about $4 \times 10^{-4} \text{ rad}$ [66], whereas the loss angle of silica has been shown to be about $1 \times 10^{-4} \text{ rad}$ [61].

Future upgrades of Advanced LIGO, including the A+ upgrade, shown in Figure 12 and A# upgrade require a reduction in thermal noise to allow full use of the increased sensitivity and bandwidth provided by advances in quantum squeezing [54, 67]. For both, the required reduction in coating thermal noise is a factor of 2. In order to achieve this reduction, it is

necessary to reduce the total loss from the coating stack by a factor of 4. Given the low loss properties of silica, investigations into lower-loss optical coatings focus on the high index of refraction material.

3.2 Quantifying the Mechanical Loss of Coating Materials

3.2.1 Vibrational Mode Ringdowns

To evaluate the thermal noise of a material and its usability as a low-noise optical coating in gravitational wave detectors, properties like the mechanical loss angle ϕ must be measured. As demonstrated in Equation 3.61, the coating thermal noise scales directly with loss angle. An experiment that measures the vibrational mode ringdown to quantify ϕ is therefore an indirect measurement of the thermal noise of the system. Alongside ϕ , other parameters must be determined, such as the Young's modulus and Poisson ratio, to fully determine the thermal noise present in the material.

The dissipational loss angle ϕ defines the rate of energy lost per cycle for a given resonant mode. Specifically, if a system is oscillating at one of its eigenfrequencies, and there is any form of internal friction to damp the oscillation, the vibrational amplitude will decay with time. If the system is of sufficiently high quality (low internal friction), we can assume that the mode amplitude is relatively constant over one period of oscillation, T . Therefore, the change in energy between two cycles of oscillation ΔE can be written as [68]

$$\Delta E = E(t) - E(t - T) = T \frac{dE}{dt} = \phi E(t). \quad (3.62)$$

If the frequency of the mode is measured, and the loss of energy in that mode (the ring down of the mode) is measured over time, the loss angle can be determined. The loss angle of the n th mode of a system [69] is

$$\phi_n = \frac{1}{Q_n} = \frac{1}{\pi f_n \tau_n} \quad (3.63)$$

where Q_n is the quality factor of the mode, determined by the frequency of the mode f_n and the decay time constant of the mode τ_n .

To test the loss angle of a particular coating material, a thin film may be deposited on a larger disk, usually referred to as the “coating” and “substrate”, respectively. The substrate

is usually a high quality resonator, such as a disk of fused silica, and the coating some material with promising optical properties.

In such a system, the total energy is a sum of the energy contained within the coating and the disk, $E = E_c + E_s$. The loss angle of the whole system will be a combination of the individual loss angles of the coating and substrate,

$$\phi E(t) = \phi_c E_c(t) + \phi_s E_s(t). \quad (3.64)$$

This is often recast in terms of the coating loss angle, and defining the *dilution factor*, $D = \frac{E_c}{E_s + E_c}$,

$$\phi_c = \frac{\phi - (1 - D)\phi_s}{D}. \quad (3.65)$$

Therefore, an experiment that drives vibrational modes of a system and observes their decay time τ_n , can quantify the loss angle. A Fourier transform of the time domain data can provide information about each eigenfrequency f_n of the system. If this measurement is performed first on an uncoated disk to determine ϕ_s of each resonant mode, and then again to determine ϕ of those modes when coated, the loss of the thin coating can be determined. One such method of driving vibrational modes of a system and observing their ringdowns is a *Gentle Nodal Suspension* experiment.

3.2.2 Gentle Nodal Suspension




Figure 13: A simple diagram of a Gentle Nodal Suspension. A disk is balanced on a clamped sphere at its center. Image taken from Cesarini et al [60].

A Gentle Nodal Suspension (GeNS) can be used to characterize the mechanical loss of substrates and coatings, and is a method first demonstrated by Cesarini et al [60]. A thin optical sample is balanced on a clamped lens at its center within a vacuum chamber as in Figure 13. An electrostatic comb exciter is placed near the surface of the sample to drive the vibrational modes of the sample. The resulting motion of the vibrational resonances can be read out using an optical lever scheme, see Figure 14. This suspension type is considered “gentle” as no welding or other attachment to the sample is required.




Figure 14: A schematic diagram of the optical lever scheme of the Gentle Nodal Suspension. A comb exciter drives mechanical modes of a sample, and the optical lever beam is steered to a QPD that detects the beam motion due to the excited vibrational modes. Image taken from Didio PhD thesis [70].

The ease and repeatability of the gentle suspension is one benefit to this experimental set up. One drawback is that the contact of the suspension with the sample center precludes driving any vibrational mode that requires motion of the center point, such as drumhead modes, and any mode that requires a nonzero first derivative of the motion.

Figure 15 shows a chart of the modal basis of a fused silica disk 76 mm (3 inches) in diameter and 1 mm in thickness. Modes with $n = 0$ or $n = 1$ cannot be driven in the GeNS experiment. The modes where $n = 0$ require motion at the center point, and modes with $n = 1$ require a nonzero first derivative of motion at the center point. However, many other vibrational modes are accessible in this set up. Only modes less than 30 kHz are demonstrated in this chart, as the highest sampling rate in the GeNS experiment is 60 kHz.

For the Syracuse GeNS design, a 1-inch polished fused silica lens is clamped in the center of a large vacuum chamber. The vacuum chamber is controlled via a turbo pump that brings

Modes below 30 kHz

Silica disks: $\varnothing 76 \text{ mm} \times 1.0 \text{ mm}$

Images show absolute value of displacement.




Figure 15: A depiction of the modes of a 76 mm diameter (3 inches), 1 mm thick fused silica disk. The lighter areas show greater displacement, while the darker areas show less displacement. The darkest shading indicates no motion, and are the location of the nodes of each mode. Modes with $n = 0$ or $n = 1$ cannot be driven in a Gentle Nodal Suspension experiment, as the placement of the center of the sample on the lens damps any motion at the center.

the vacuum level to around 1×10^{-7} torr. This pump is turned off for data collection, and the vacuum level is stable for over 24 hours. A 3-inch optical sample is placed on a metal tray centered over the suspension lens. When the chamber is closed, a motor lowers the sample tray down until the sample is suspended directly on the lens (and no longer in contact with the tray). Steering mirrors point a 2 mW HeNe laser beam into the chamber to be incident on the sample, and then guide the reflected beam back out of the chamber to be incident on a quadrant photodetector (QPD). The resulting lever arm around is about 1 m. Once the sample is suspended, steering mirrors are adjusted to ensure that the beam is well centered onto the QPD. An electrostatic comb exciter is suspended a few millimeters away from the surface of the optical sample.

A LabVIEW program called MultiQ, written by Steven Penn, is used to take measurements of mechanical ring downs of the optical sample [71]. The program generates a band-limited white noise spectrum of a given amplitude that is applied to comb exciter through a high voltage amplifier for up to 5 minutes. Typical excitation amplitudes are around 1600 V, and are adjusted depending on the thickness of the sample and the resulting signal-to-noise ratio of excited modes. The readout scheme of the GeNS includes a quadrant photodiode whose four outputs are sent to a summing board that calculates X, Y and SUM signals via combination of the four quadrants. The resulting signals are then sent to a BNC 2090A analog-to-digital board. The data is collected using MultiQ and saved along with metadata that describe the parameters of each measurement.

The ease of sample placement and stability of the vacuum ensures that measurements can be easily repeated to achieve good measurement statistics. If the vacuum level degrades significantly, the sample can be re-captured by the sample tray while the turbo pump is run, and then re-suspended once the desirable vacuum level is achieved. It is necessary to maintain a stable high vacuum to avoid gas damping [28]. Measurements are made with a vacuum pressure level of at least 1×10^{-7} torr.

Disk Stability on a Sphere

To determine the stability of a disk with diameter D and thickness t balanced on a sphere, we can define θ as the angle between the horizontal plane and the position of the disk, see diagram in Figure 13. The potential energy of the disk is

$$V = mg \left[\frac{D+t}{2} \cos \theta + \frac{D\theta}{2} \sin \theta \right] \quad (3.66)$$

where m is the mass of the disk and g is the gravitational acceleration of the earth.

The equilibrium point occurs when potential is minimized with respect to angular position θ ,

$$\begin{aligned} 0 &= \frac{dV}{d\theta} \\ &= mg \left[-\frac{D+t}{2} \sin \theta + \frac{D}{2} \sin \theta + \frac{D\theta}{2} \cos \theta \right] \\ &= D\theta - t \tan \theta. \end{aligned} \quad (3.67)$$

Applying a Taylor expansion of $\tan \theta$,

$$0 = D\theta - t \left(\frac{\theta^3}{3} + \theta \right), \quad (3.68)$$

which gives the solution

$$\theta = \pm \sqrt{3 \left(\frac{D-t}{t} \right)}. \quad (3.69)$$

Typical samples measured in this experiment have a diameter of 76 mm and thickness of 1 mm. This gives a stability range of

$$\begin{aligned} \theta &= \pm \sqrt{3 \left(\frac{76 \text{ mm} - 1 \text{ mm}}{1 \text{ mm}} \right)} \\ &= \pm 15^\circ. \end{aligned} \quad (3.70)$$

Therefore, this type of suspension is stable over small perturbations about the center, and has good tolerance for sample placement.

Electrostatic Exciter Design

The electrostatic exciter is designed to apply a variable electric force to the sample that will drive its mechanical eigenmodes. The optical samples are assumed to be dielectrics, therefore the exciter will induce a force via coupling an induced dipole with the electric field

generated by the voltage passing through the exciter. The basic design of the exciter is a comb, with alternating ground and high voltage electrodes. The exciter is placed parallel to the sample at a short distance. The distance between the exciter and the sample is set mainly to avoid contact with the optical surface. The force exerted on the sample is optimized by the electrode spacing. To calculate the optimal electrode spacing, we begin with the form of the potential at the surface of the exciter which is

$$V(x, z) = \sum_0^{\infty} A_n \sin k_n x e^{-k_n z} \quad (3.71)$$

where $k_n = \frac{(2n+1)\pi}{s}$, using s to represent the electrode spacing, and z the distance of the exciter to the sample. This form assumes the electrodes extend uniformly in the y direction.

The resulting electric field from this potential is calculated as

$$\vec{E}(x, z) = -\vec{\nabla}V(x, z) \quad (3.72)$$

$$= -\sum_0^{\infty} A_n k_n e^{-k_n z} [\cos k_n x \hat{x} - \sin k_n x \hat{z}] \quad (3.73)$$

To calculate the force exerted on the sample, we use $\vec{F} = p_0(\vec{E} \cdot \vec{\nabla})\vec{E}$. For simplicity, we can separately calculate the forces along \hat{x} and \hat{z} . The resulting average force along the \hat{x} direction can be shown to be

$$\begin{aligned} F_x &= p_0 \left(\sum_{n=0}^{\infty} A_n k_n e^{-k_n z} \left[\cos k_n x \frac{\partial}{\partial x} - \sin k_n x \frac{\partial}{\partial z} \right] \right) \sum_{m=0}^{\infty} A_m k_m e^{-k_m z} \cos k_m x \\ &= p_0 \sum_{n=0}^{\infty} \left[A_n k_n e^{-k_n z} \sum_{m=0}^{\infty} \left(A_m k_m^2 e^{-k_m z} \sin(k_n - k_m)x \right) \right]. \end{aligned} \quad (3.74)$$

Meanwhile, the resulting force along the \hat{z} direction is

$$\begin{aligned} F_z &= -p_0 \left(\sum_{n=0}^{\infty} A_n k_n e^{-k_n z} \left[\cos k_n x \frac{\partial}{\partial x} - \sin k_n x \frac{\partial}{\partial z} \right] \right) \sum_{m=0}^{\infty} A_m k_m e^{-k_m z} \sin k_m x \\ &= -p_0 \sum_{n=0}^{\infty} \left[A_n k_n e^{-k_n z} \sum_{m=0}^{\infty} \left(A_m k_m^2 e^{-k_m z} \cos(k_n - k_m)x \right) \right]. \end{aligned} \quad (3.75)$$

Averaging over x , the force along the \hat{x} direction becomes zero, and the force along the \hat{z} direction becomes

$$\bar{F}_z = -p_0 \sum_{n=0}^{\infty} A_n^2 k_n^3 e^{-2k_n z}, \quad (3.76)$$

which, in terms of separation distance, s , is

$$\bar{F}_z = -p_0 \sum_{n=0}^{\infty} A_n^2 \left(\frac{(2n+1)\pi}{s} \right) e^{-2(2n+1)\pi z/s}. \quad (3.77)$$

For the GeNS experiment, the distance of the exciter to the sample is fixed, and so the electrode spacing must be set to maximize the force on the sample. Using only the $n = 0$ term, and assuming a distance of z_0 between sample and exciter, the optimal electrode spacing s is,

$$\begin{aligned} 0 &= \frac{\partial F_{z_0}}{\partial s} \\ &= -p_0 \frac{\partial}{\partial s} \left[A_0^2 \left(\frac{\pi}{s} \right)^3 e^{-2\pi z_0/s} \right] \\ &= -\frac{3}{s^4} + \frac{2\pi z_0}{s^5} \end{aligned} \quad (3.78)$$

Therefore, the maximum force occurs when $s = \frac{2\pi z_0}{3}$.

For the GeNS, $z_0 = 1$ mm, so the exciter design uses 2 mm electrode spacing. The exciter is a piece of milled PCB, half an inch in width and four inches long to extend over the entire surface of the sample. In the middle, one inch of the copper is milled off completely to ensure no excitation is driven over the mounting point of the sample. One set of electrodes are connected to ground, and the other set are connected to a high voltage amplifier. The high voltage amplifier has a maximum amplitude output of ± 2000 V.

Quadrant Photodiode Board Design

The readout of the optical lever is a quadrant photodiode (QPD), mounted with a transimpedance amplifier circuit. The schematic of one quadrant of the circuit is shown in Figure 16. The original design includes two amplification stages and a differential signal amplifier. It was determined that a differential output was not required for the noise performance of the experiment, so the board was redesigned with a single-ended output for each quadrant. A modeled frequency response of the two amplification stages is shown in Figure 17.




Figure 16: A schematic diagram of the optical lever quadrant photodiode transimpedance amplifier.




Figure 17: Modeled frequency response of the two quadrant photodiode board amplification stages.

3.2.3 Multimodal Measurements

The experimental set up is designed to achieve a multimodal excitation and measurement of the optical sample, that is, multiple mechanical modes are excited and measured at once. The band-limited excitation is set to extend over the expected band of mechanical modes on the sample, parameters which depend on the sample dimensions and material.

Typical samples measured in the Syracuse University GeNS were 3 inches in diameter, and had a thicknesses ranging from 0.3 to 1 mm. This sets the lowest frequency mechanical mode to be between 300 Hz to 1000 Hz. In principle, the eigenmodes of the sample extend indefinitely with increasing frequency, but depending on the quality of the sample and the strength of the excitation, only a few or more modes may be excited. Typical excitations drive anywhere from 5 to 13 mechanical modes at any given time. The band limit generally extends to just below the first expected mode, and up to a few tens of kilohertz.

Figure 18 shows the resulting amplitude spectral density of an uncoated optical sample before and after the sample is excited. The detected mechanical mode peaks are marked. The benefit of a multimodal measurement is that several modes of different geometries are captured in one measurement with decent signal-to-noise ratio to use them for accurate geometrical analyses and loss measurements.

3.3 Extracting Coating Loss From Coated Sample Loss

As previously stated, the substrate loss angle is obtained by simply measuring the loss of the sample before coating. The total loss angle is then measured after the sample is coated. Considering Equation 3.65, further information is required: the dilution factor, or energy ratio, between the coating and the whole sample. There are a few different methods to determine these values.

The samples measured in the GeNS experiment are thin, ranging from 0.3 mm - 1 mm in thickness. The deposited coatings are also thin, ranging from 100 nm - 500 nm. Therefore, thin-plate or thin-film approximations can be made when determining the loss present within the coating or substrate.

The eigenfrequencies of the sample before and after coating will shift, according to the properties of the coating applied to the sample. As found in Li [72], the dilution factor of the coating can be determined via the substrate eigenfrequencies $f_{s,n}$ and coated eigenfrequencies $f_{coated,n}$ via




Figure 18: A spectrum of excited modes on an uncoated 1 mm thick optical sample. The unexcited spectrum is compared with the excited spectrum to identify which peaks are vibrational eigenmodes of the sample. Identified modes are marked.

$$D_{c,n} = 1 - \left(\frac{f_{s,n}}{f_{coated,n}} \right)^2 \frac{m_s}{m_{total}}, \quad (3.79)$$

where m_s and m_{total} represent the mass of the substrate and mass of the coated sample respectively. Using the simple relation $D_{s,n} = 1 - D_{c,n}$, the dilution factor of the substrate is also obtained. Therefore, all the information required to decompose the substrate and coating loss in the gentle nodal suspension experiment is contained in the determination of the sample loss ϕ_n and sample eigenfrequencies f_n before and after coating.

However, use of Equation 3.79 makes several assumptions, namely [68],

1. The substrate and coating are thin films, such that they are planar and have thicknesses much smaller than other sample dimensions
2. The substrate material and thickness are uniform
3. The coating thickness and density are uniform
4. The coating is thin enough such that its presence on the substrate disk does significantly change the shape of the resonant modes.

Also, using this method to determine dilution factors requires knowledge of the thickness and density of the coating.

If any of the above assumptions do not hold, or the other parameters such as thickness or density are unknown, a more complex method for determining the dilution factors is required. A Finite Element Analysis of the substrate and coated sample can provide dilution factors based on known properties of the sample. The software used for this purpose is COMSOL [73].

Similar to the measurement procedure, first the substrate disk must be modeled. The generated parameters of the fit, such as the eigenfrequencies of the sample, can be compared to experimental data. Some iteration of the COMSOL model will be required to match the modeled frequencies with the measured frequencies. Making small adjustments to parameters such as the material density, disk thickness, Young's modulus, or Poisson ratio adjusts the frequencies as necessary. Given that the density of fused silica (the substrate material used) is well known, as well as the sample thickness within some error, the parameters usually adjusted are the Young's modulus and Poisson ratio.

A similar process is followed in modeling the coated sample parameters. Unlike the substrate, thickness and density parameters may not be well known for the coating, so they can be varied along with Young's modulus or Poisson ratio. However, these parameters are somewhat degenerate with each other. Once the resulting frequencies of the modeled sample again match with experimental results, the dilution factors can be calculated from the knowledge of the substrate and coating Young's modulus.

A study performed in [68] compared the validity of dilution factors calculated analytically using Equation 3.79 and those generated by a COMSOL model. Figure 19 demonstrates the calculated dilution factors from both methods. There is mostly good agreement between the methods. However, the COMSOL model is better able to account for the material properties of the system, as a Young's modulus and Poisson ratio must be fit as a part of the model.




Figure 19: A comparison of dilution factors as calculated analytically, using Equation 3.79, and using COMSOL to generate an FEA model, taken from [68]. The error bars in the COMSOL-generated points result from the confidence intervals of the Young's modulus and Poisson ratio.

3.4 A Digital Lock-In Amplifier for Multimodal Analysis

Because the design of the Syracuse GeNS involves multimodal excitation the resulting measurement of the voltage measured in the QPD will be a time series that contains information about every mode excited in the experiment. This data output is large- a 1 mm sample contains eigenmodes from 1 kHz and up, excited up to 30 kHz. To analyze each mode excited in the measurement, it is necessary to identify each mode and track the mode throughout the measurement period as it decays until it is indistinguishable from the background noise of the measurement. The amplitude of each mode will follow a damped sinusoidal trend with time t ,

$$A(t) = A_{0,n}e^{-t/\tau_n} \sin(2\pi f_n t + \varphi). \quad (3.80)$$

By fitting each mode to this trend, the τ_n and f_n can be extracted, and the sample loss ϕ_n calculated.

To properly fit the damped sinusoid ringdown of each mode, a digital lock-in amplifier technique is applied. First, the frequency of the mode must be identified within the excited spectrum, as in Figure 18. Then, the data is bandpassed around this mode to isolate the mode for analysis. A heterodyne scheme is applied to beat the mode frequency with a reference frequency that is slightly offset from the mode frequency, usually around 0.2 Hz. The resulting data is then lowpassed and decimated to produce a data output that contains only the ringdown information of a specific mode. Through this process, the large MultiQ output file can be reduced into smaller files for each mode.

There are two anomalies that must be accounted for in this analysis. First, due to the suspension method of the GeNS, there is a degeneracy between the mode axis along the x- and y-direction. Due to small inhomogeneities in the sample, the resulting modes appear quasi-degenerate, in that they are separated by tens of millihertz, as seen in Figure 20a. When applying the digital lock-in analysis, each mode must be treated separately. Therefore, mode identification must identify each quasi-degenerate mode. Then, the analysis applies a notch filter to one mode during the heterodyne, and the process is repeated similarly for the other mode, Figure 20b.

Second, the mode may shift in frequency slightly during the measurement, due to temperature or other environmental changes. By fitting the φ term in Equation 3.80 with a cubic spline, these shifts in frequency can be applied to the reference frequency, such that





Figure 20: Due to the symmetry of the GeNS experiment, modes along the x- and y-direction will be generated when excited. These modes appear quasi-degenerate, and one mode must be notched at a time to properly apply the digital lock-in analysis. One such quasi-degenerate mode, (0,3), with frequency of 2434.5 Hz is shown here.

the digital heterodyne shifts in frequency with the mode frequency during the ringdown.



(a) Ringdown fit of (0,3) mode

(b) Ringdown fit of (2,2) mode

Figure 21: A fit of the damped sinusoid ringdown of two modes of a 1 mm coated sample. The light blue traces represent the fit, the dark blue traces represent the experimental data, and the red is the residual, or experiment - fit. The left plot shows a fit of mode (0,3) at 2424 Hz with Q of 1.5 million and the right plot a fit of mode (2,2) at 15358 Hz with a Q of 0.7 million.

The resulting ringdown data of each mode is then fit with a damped sinusoid function, as in Equation 3.80. This fit determines the decay time τ_n of the mode, which is then used to calculate the Q of the sample. Figure 21 shows two such damped sinusoid fits, their frequencies, and the resulting Q value.

3.5 Zirconia Titania Measurements

Amorphous oxides, such as the current aLIGO coating materials, have previously been shown to have lower mechanical loss compared to other types of coating materials [74, 75]. Accompanying this material type is the use of ion beam sputtering (IBS) as well as post-depositing annealing to achieve the lowest loss and other favorable optical properties [76, 77, 78]. In particular, tetrahedral metal dioxides like silica achieve the most low-loss bonds, due to the elasticity of the molecular structure [74, 75, 76]. Post-deposition annealing of these coatings at hundreds of degrees Celsius encourage the development of these favorable bonds, further

lowering the mechanical loss, optical absorption, and perhaps even the scatter [61, 76, 62, 79]. Based on this knowledge, it is beneficial to investigate other amorphous materials of this type.

Amorphous zirconia and titania, ZrO_2 and TiO_2 , are two such materials. Loss of these materials has been investigated previously, either independently, or added as a dopant to tantala coatings [80, 81]. When deposited alone, zirconia has been shown to be unstable, and titania has been shown to crystallize quickly when annealed [80]. These are unfavorable properties, but it may be possible to achieve a stable and beneficial mixture of these two materials. A commonly used mixture of zirconia is yttria-stabilized zirconia, or ysz. A ternary mixture of zirconia, yttria, and titania could also have favorable results. Furthermore, these materials are known to have a high index of refraction. This property is particularly beneficial, as the higher difference of indices of refraction between the low and high- n coatings, the less layers required to meet the reflectivity requirements of the aLIGO HR stack.

3.5.1 Measurement Procedure

Before any coating deposition, blank fused silica substrates were measured in the GeNS. This measurement ensured the low loss of the substrates prior to deposition, and provided the required information about the mechanical modes of the substrate. A set of six fused silica samples were annealed at 950°C and then measured in the Syracuse GeNS. Post-measurement, these samples were coated with three different mixtures of zirconia and titania. The coated samples were measured in the GeNS unannealed, and then annealed and remeasured after every annealing step, starting at 300°C and stepping up to 600°C . Measurements after each annealing step were repeated 6-10 times to obtain measurement statistics. One-inch witness samples with the same coating deposition were annealed alongside the 3-inch samples. These witness samples were measured in an XRD to monitor possible crystallization in the samples due to annealing.

3.5.2 Coating Loss of Zirconia-Titania Mixtures

Three mixtures of zirconia-titania coated samples were prepared, in 1:2, 1:1, and 2:1 ratios of Zr:Ti to provide insight into ideal cation ratios, given the previously measured unstable behavior of these oxides alone. Table 2 summarizes the six samples prepared for this experiment by collaborators at the Colorado State University in Fort Collins. The table includes sample thicknesses and indices of refraction, determined through ellipsometry after




Figure 22: Zirconia titania coating loss versus annealing temperature. All samples were measured as deposited, and then annealed at increasingly higher temperatures. Annealing decreased loss in almost every measurement. Coating loss angle calculated using Young's moduli of zirconia and titania assumed from literature. Samples are labeled according to the target composition as listed in Table 2.

Sample Numbers	Target Ratio (Zr:Ti)	Thickness (nm)	n @ 1064 nm
5, 7	1:2	426	2.497
13, 17	1:1	375	2.37
6, 16	2:1	375	2.315

Table 2: Deposited sample parameters measured via ellipsometry

deposition. By measuring the mechanical loss of the samples and tracking the crystallization of the samples after each annealing step, the ideal mixture of zirconia and titania can be determined from the three sets of samples.

Each of these samples were measured and then annealed for several hours at a high temperature. All samples were annealed up to 600°C except for samples 6 and 16, which were found to be damaged after annealing at 550°C. The coating loss was determined using the procedure described in Section 3.3. Figure 22 shows the trend coating loss angle of the lowest frequency mode with annealing temperature. With one exception, coating loss decreases with increasing annealing temperature for all samples. However, the samples with the lowest loss were samples with approximately 33% zirconia as the target composition. The error in each measurement is estimated via a combination of the statistical error from the measurements and the fit error from the damped sinusoid fit.

The dilution factors used to calculate the coating loss in Figure 22 were generated by assuming the Young’s modulus and Poisson ratio based on the target ratios for zirconia and titania as listed in Table 2. Estimates from literature give an approximate zirconia Young’s modulus of 215 GPa [82] and titania Young’s modulus of 147 GPa [83]. Similarly, the density was assumed according to the known values for zirconia and titania density.

By capturing several modes at once, an understanding of the frequency-dependence of zirconia-titania loss can be determined, especially during the annealing process. Figure 23 shows the evolution of the frequency dependence for all measured modes of each sample. The error in each measurement is estimated via a combination of the statistical error from the measurements and the fit error from the damped sinusoid fit. The known titania-doped tantala loss is also shown for comparison. Again, loss is shown to decrease with higher annealing temperature, some samples reaching the titania-doped tantala loss level. Similarly, these loss angles were calculated using the dilution factors with assumed Young’s moduli from literature.




Figure 23: A look at the frequency dependence of the zirconia titania mixture loss. The samples are labeled according to target composition. The loss of titania-doped tantala at 100 Hz is shown as a dotted line in each plot. Two mixtures show little frequency dependence in the loss, while the Zr:Ti 1:2 mixture shows a significant dependence on frequency. All samples show a general trend of improvement with increasing annealing temperature.

3.5.3 Constraining Cation Concentration

To further determine the composition of the samples, collaborators at the University of Montreal performed a Rutherford Backscattering Spectroscopy (RBS) on the zirconia-titania samples. These results are summarized in Table 3, and show a very different amount of zirconia present than originally expected. Furthermore, these samples also showed a significant level of other contaminants, such as platinum and tungsten, further reported in Table 4. These results also indicated abnormally high extinction coefficients, k , possibly as a result of these contaminants.

Target Zr	RBS Zr	RBS Thickness (nm)	n @ 1064 nm	k @ 1064 nm
33%	19%	451	2.50	0.02
50%	21%	377	2.35	0.004
67%	25%	343	2.28	0.004

Table 3: Rutherford Backscattering Spectroscopy results: zirconia concentration, density, index, extinction coefficient

As a result of these measurements, a more accurate value of the thickness and density of the samples were determined, making use of the measured atomic density and thickness from the RBS results.

Target Zr	Zr	Ti	O	Ar	W	Pt	Density (atoms/cm ²)
33%	0.06	0.258	0.66	0.007	-	0.013	3.90×10^{18}
50%	0.068	0.256	0.66	0.018	0.00043	-	3.23×10^{18}
67%	0.082	0.245	0.66	0.018	0.00043	-	2.96×10^{18}

Table 4: Rutherford Backscattering Spectroscopy results: elemental breakdown, atomic density

3.5.4 Determining the Young's Modulus

Using the accurate density values obtained from the RBS results, the zirconia-titania samples were modeled in COMSOL. Using the known density and thickness of each sample, the model was run to constrain the Young's modulus of the mixtures. Given the contaminants in the

samples, assuming literature values of Young's moduli is likely not accurate enough to provide correct dilution factors for these materials.

Instead, Young's moduli were fit using the COMSOL procedure described in Section 3.3. The RBS results listed in Table 4 allow an accurate calculation of the density. For these samples, this value was approximately 4000 g/cm^3 . To fit the data, the Young's modulus Y and Poisson ratio ν ¹ were adjusted. Shown in Figure 24, the slope of the frequency shift from the three mixtures predicts a very high value for the Young's modulus, which is unexpected.




Figure 24: Measured frequency shifts for three samples of differing zirconia concentrations. These values are compared to two COMSOL models of frequency shifts from given Young's moduli and Poisson ratios.

Overall, a Young's modulus of over 200 GPa represents an incredibly stiff material, and one unlikely to yield decent results. Furthermore, despite relatively similar zirconia concentrations, the deviation of slopes of the frequency shifts from sample to sample is large, and requires possibly even higher Young's modulus values to properly fit. These results indicate some problem with the samples that make them impossible to fit, or predict their properties.

¹Equation 3.61 uses σ to represent the Poisson ratio, while the symbol ν is used here.

Given the similarity in cation concentration measured by the RBS analysis, the difference in slope for the three types of samples is unexpected. With similar cation concentrations and uniform sample properties, the Young's moduli and Poisson ratio fit should be similar for each sample. The deviation from sample to sample indicates a larger issue with the samples overall.

The presence of other contaminants such as platinum and tungsten in the concentrations provides some hint that the deposition did not completely account for all parameters to properly coat the samples. These contaminants could be present in coating chamber and inadvertently sputtered onto the samples during the deposition process. Contaminants could help explain why the material properties of the samples are so difficult to fit. One of the key assumptions in the FEA model is that the coating density, thickness, and composition is uniform. Nonuniformities in these parameters due to poor deposition control could explain this strange behavior.

As such, no further work was done to analyze these samples. However, the low loss observed in these samples indicates a possible avenue for further study, if the deposition parameters can be better controlled when new coating samples are produced. The rest of this chapter will describe final procedural steps towards determining the validity of new materials for optical coatings, although these steps were not completed as a part of this work.

3.5.5 Bulk and Shear Loss: fitting frequency dependence

Once the coating loss is properly decomposed from the total sample loss, the frequency dependence of the coating loss must be determined. Two components make up the total coating loss, each possibly having their own frequency dependence: bulk and shear loss. Similarly to the decomposition of substrate and coating loss demonstrated in Equation 3.64, the bulk and shear loss components of the coating loss are related linearly,

$$\phi_c = D_c^{bulk} \phi_c^{bulk} + D_c^{shear} \phi_c^{shear}. \quad (3.81)$$

The benefit of the multimodal measurement is that each measurement easily observes several vibrational modes, each having their own bulk and shear components. The ability to measure ringdowns of several modes from different mode families, as designated by the index m , seen in Figure 15, allows a better understanding of the frequency dependence of the bulk

and shear loss. Modes of the $m = 0$ family, where the nodes of vibration occur only as radial lines, are heavily dominated by shear stress as the greatest amplitude of motion occurs at the edge of the sample, while modes of the $m = 1$ family will have a greater ratio of bulk loss due to high amplitude of motion within the bulk of the sample. In fact, this is again one drawback of the GeNS experiment- the modes of highest bulk energy are modes within the $n = 0$ family, with no radial lines, where the highest amplitude of motion is confined to the center of the sample.




Figure 25: Frequency dependence of the total sample loss is evident across different mode families of one sample with 19% zirconia concentration. Three such mode families are identified here. Note: this plot shows the total sample loss of the substrate and coating combined.

Figure 25 shows the evident frequency dependence for different mode families. Three different mode families are identified, with at least two distinct frequency-dependent trends. A bulk/shear decomposition would allow a better understanding of the distinct bulk and shear frequency dependencies. Based on the behavior of the $m = 0$ family, it is possible that the shear loss angle is low, while the bulk loss angle dominates.

3.5.6 Building a High Reflective Stack

The coatings tested in these experiments were single-layer films. To fully realize a coating of this type in a gravitational wave interferometer, a full stack must be deposited. The zirconia-titania mixtures tested in these experiments show a very high index of refraction, therefore, they must be paired with a low index material to create a high reflective optical coating. The low index material in the current Advanced LIGO detectors, silica, has $n = 1.45$, while the high index of material, titania-doped tantala, has $n = 2.09$. Each layer of the coating stack has thickness d determined by $nd = \frac{\lambda}{4}$, where λ is the laser wavelength and n the index of refraction. The coating layers are deposited on a fused silica mirror substrate.

As demonstrated in [84], the reflectivity of an alternating stack of material with high- and low-indices of refraction n_H and n_L at normal incidence is

$$R = \left[\frac{n_0(n_H)^{2N} - n_s(n_L)^{2N}}{n_0(n_H)^{2N} + n_s(n_L)^{2N}} \right]^2, \quad (3.82)$$

where n_0 represents the index of refraction of the medium outside the reflective stack, which for these purposes is vacuum, where $n_0 = 1$. The total stack will have N layers, and be placed on a substrate of index of refraction n_s .

One possible low loss material for future gravitational wave detector upgrades is titania-doped germania (GeO_2) [85]. While the loss values for germania are favorable, achieving loss comparable with silica, the index of refraction is lower, requiring thicker high index of refraction layers, and more layers overall to achieve the reflectivity requirements. Table 5 lists parameters relevant to creating a full high reflective stack. The index of refraction reported for Ti:Zirconia represents the index of the coating with 19% zirconia, which achieved the lowest loss after 600°C annealing in the GeNS measurements. The layer thickness is the thickness of one individual layer in the HR stack, and the number of high/low layer pairs calculated based on the ETM transmissivity requirement (5 ppm).

While the current results for Ti:Zirconia do not provide evidence that it would be a successful material for an optical coating, further investigation could be fruitful. Preliminary results indicate that the loss angle could be low, possibly comparable with the current Ti:Tantala level. The largest benefit of this material is in the high index of refraction. Even if the loss angle is comparable with the current Advanced LIGO material, the higher index of refraction would allow a thinner HR stack to be deposited. The effective total loss angle of such a stack would therefore be smaller.

Material	Index	Layer Thickness (nm)	Number of Doublet Layers for ETM
Ti:Tantala	2.09	127	19
Ti:Germania	1.88	141	27
Ti:Zirconia	2.5	106	13
Silica	1.45	183	-

Table 5: Coating material, HR layer thickness, total layers

3.6 Conclusions

The Gentle Nodal Suspension experiment is an effective tool for investigating coating loss at room temperature. Investigating future materials for gravitational wave detectors is a lengthy process, often requiring significant iteration of deposition or annealing parameters and coating mixtures. The ease of measurement in a multimodal GeNS, combined with streamlined data reduction and fitting methods using the digital lock-in amplifier, makes measuring sample eigenmodes and fitting ringdown parameters an efficient process. This system, while applied only to the study of one possible future coating material, could be applied to many such samples for testing.

The results reported for zirconia-titania mixtures are inconclusive. Some evidence supports low loss angle results, while further investigation into other material parameters indicate the deposition parameters of these materials must be reinvestigated. As such, the results indicate the importance of quantifying cation concentration with methods such as Rutherford Backscattering Spectroscopy, to further aid in the understanding of the material properties. The zirconia-titania mixtures were largely shown to be stable, even annealing to high temperatures. Further iteration on cation ratios, improved uniformity, and reduction of impurities could provide more favorable results. Beyond the investigation of thermal noise, other optical properties such as scatter and absorption must be studied in this material before it can be deemed fit for the stringent requirements of gravitational wave detector optics. Given the importance of the study of coating thermal noise in room temperature detectors such as Advanced LIGO, all avenues towards possible new low loss materials should be pursued.

Chapter 4

Overview of Alignment Sensing and Control

An active alignment control scheme is key to the operation and high sensitivity of a complex interferometer such as Advanced LIGO. Cavity resonance conditions require stable alignment, and alignment control reduces noise couplings that can limit the gravitational-wave band. This chapter provides a basis for understanding alignment requirements.

This chapter begins with some basic principles of beam propagation through optical components in Section 4.1, as a way to understand the interaction of different optical modes in cavities. Then, it describes the basics of alignment sensors in Section 4.2. Next, Section 4.3 presents an overview of the different alignment degrees of freedom in a aLIGO-like interferometer. With that basis, the current topology for alignment control in Advanced LIGO is outlined. A useful method of measuring alignment controls is described, and several measurements of the current Advanced LIGO alignment controls are presented in Section 4.4. Finally, Section 4.5 describes an argument for an improved alignment configuration for future detector upgrades.

4.1 Basics of Optical Cavities

The propagation of a beam through free space or other optical components can be determined by a ray-transfer matrix, often referred to as an *ABCD* matrix [8].

$$\begin{bmatrix} y_2 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} y_1 \\ \theta_1 \end{bmatrix} \quad (4.83)$$

For example, a ray propagated through free space of distance L would acquire

$$\mathcal{M} = \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix}. \quad (4.84)$$

If the beam propagation occurs in the same medium, such that $n_2 = n_1$, the determinant of such a matrix must be one,

$$AD - BC = 1. \quad (4.85)$$

With this criterion, for a periodic system of optical components, a harmonic (stable) trajectory is possible if

$$\frac{|A + D|}{2} \leq 1. \quad (4.86)$$

For our purposes, we consider the Fabry-Pérot cavity, formed with an input and end mirror some distance L apart. Each mirror has some radius of curvature, R_1 and R_2 . The ray-transfer matrix of reflection from a spherical mirror is [8]

$$\mathcal{M} = \begin{bmatrix} 1 & 0 \\ -\frac{2}{R} & 1 \end{bmatrix}. \quad (4.87)$$

We propagate a ray through our Fabry-Pérot cavity by cascading our transfer matrices for spherical mirrors and free space,

$$\begin{aligned} \mathcal{M}_{FP} &= \begin{bmatrix} 1 & 0 \\ -\frac{2}{R_1} & 1 \end{bmatrix} \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -\frac{2}{R_2} & 1 \end{bmatrix} \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 - \frac{2L}{R_2} & 2L - \frac{2L^2}{R_2} \\ -\frac{2}{R_1} - \frac{2}{R_2} + \frac{4L}{R_1 R_2} & 1 - \frac{4L}{R_1} - \frac{2L}{R_2} + \frac{4L^2}{R_1 R_2} \end{bmatrix}. \end{aligned} \quad (4.88)$$

To calculate the stability condition of a Fabry-Pérot, we apply Eq. 4.86 and find

$$0 < \left(1 - \frac{L}{R_1}\right) \left(1 - \frac{L}{R_2}\right) < 1. \quad (4.89)$$

A key Fabry-Pérot cavity parameter is the *g-factor*, where $g = 1 - \frac{L}{R}$, such that the Fabry-Pérot stability condition can be described as

$$0 < g_1 g_2 < 1. \quad (4.90)$$

The geometry of the Advanced LIGO Fabry-Pérot cavities is chosen to satisfy this condition.

4.1.1 Cavity Eigenmodes

One solution of the paraxial Helmholtz equation is known as the Gaussian beam [8]. The optical intensity of such a beam is a function of the axial and radial distances z and $\rho = \sqrt{x^2 + y^2}$, [8]

$$I(\rho, z) = \frac{2P}{\pi w^2(z)} \exp\left(-\frac{2\rho^2}{w^2(z)}\right). \quad (4.91)$$

P is used to represent the total optical power of the beam, and $w(z)$ the beam width, which is also dependent on axial position. The beam width is defined such that the intensity drops by $1/e^2$ when $\rho = w(z)$.

A Gaussian beam is often described completely via q -parameter, or a complex parameter that contains the beam width wz and radius of curvature $R(z)$

$$\frac{1}{q} = \frac{1}{R(z)} - \frac{i\lambda}{\pi w^2(z)} \quad (4.92)$$

for a laser of given wavelength λ [8].

The full expressions of the radius of curvature and beam width are

$$w(z) = w_0 \left[1 + \left(\frac{z}{z_0} \right)^2 \right]^{1/2} \quad (4.93)$$

$$R(z) = z \left[1 + \left(\frac{z}{z_0} \right)^2 \right] \quad (4.94)$$

where z_0 is known as the Rayleigh range.

As discussed in [86], the Gaussian beam can be expressed in a simple form depending on q

$$\Psi(\rho, q) = A(x, y, q) e^{-ikz} \quad (4.95)$$

$$A(\rho, q) = \frac{A}{q} e^{-ik\rho^2/q^2} \quad (4.96)$$

where A is some complex amplitude. By introducing a term $\psi = \frac{A}{q}$, the Gaussian beam can be described everywhere by (q, ψ) . This is useful for applying ray-transfer matrices, as the beam propagation can be performed via

$$\mathcal{M} \begin{pmatrix} \frac{1}{\psi_i} \\ \psi_i q_i \end{pmatrix} = \begin{pmatrix} \frac{1}{\psi_f} \\ \psi_f q_f \end{pmatrix}. \quad (4.97)$$

This transverse spatial mode is just one possible cavity eigenmode, and it is the spatial beam profile that we would like to achieve in the Advanced LIGO optical cavities. However, it is not the only cavity eigenmode possible in a Fabry-Pérot; in fact the Hermite-Gauss (HG) and Laguerre-Gauss (LG) modes provide complete orthonormal bases for the electric field within an optical cavity [87]. For each basis, the Gaussian beam provides the 00 mode, often referred to as the TEM₀₀ mode.

The *Gouy phase* is an additional phase acquired by a Gaussian beam as it propagates from z_0 to z [8]

$$\zeta(z) = \tan^{-1} \frac{z}{z_0}. \quad (4.98)$$

Any higher order mode, described by indices n and m , accumulates Gouy phase faster than the fundamental mode

$$\zeta(z) = (n + m + 1) \tan^{-1} \frac{z}{z_0}. \quad (4.99)$$

Remembering our Fabry-Pérot cavity geometry, Kogelnik and Li provide a relation between the accumulated Gouy phase in a cavity and the cavity g-factors [87]

$$\zeta(2L) = 2 \arccos(\pm g_1 g_2). \quad (4.100)$$

Therefore, the propagation of spatial modes within a cavity is directly related to the choice of cavity geometry.

The beam propagation method shown in Equation 4.97 becomes very useful when considering the evolution of Gouy phase in a cavity. The change in Gouy phase of a system $\Delta\zeta$ can be related to the terms (q, ψ) via [86]

$$e^{i\Delta\zeta} = \sqrt{\frac{\psi_f \psi_i^*}{\psi_f^* \psi_i}}. \quad (4.101)$$

4.1.2 Higher Order Mode Spacing

As discussed in Section 1.2.1, a Fabry-Pérot cavity achieves resonance when the round trip phase of the light in the cavity is $\phi_{rt} = 2\pi N$. A more complete description accounts for cavity losses or other propagation effects

$$\phi_{rt} + \Delta\phi = 2\pi N. \quad (4.102)$$

The *free spectral range*, FSR, of the cavity is defined as the spacing between cavity resonances occurring at N and $N + 1$, for example. Using the relation $\lambda\nu = c$ and Equation 4.102,

$$\frac{4\pi L\nu_0}{c} + \Delta\phi = 2\pi N \quad (4.103a)$$

$$\frac{4\pi L\nu_1}{c} + \Delta\phi = 2\pi(N + 1) \quad (4.103b)$$

By subtracting these two equations, we find

$$\nu_1 - \nu_0 = \nu_{FSR} = \frac{c}{2L}. \quad (4.104)$$

A TEM _{nm} mode acquires an additional Gouy phase $\zeta(z)$ as it traverses the cavity, which can be expressed as an additional phase shift $\Delta\phi$ in Equation 4.102,

$$2\pi N + (n + m + 1)\zeta(2L) = \frac{4\pi L\nu}{c} \quad (4.105)$$

We would like to know the frequency spacing of higher order mode (n, m) away from the main resonance at ν_{FSR} . We can apply Equation 4.104

$$2\pi N = \frac{2\pi\nu_{nm}}{\nu_{FSR}} - (n + m + 1)\zeta(2L) \quad (4.106)$$

which can be rearrange to find

$$\nu_{nm} = \nu_{FSR} \left[N + (n + m + 1) \frac{\zeta(2L)}{2\pi} \right]. \quad (4.107)$$

Therefore, ν_{nm} is the frequency at which some higher order mode will resonate, offset from the fundamental cavity mode at ν_{FSR} .

A cavity scan can demonstrate the mode spacing through one FSR. Figure 26 shows the results of a scan of the output mode cleaner cavity length in “single-bounce”, such that the beam traverses the input mode cleaner and recycling cavities, but not the arm, demonstrating the mode spacing of the input beam of the Advanced LIGO interferometer [88, 89, 90]. Higher order modes of the carrier (CA) as well as the sidebands (SB) are demonstrated. The scan is fit to begin and end at the FSR, shown with the carrier 00 resonances at 0 Hz and ~ 260 Hz.

Overall, these concepts tell us a few things that are important to understand when considering alignment control:




Figure 26: A single-bounce mode scan of the Advanced LIGO input beam [88, 89, 90]. The scan is taken over one free spectral range of the output mode cleaner cavity. The carrier (CA) and sidebands (SB) are shown here, as well as their first several higher order modes. The higher order mode spacing of the carrier and sidebands are defined via the acquired Gouy phase of each mode.

- There are several possible eigenmodes of the field within our optical cavities, but the desired profile is that of Gaussian beam
- Different eigenmodes propagate with different Gouy phases within the optical cavities
- There is a relationship between the cavity geometry and accumulated Gouy phase of each eigenmode within the cavity.

In this chapter and beyond, these concepts will aid in our understanding of many requirements and challenges of interferometric controls.

4.2 Alignment Sensing Concepts

Before we begin the discussion of the alignment configuration of the Advanced LIGO interferometer, it is important to understand the basic concepts of alignment sensors and alignment degrees of freedom.

4.2.1 Alignment Sensors




Figure 27: The basics of wavefront sensing are demonstrated here. A laser generates a carrier beam, shown in red. The electro-optic modulator (EOM) generates a radio frequency sideband, shown in blue. The carrier beam is resonant in the Fabry-Pérot cavity, while the sideband is not. If a mirror, such as the ETM, is misaligned, the corresponding carrier 01/10 mode (depending on whether it is a pitch or yaw misalignment) beats with the sideband 00 mode. Both of these propagate to two wavefront sensors, WFS A and B, placed 90° in Gouy phase apart. The misalignment signal is then derived via subtraction of the upper and lower or left and right quadrants.

Methods of sensing alignment degrees of freedom can generally be divided into interferometric schemes and pointing schemes. Interferometric schemes rely on interfering a carrier field with radio frequency (RF) sidebands at a photodiode (PD). By placing four PDs in a quadrant, the upper and lower quadrants can be subtracted to detect a pitch misalignment,

and the left and right segments can be subtracted to detect a yaw misalignment. This scheme is known as a wavefront sensor, or WFS, pictured in Figure 27 [91]. The RF sidebands do not resonate within the cavity, so they directly reflect off the input mirror, while the carrier beam resonates within the cavity and detects a misalignment. A cavity misalignment will generate a first-order TEM mode, 01/10, and the resulting carrier misalignment will propagate to the WFS along with the 00 modes of the RF sidebands. Therefore, the WFS generally detect the beat of a first-order carrier misalignment relative to the RF sidebands [91]. Furthermore, due to the Gouy phase shift of the 01/10 mode relative to the 00 mode, two WFS, placed about 90° in Gouy phase apart, are best used to fully resolve the resulting signal. Because WFS operate at RF, a demodulation phase must be specified similar to the longitudinal PDH signal [91].

The RF sidebands used by the WFS are also used as a part of the Pound-Drever-Hall locking scheme for the length degrees of freedom in the interferometer. Because the interferometer is locked in length, the WFS become less sensitive to DC spot drifts on the sensor. To fully ensure the WFS signals remain stable, low bandwidth centering loops are designed to consistently center the beam on each WFS, using the DC power level on each quadrant as the error signal, and small tip/tilt mirrors on each table for feedback. In Advanced LIGO, these are referred to as the “DC centering” loops.

Quadrant photodiodes, or QPDs, are useful for pointing degrees of freedom, where the beam position on a mirror must be sensed. In this case, RF sidebands do not play a part of this sensing. However, pitch and yaw signals are derived in a similar manner via subtracting upper/low or left/right diodes respectively. Other methods for detecting pointing degrees of freedom include dither schemes and camera servos. Dither schemes rely on minimizing an angle-to-length coupling that results from beam miscentering. By injecting a sine wave of known frequency into the angular degree of freedom of the mirror, a length signal is demodulated at that frequency and the resulting servo minimizes the coupling by centering the beam on the mirror. However, injecting a line directly into the differential arm length creates excess noise in the gravitational wave band as significant lines appear directly in the spectrum. Therefore, a new method of beam centering relies on calculating the centroid position of a beam on a camera image, and using that signal to feedback the beam position on the mirror. This method is of similar bandwidth to a dither scheme and does not inject excess noise in the form of lines. At this time, no other excess noise from this control method has been observed.

4.2.2 Alignment Degrees of Freedom




Figure 28: A diagram of the HARD and SOFT mode degrees of freedom in a Fabry-Pérot cavity. In a HARD mode, the mirrors tilt in the same direction, shifting the cavity axis angle, represented with the grey dashed line, relative to the input axis of the beam, represented with the red solid line. In a SOFT mode, the mirrors tilt opposite each other, translating the cavity axis, again grey dashed line, relative to the input beam, solid red line. Both of these degrees of freedom must be controlled in a Fabry-Pérot cavity.

In a complex interferometer like Advanced LIGO, there are several alignment degrees of freedom that must be controlled. Beginning in the Fabry-Pérot arm cavities, the cavity alignment basis contains two degrees of freedom, referred to as the HARD and SOFT mode. The reasoning for this nomenclature will become evident in the discussion in Chapter 6. Figure 28 demonstrates these two cavity modes. The HARD degree of freedom involves a tilt of the ITM and ETM (input and end test mass) such that the the cavity axis becomes tilted relative to the beam input into the cavity. The other cavity degree of freedom is the SOFT mode, where ITM and ETM tilt opposite of each other such that the cavity axis becomes translated relative to the input axis. Alignment control is needed to overlap the two optical axes, such that the power buildup in the TEM_{00} cavity mode is maximized. An interferometer like Advanced LIGO involves two Fabry-Pérot arm cavities, therefore this creates two interferometric degrees of freedom that must be controlled.

Along with these two interferometric degrees of freedom, both the input axis and the cavity axis could shift, in essence shifting the beam position on one of the mirrors. While

this does not effect the power buildup in the cavity, it has other unfavorable effects. Angle-to-length coupling depends on beam mis-centering, and it is best for noise performance to minimize the amount of angular noise coupling in the arm length control. Additionally, beam mis-centering can move the beam spot away from a favorable place on the mirror. In the case of nonuniform mirror qualities, like point defects that absorb excess power or scatter light, the resulting wavefront can become distorted, scattering light into higher order modes. As such, it is important to control “pointing” degrees of freedom. For the arms, this adds two more alignment degrees of freedom that must be controlled.




Figure 29: A diagram of the Michelson alignment degree of freedom. Assuming the input axis and arm cavity axis are overlapped, the Michelson cavity axis must then overlap as well. With the movement of an ITM, the beamsplitter alignment must change to bisect the norm of each ITM.

Moving out from the arm cavities, we add a beamsplitter between the ITMs of the X and Y arms to create the Michelson cavity. This adds another interferometric degree of freedom, as the Michelson cavity axis must overlap the input axis into each arm. This assumes that the arm alignment control has already overlapped the arm cavity axis with the input axis; therefore the beamsplitter alignment must move to bisect the norm of each ITM. Figure 29 demonstrates this degree of freedom, with the red line indicating the input beam axis/Fabry-Pérot cavity axis and the gray dashed line representing the Michelson cavity axis.

Next, we add a recycling mirror, RM, that can represent a mirror of the signal-recycling or power-recycling cavity. For simplicity, we assume the recycling cavity is a linear cavity




Figure 30: A diagram of the recycling cavity alignment degree of freedom. One RM is shown for simplicity, although a full picture of the recycling cavities would demonstrate the three mirror folded cavity. The recycling cavity axis must overlap with the input axis of the arm cavities to allow for maximized signal buildup in the recycling cavities.

(in reality they are folded cavities with three mirrors). With the arm cavity axis overlapped with the input axis, the recycling cavity axis must now overlap the cavity axis in order to maximize power- or signal-recycling cavity buildup. Figure 30 demonstrates this degree of freedom. With both a power-recycling and signal-recycling cavity, this adds two more interferometric degrees of freedom. Similarly to the arms, it is possible for these axes to move together, shifting a beam position on a mirror. Therefore, this also adds two more pointing degrees of freedom.

Finally, the input axis to the interferometer must also be aligned with the power-recycling cavity axis. Although not pictured, this is a similar alignment concept to the RM alignment. This adds one more interferometric degree of freedom.

Importantly, due to the folded cavity layout of the recycling cavities, there is a choice of mirror to use in control of the recycling cavity alignment for both interferometric and pointing DOFs. Therefore, alignment actuator choice comes down to the layout of sensors and signal-to-noise for mirror motion. However, even with these choices, there can be a degeneracy between alignment DOFs that makes the control topology challenging.

In Advanced LIGO, arm alignment control is chosen to follow the HARD/SOFT cavity basis as well as a common and differential basis similar to length control. In essence, the four alignment DOFs of the arms follow $X+Y$ and $X-Y$ for common and differential HARD and SOFT modes. These four degrees of freedom are referred to as common HARD (CHARD), common SOFT (CSOFT), differential HARD (DHARD), and differential SOFT (DSOFT).

Figure 31 demonstrates the mirror motions for each of these four degrees of freedom in yaw. The dotted lines demonstrate the resulting cavity axes due to these motions. Considering, for example, the DHARD alignment case, the cavity axes from each arm shift away

from each other on the beamsplitter. The Michelson alignment must then also shift to account for this change, such that the beamsplitter properly overlaps the Michelson cavity axis with the arm cavity axis. Therefore, the Michelson and DHARD alignment controls can be degenerate. This adds a requirement to develop a sensing scheme that can appropriately distinguish the alignments of the two. Similar degeneracies can exist in almost all the alignment controls, and difficulties in diagonalizing this motion in the sensing scheme can add controls challenges.

For completeness, the pitch degrees of freedom for the arms are also included in Figure 32. In principle, the cavity basis is the same for pitch and yaw, but the definition of the vectors is slightly different, leading to a sign change in the pitch and yaw drive matrices. For the yaw degree of freedom, a counter-clockwise motion represents positive yaw motion, while a tilt downward represents a positive pitch motion.

4.3 Alignment Sensing in Advanced LIGO

In the dual-recycled, Fabry-Pérot Michelson layout of the Advanced LIGO detector, there are a total of ten alignment degrees of freedom to control: six interferometric and four pointing degrees of freedom. These must be controlled in both pitch and yaw, resulting in a total of twenty controlled DOFs. Furthermore, the AC motion of the arm pointing degrees of freedom must be suppressed, and must be sensed separately from the DC alignment for that same DOF, for reasons that will soon be explained. This adds an additional four control DOFs. The resulting alignment control topology quickly becomes complex, and often cross-coupled due to sensing and control limitations.

A measurement of the sensing matrix for each alignment degree of freedom helps determine the best sensors to use in control. One method of determination involves injecting a sine wave of known frequency and magnitude into each degree of freedom. For arm control, this is injected according to the common/differential HARD/SOFT drive matrix. For auxiliary degrees of freedom, this is injected to the mirror that controls that degree of freedom, such as PR2 for the power-recycling cavity alignment. This type of measurement is best performed when the loops controlling each degree of freedom are not closed so that the injection does not encounter any loop effects. However, if the loop is closed, a notch filter located at the injection frequency can be applied in the control loop to remove loop effects.

Local suspension shadow sensors are well-calibrated into microradians, and provide a




Figure 31: A diagram of the arm cavity alignment basis in yaw. Four degrees of freedom are depicted, diagonalized according to the cavity basis (HARD/SOFT) and the common/differential basis. This diagram provides a top-down view, where a positive yaw motion is counter-clockwise in the local coordinate system of the optic. Using this coordinate reference, the resulting drive matrix for each degree of freedom is also included, defined by some factor r_x or r_y that depends on the cavity geometry. The beamsplitter is also shown, so the overall effect on each DOF is evident for the Michelson DOF. Made by D. Brown [92].




Figure 32: A diagram of the arm cavity alignment basis in pitch. Four degrees of freedom are depicted, diagonalized according to the cavity basis (HARD/SOFT) and the common/differential basis. This diagram provides a side view, where a positive pitch motion tilts the mirror downward in the local coordinate system of the optic. Using this coordinate reference, the resulting drive matrix for each degree of freedom is also included, defined by some factor r_x or r_y that depends on the cavity geometry. The beamsplitter is also shown, so the overall effect on each DOF is evident for the Michelson DOF. Made by D. Brown [92].

good witness to the induced mirror motion when an excitation is run. Resulting Watts on the sensors is calibrated through an understanding of the photodiode sensing chain. More information on this calibration can be found in Section 5.1. The sensing matrix results presented here are calibrated using these witnesses and methods in milliWatts/microradians.

In principle, these results form a sensing matrix that can be inverted to create the input matrix for each control degree of freedom. In practice, this matrix inversion often does not result in a stable control scheme, likely to due unaccounted-for cross-coupling. The sensing matrix measurement instead provides a guide for how to best combine the sensors to maximize the desired signal while minimizing other degrees of freedom. The chosen combination of sensors should then be checked for an appropriate zero crossing when the power buildup in the cavity is maximized. When possible, the sensitivity to other degrees of freedom should be minimized. Even with some cross coupling to other DOFs, a reliance on the gain hierarchy of control ensures that dominant signals will be driven to zero quickly compared to other lower-bandwidth degrees of freedom.

4.3.1 Interferometric DOFs

Interferometric degrees of freedom are usually sensed at two ports of the interferometer: the reflection and antisymmetric port. The reflection port (REFL) is created via sensors placed on reflection of the interferometer—behind the AR side of the power-recycling mirror (PRM) (see cartoon layout for reference, Figure 5). At REFL, wavefront sensors best detect common mode misalignments of the interferometer, input beam pointing, and the alignment of the power-recycling cavity. The antisymmetric port (AS), is created via sensors placed on transmission of output steering mirror 3 (OM3) (see cartoon layout for reference 5). At AS, the differential mode, Michelson cavity and signal-recycling cavity alignments are best detected. During observing run four, all interferometric DOFs were sensed at these ports, but in principle, the pick-off port of the power recycling cavity (POP), on transmission of power-recycling mirror 2, can be used for interferometric alignment control. This will be discussed further in Section 4.5.

At the reflection port, two WFS (A and B) sense the beat of the carrier beam with the 9 and 45 MHz sideband signals. The resulting signals are referred to as “REFL_A/B_RF9/45_I”, A/B denoting the two WFS placed at different Gouy phases, 9/45 denoting the RF sideband demodulation, and I denoting the demodulation is performed In-phase. Each quadrant of the WFS is individually phased to maximize common arm length motion in I-phase. No

Q-phase demodulated signal is used from the reflection port. This results in four channels each for pitch and yaw to provide signals to control the common alignment.

Figures 33 and 34 show the results of a sensing matrix measurement taken at the REFL port of the Hanford interferometer. The strongest signal that appears in these sensors is the common cavity axis motion of the arm cavities, referred to as **CHARD**. A linear combination of all four WFS signals are used to control this DOF, diagonalized in such a way to minimize the signal from other DOFs while maximizing the CHARD signal. The CHARD DOF is controlled with a bandwidth between 3-4 Hz.

The alignment of power-recycling cavity relative to the arm cavity axis is sensed next, controlled from power-recycling mirror 2, referred to as **PRC2**. This alignment shows up stronger in the 9 MHz WFS signals. A linear combination of the 9 MHz WFS signals are used to control this DOF, in a way that subtracts the CHARD signal from the combination while maximizing PRC2. This alignment is controlled with a bandwidth close to 0.2 Hz.

Finally, the input beam pointing to the interferometer, controlled from input mirror 4 (IM4), is referred to as **INP1**. This signal shows up weakest, but a linear combination of the 45 MHz WFS signals best diagonalizes this signal. This is the slowest loop, controlled with a bandwidth of 0.04 Hz.

Table 6 summarizes the input matrix values used at the reflection port for the CHARD, PRC2 and INP1 degrees of freedom.

	REFL A 9 I	REFL A 45 I	REFL B 9 I	REFL B 45 I
CHARD P	0.175	1.337	0.232	1.356
PRC2 P	0.056	0	0.034	0
INP1 P	0	1.74	0	-2.93
CHARD Y	-0.72	3.158	-0.72	6.317
PRC2 Y	0.046	0	0.091	0
INP1 Y	0	1.6	0	-1.6

Table 6: The O4a H1 values for the reflection port input matrix. No Q phase signals are used.

At the AS port, the WFS are phased to maximize differential arm length motion in each quadrant in **Quadrature-phase**. While two WFS, A and B, are placed at AS, only WFS A is used for main interferometer control (WFS B is used for the quantum squeezing




Figure 33: A log-scale radar plot of the Hanford detector reflection port sensing matrix in pitch measured from two wavefront sensors (WFS A and WFS B), approximately 90° in Gouy phase apart, demodulated at 9 MHz (RF9) and 45 MHz (RF45). CHARD denotes the common hard mode degree of freedom in the arms, controlled by all four test masses. INP1 denotes the input pointing degree of freedom, controlled by input mirror 4. PRC2 denotes the power-recycling cavity degree of freedom controlled by power-recycling mirror 2. The units are in milli-Watts of photodiode power per micro-radian of pitch mirror motion. Data and procedure found in [93, 94].




Figure 34: A log-scale radar plot of the Hanford detector reflection port sensing matrix in yaw measured from two wavefront sensors (WFS A and WFS B), approximately 90° in Gouy phase apart, demodulated at 9 MHz (RF9) and 45 MHz (RF45). CHARD denotes the common hard mode degree of freedom in the arms, controlled by all four test masses. INP1 denotes the input pointing degree of freedom, controlled by input mirror 4. PRC2 denotes the power-recycling cavity degree of freedom controlled by power-recycling mirror 2. The units are in milli-Watts of photodiode power per micro-radian of yaw mirror motion. Data and procedure found in [93, 94].

subsystem alignment). These WFS are demodulated at 36 MHz, 45 MHz, and 72 MHz. Here, the 36 MHz detects the beat between the 9 and 45 MHz sidebands, and the 72 MHz detects the beat of the 118 MHz and 45 MHz sideband. These signals are referred to as “AS_A_RF36/45/72_Q”. No I-phase signals are used at the AS port.

The differential cavity axis motion of the arm cavities appears strongly in the RF45 signal (beat between carrier and 45 MHz sideband). This degree of freedom is referred to as **DHARD**, and it is controlled with a bandwidth between 4-5 Hz. The Michelson cavity alignment (**MICH**), controlled from the beamsplitter, appears in both the RF36 and RF45 signals. However, since DHARD is controlled on the RF45 demodulation, the RF36 is used for MICH, as the two signals cannot be distinguished on RF45. MICH alignment is controlled with a bandwidth of about 1 Hz. Finally, the alignment of the signal-recycling cavity axis with the differential arm cavity axis is referred to as **SRC1**, and controlled using the signal-recycling mirror, SRM. This DOF is controlled via the RF72 channel with a bandwidth around 0.01 Hz.

4.3.2 Pointing DOFs

Two of the four pointing DOFs of the interferometer control the beam position on the end test masses. In principle, the QPDs on transmission of each arm could be used to sense this beam position. However, these QPDs are sensitive to the drift of the platform they are mounted on, and therefore cannot provide an absolute reference for the beam position. Therefore, a dither scheme or camera servo scheme must be used to control the beam position.

For the dither scheme, the end test masses are dithered with a sine wave near 20 Hz, and the resulting angle-to-length coupling in the differential arm length signal is minimized by feeding back to their respective input and end test masses [44]. A similar process is used to control the beam pointing within the power-recycling cavity, except the input test mass of the X arm is dithered, and control is fed back to the power-recycling mirror. This controls a third pointing degree of freedom: the pointing of the power-recycling cavity. This method relies on injecting sine waves with significant signal-to-noise to demodulate the DARM signal at these frequencies, and due to angle-to-length coupling behavior these lines must be injected within the gravitational wave band [44]. (For a deeper discussion of angle-to-length coupling, see Section 5.5.) While effective at stabilizing the beam position, this purposefully injects alignment noise into the DARM spectrum at low frequency and contributes excess noise.

To avoid noise coupling during science mode, the dither scheme is transitioned to a camera

servo, where the beam position on the end test masses and input test mass is monitored via camera and fed back to the corresponding control point [95]. This method requires no line injection and contributes no measureable excess noise into the differential arm length.

In order to avoid any uncertainties in camera position from lock to lock, the alignment dither system is applied during lock acquisition. Once the beam position is stabilized, the camera servo error signal is compared with the dither signal, and offsets and gains of the servos are adjusted accordingly. Then, the beam control is transitioned from the dither control to the camera servo just before science mode is activated. The bandwidth of both schemes is very low, on the order of milliHertz [44, 95].

Lastly, the pointing of the signal-recycling cavity axis is controlled by an AS QPD on transmission of output steering mirror 1 (OM1) (see cartoon layout for reference, Figure 5). This DOF is referred to as **SRC2** and is controlled via SRM and SR2. The SRC2 bandwidth is approximately 0.02 Hz.

The beam position on both input test masses cannot be controlled at the same time, so a choice of which mirror to control must be made. For observing run four, the Hanford detector chose to control the beam position on the X arm input test mass due to the presence of a point absorber [96]. The beam is purposefully steered to avoid the point absorber and minimize thermal distortions that result from nonuniform absorption.

Table 7 summarizes the input matrix of the AS port, including both interferometric and pointing DOFs.

	AS A 36 Q	AS A 45 Q	AS A 72 Q	AS C
MICH	1	0	0	0
DHARD	0	1	0	0
SRC1	0	0	1	0
SRC2	0	0	0	1

Table 7: The O4a H1 values for the antisymmetric port input matrix. No I phase signals are used. The same input values are used for both pitch and yaw.

4.3.3 Arm Transmission Sensing

While the DC position of the beam position of the arm is controlled via other means, arm transmission QPDs are used for AC control of the beam position. Transmission QPDs, A and

B, are mounted on the transmission monitor platform behind each end test mass, referred to as “TR_X/Y_A/B”. The suspension of the platform is servo-ed to center the beam on QPD B of each arm. Two degrees of freedom are controlled with these sensors: **CSOFT**, the common cavity axis translation of the arm cavities, and **DSOFT**, differential cavity axis translation of the arm cavities. The CSOFT DOF is controlled with a bandwidth of about 0.8 Hz, while DSOFT is controlled at about 0.6 Hz. The control is AC coupled, so as not to interfere with the DC alignment performed by the camera servos. Table 8 shows the LIGO Hanford input matrix for the arm transmission control.

	TR X A	TR X B	TR Y A	TR Y B
CSOFT P	-0.108	0.115	-0.041	0.072
DSOFT P	-0.097	0.103	0.037	-0.064
CSOFT Y	0.532	-0.843	-0.728	1.041
DSOFT Y	-0.091	0.147	0.806	-1.154

Table 8: The O4a H1 values for the arm transmission QPDs input matrix.

The current input matrix for the CSOFT/DSOFT control is designed to diagonalize the hard coupling (arm cavity axis angle) from soft coupling (arm cavity cavity translation). Previously, no centering servos were ran during operation, so the matrix was designed to minimize the sensitivity to the drift of the transmission monitor platform. However, with the consistent servo of the platform to center the beam, the matrix can minimize cross coupling between degrees of freedom that exacerbate alignment control issues at high power [97, 98].

A full depiction of the Advanced LIGO optical layout and the placement of alignment sensors is shown in Figure 35. Not all alignment sensors displayed on the diagram have been discussed in this section. This diagram indicates which alignment sensors are in-air or in-vacuum. During observing run four, all alignment WFS and QPDs used during observation mode were in-vacuum.

4.4 Open Loop Transfer Functions

Measuring open loop transfer functions of the alignment degrees of freedom is an important part of understanding the stability and performance of the alignment controls system. This is often a challenging task—the low bandwidth of these loops requires a long integration time




Figure 35: The Advanced LIGO alignment sensor layout for observing run four. No squeeze alignment sensors are pictured. Adapted from LIGO-G1601619 [99].

to resolve the features of the loop, and the marginal stability of some loops means that large excitations while in-lock can destabilize the entire system. There are a few techniques that one can apply to accurately measure the open loop transfer function, and a way to quantify the error of the measurement.

4.4.1 Measurement Principles

The transfer function of some linear, time invariant system, $H(f)$, is performed by injecting some signal $x(f)$ and measuring $y(f) = H(f)x(f)$ at the output. The estimate of $H(f)$ is calculated via

$$\hat{H}(f) = \frac{S_{xy}(f)}{S_{xx}(f)}, \quad (4.108)$$

where $S_{xy}(f)$ represents the the cross spectral density of x and y and $S_{xx}(f)$ represents the power spectral density of x . Note that $H(f)$ is a complex quantity.

Another quantity related to this measurement is the *coherence*, a measure of the correlation between x and y ,

$$\gamma_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)} \quad (4.109)$$

which ranges from 0 (no correlation) and 1 (perfect correlation). This is often referred to as the “coherence in power”, hence the γ^2 . The coherence is a useful quantity in this type of measurement, as it provides a diagnostic for quality of the measurement.

A transfer function measurement is a useful tool to characterize control systems. A basic control system design will have some plant, P , and control filter, K , such that the *open loop gain* of the system can be defined as $H(f) = PK$. See Figure 36 for a block diagram depiction of such as system. A measurement, as described above, involves injecting some signal into the system, measuring the output, and calculating the resulting frequency response. For such a system, there are three quantities that may be useful in our understanding,

- The *open loop gain*, which we can define as H ,
- the *closed loop gain*, which is defined as

$$\frac{H}{1 - H},$$

- and the *loop suppression function*

$$\frac{1}{1 - H}$$

All of these quantities provide information about the system, but quick examination shows that given one quantity, the other two are easily derivable via algebra. Note: the diagram depicted in Figure 36 shows the output of the control being summed into the feedback system. A different convention involves *subtracting* the control output at this point, in which the resulting signs in the closed loop gain and loop suppression function will flip: $\frac{H}{1+H}$ and $\frac{1}{1+H}$.

There are several ways to make a measurement of a transfer function, all with varying benefits. In the context of alignment control loop measurements, certain methods are more preferable than others. In particular, alignment control loops are low bandwidth, meaning that transfer function measurements will require significant time to properly resolve and average all of the relevant frequencies of the measurement. An additional challenge lies in that these loops must be measured when closed, i.e. when the interferometer is locked. However, the loops are often marginal and cannot handle strong excitations without causing a lockloss. Therefore, these loops must be measured carefully and over long periods of time to resolve the features.




Figure 36: A block diagram of our general control loop. A plant P experiences some external fluctuation x . The control filter K is applied to the error signal output from P , which is then fed back to the input of P to suppress n . To measure the response of such a system, two measurement points b and a are defined, while u represents the excitation injected into the system to resolve the measurement.

One method of measurement involves measuring the “unbiased” open loop gain. Each loop has some external disturbance noise x passing through the loop that the control filter is

designed to suppress. As described in Appendix A in [45], we measure at two points, **before** our excitation and **after** our excitation u , such that we obtain

$$a = \frac{u + x}{1 - H} \quad (4.110)$$

$$b = \frac{Hu + x}{1 - H}. \quad (4.111)$$

The resulting transfer function estimate then becomes

$$\hat{H}(f) = \frac{HS_{uu}(f) + S_{xx}(f)}{S_{xx}(f) + S_{xx}(f)}, \quad (4.112)$$

and $\hat{H}(f) = H(f)$ only if $|u| \gg |x|$ (note that we assume u and x are uncorrelated, so that $S_{ux}(f) = 0$). Otherwise, the result is biased towards one.

However, if we instead measure our transfer function at points a and b with our excitation u , we obtain an unbiased estimation.

$$\begin{aligned} \hat{H}(f) &= \frac{S_{ub}(f)}{S_{uu}(f)} = \frac{HS_{uu}(f)}{(1 - H)S_{uu}(f)} \\ &= \frac{H}{1 - H} \end{aligned} \quad (4.113)$$

and

$$\begin{aligned} \underline{H}(f) &= \frac{S_{ua}(f)}{S_{uu}(f)} = \frac{S_{uu}(f)}{(1 - H)S_{uu}(f)} \\ &= \frac{1}{1 - H}. \end{aligned} \quad (4.114)$$

Notice that these values correspond to the closed loop gain and loop suppression function respectively. Once these are measured, the ratio $H(f) = \hat{H}(f)/\underline{H}(f)$ provides the full open loop gain.

Therefore, the method for measuring the unbiased open loop gain of an alignment control loop involves these steps:

- Shape a broadband excitation similar to the quiescent noise in the loop,
- Inject this shaped excitation for a long time to resolve features with good coherence,

- Measure the closed loop gain and loop suppression function during the injection,
- Calculate the open loop gain via the ratio of the two measurements.

This process still involves long measurement periods to provide good resolution of the low frequency features. The desired frequency resolution will help set the length of the measurement. Assuming that the desired resolution is $\Delta f = 0.01$ Hz, each average will take approximately 60 seconds to measure. With a standard 50% overlap for each average, 100 averages requires a measurement period of nearly one hour. These parameters can vary depending on the measurement target. Overall, this technique involves excitations low enough to maintain lock stability while resolving features over a broad range of frequencies, making it useful despite the time requirement.

Finally, the coherence can also provide an error on the measurement. As derived in Bendat and Piersol [100], the measurement uncertainty σ is related to the coherence in power γ^2 and the number of measurement averages N ,

$$\sigma^2 = \frac{1 - \gamma^2}{2N\gamma^2}. \quad (4.115)$$

Therefore, good measurements require high coherence and many averages, as long as measurement bias is accurately accounted for.

Once an accurate measurement of the loop is made, there are a few characteristics that are used to determine the efficacy of the loops. First, a loop has a *unity gain frequency* wherever $|H(f)| = 1$. To ensure that the loop suppression function is not driven to infinity at this point, the phase must not be zero, $\arg(H(f)) \neq 0^\circ$. When the phase is small near unity gain, the loop suppression becomes greater than one and injects noise in the system. Therefore it is best practice to keep the phase at unity gain within a few tens of degrees at least. The resulting phase at unity gain is referred to as the *phase margin*, and the gain at zero phase is referred to as the *gain margin*. Note: in the previously mentioned case where the sign convention is flipped, when $|H(f)| = 1$, $\arg(H(f)) \neq 180^\circ$. In this case, the phase must be kept a few tens of degrees less than 180° .

4.4.2 Alignment Control Actuation

As listed in Section 4.3, the control of each alignment degree of freedom is performed from a chosen suspended mirror, based on the alignment topology and signal-to-noise of that optic

in an alignment sensor. Most mirrors used in alignment control are triple- or quadruple-suspended mirrors within the auxiliary or arm cavities. This means that the control scheme has an additional complexity to account for the suspension dynamics.

For all multiple-suspensions, actuation is divided between two stages: the top stage for DC actuation, controlled by an integrator and an intermediate or final stage for the AC actuation. The choice of an intermediate or final stage depends the suspension actuation design, the actuator bandwidth, or other noise considerations.

For example, the triple suspension for the beamsplitter has no actuators on the final stage, due to space considerations from the angle of incidence on the mirror. All length and alignment actuation is distributed between the intermediate and top stage of the suspension. For other triple suspensions such at the power- and signal-recycling mirrors, length and alignment actuation is distributed between the bottom stage and the top stage, as the bottom stage has the widest bandwidth for actuation at AC.

For the quadruple suspensions (QUADs), no alignment actuation is performed at the bottom (test mass) stage to avoid actuation noise contamination at the final stage. The penultimate mass stage (PUM) is therefore the best stage to perform AC actuation for alignment, with a low frequency offloading to the top mass for DC control. There is no alignment actuation from the upper intermediate mass.

Independently of the alignment control, all optics have a local control design that damps all six degrees of freedom at the top stage (length, pitch, yaw, transverse, roll, vertical). This damping control should suppress the suspension resonances by some amount before global control (for length and alignment control) is set.

The alignment control actuation is defined via a digital drive matrix, similar to the digital input matrix listed in previous sections. The pitch and yaw drive matrices for the arm cavity control are defined as

$$\begin{bmatrix} \text{CHARD} \\ \text{CSOFT} \\ \text{DHARD} \\ \text{DSOFT} \end{bmatrix}_P = \begin{bmatrix} 0.74 & -0.74 & 1 & 1 \\ 1 & 1 & 0.74 & 0.74 \\ -0.74 & 0.74 & 1 & -1 \\ 1 & -1 & 0.74 & -0.74 \end{bmatrix} \begin{bmatrix} \text{ITMX} \\ \text{ITMY} \\ \text{ETMX} \\ \text{ETMY} \end{bmatrix}_P \quad (4.116)$$

and

$$\begin{bmatrix} \text{CHARD} \\ \text{CSOFT} \\ \text{DHARD} \\ \text{DSOFT} \end{bmatrix}_Y = \begin{bmatrix} 0.72 & -0.72 & 1 & -1 \\ 1 & -1 & -0.72 & 0.72 \\ 0.72 & 0.72 & 1 & 1 \\ 1 & 1 & -0.72 & -0.72 \end{bmatrix} \begin{bmatrix} \text{ITMX} \\ \text{ITMY} \\ \text{ETMX} \\ \text{ETMY} \end{bmatrix}_Y. \quad (4.117)$$

These matrices drive the four test masses according to the cavity basis, which will be further discussed in Chapter 6.

The actuation for the MICH alignment is sent directly to the beamsplitter with a matrix value of one. Similarly so for the PRC2 and INP1 actuation, which are sent to PR2 and IM4 respectively.

The cavity axis alignment of the signal-recycling cavity is controlled from the SRM only, but the pointing axis is distributed between SRM and SR2. The drive matrix of the SRC is then

$$\begin{bmatrix} \text{SRC1} \\ \text{SRC2} \end{bmatrix}_P = \begin{bmatrix} 1 & 0 \\ -7.6 & 1 \end{bmatrix} \begin{bmatrix} \text{SRM} \\ \text{SR2} \end{bmatrix}_P \quad (4.118)$$

and

$$\begin{bmatrix} \text{SRC1} \\ \text{SRC2} \end{bmatrix}_Y = \begin{bmatrix} 1 & 0 \\ 7.1 & 1 \end{bmatrix} \begin{bmatrix} \text{SRM} \\ \text{SR2} \end{bmatrix}_Y \quad (4.119)$$

4.4.3 Open Loop Gain Results

With this background, we are now ready to examine the results of the measured open loop transfer functions. All measurements were taken at the LIGO Hanford Observatory during the early part of the fourth observing run.

Arm Cavity

First, four measurements were made of the arm cavity interferometric alignment loops: DHARD P, DHARD Y, CHARD P and CHARD Y. Figures 37 and 38 show the results of these open loop gain measurements. More detail will be provided in Chapter 6 regarding the unique challenges of the arm cavity alignment control, so description in this section will be brief.

These are the highest bandwidth loops in the alignment control family, with unity gain frequencies (UGFs) between 3-5 Hz. The design of these loops is often a trade-off between required loop suppression and a reduction of noise injection above 10 Hz due to the strong




Figure 37: An open loop gain measurement of CHARD pitch and yaw. The yaw control is set to a slightly higher unity gain frequency to increase low frequency noise suppression.

coupling of each loop to differential arm length. As such, the phase margins at the highest unity gain frequencies range between 22-44 degrees due to the low pass filtering applied to suppress noise. For observing run four, nonlinearities related to the high yaw RMS required the the yaw loop bandwidths be set higher than the pitch loops to suppress the excess noise.

The multiple unity gain crossings in these measurements result from the complicated multi-stage suspension control. The crossover between the top and penultimate stage occurs around a few tenths of a Hz. Other UGFs can be attributed to the plant compensator design which buys phase through multiple gain crossings to ensure the low pass at 10 Hz can suppress as much noise as possible while maintaining loop stability. The gain margin on the high end causes approximately 6 dB gain peaking above the UGF. This is not ideal but kept to a low value to avoid severely impacting the DARM RMS, which causes nonlinear




Figure 38: An open loop gain measurement of DHARD pitch and yaw. The difference in uncertainty can be attributed to the number of averages. Pitch was measured with only 20 averages, while yaw was measured with 60 averages. The yaw control is set to a higher unity gain frequency to increase low frequency noise suppression.

responses in the actuators if too large. There will be more discussion of the arm cavity alignment control design in Chapter 6.

Corner Cavities




Figure 39: Open loop gain measurements of the Michelson cavity alignment loops in pitch and yaw. The differing suspension plants for the beamsplitter pitch and yaw motions drive the difference in unity gain frequency in this loop.

Next, the Michelson cavity alignment loop is shown in Figure 39. The pitch loop reaches its highest UGF at 1 Hz with a phase margin of 46 degrees, and the yaw loop 0.55 Hz with 35 degrees. The differing suspension plant of the beamsplitter in pitch and yaw defines the ability to set the unity gain between the two loops. Due to the cross-coupling with DHARD, the Michelson loop must be lower bandwidth than the arm control, but high bandwidth enough to suppress residual alignment motion of the beamsplitter.




Figure 40: Open loop gain measurements of the power-recycling cavity interferometric alignment (PRC2), input axis alignment (INP1) [left] and signal-recycling cavity interferometric (SRC1) and pointing (SRC2) alignment [right]. The measurements for pitch and yaw are shown. These loops are designed to have very low bandwidth control to manage slow cavity axis and spot position drifts with time. The INP1 bandwidth is extremely low, and the loop performance would likely benefit from a redesign targeting a unity gain frequency similar to INP1 Y.

Both the power-recycling cavity and input alignment control loops are shown on the left in Figure 40. Challenges in sensing diagonalization with CHARD result in some cross-coupling with these degrees of freedom. The loop bandwidths are set low as they are only required to control slow DC drift with time. The added benefit of the low bandwidths mean that the gain hierarchy with CHARD minimizes cross-coupling. PRC2 P has a unity gain frequency of about 0.12 Hz with a phase margin of 46 degrees while PRC2 Y is near 0.17 Hz with a phase margin of 53 degrees. INP1 P has a unity gain frequency of about 0.036 Hz and phase margin of 87 degrees and INP1 Y has a UGF of about 0.25 Hz with phase margin 35 degrees. The INP1 pitch results are strikingly low for the target of this loop, which is closer to 0.1 Hz. This is likely a loop that needs redesign in order to improve some low frequency suppression of the input axis pitch noise.

Finally, the signal-recycling cavity interferometric and pointing loops are shown on the right in Figure 40. Similar to the input alignment loops, these loop bandwidths are low as they maintain alignment against slow cavity drifts. SRC1 P has a UGF around 0.01 Hz, with a phase margin of about 90 degrees. SRC1 Y is similar, 0.01 Hz UGF and phase margin of 100 deg. SRC2 P has a UGF around 0.05 Hz with a phase margin of 42 degrees. SRC2 Y has a UGF around 0.02 Hz and a phase margin of about 46 degrees. The low pointing bandwidth is likely fine, but the SRC1 design in pitch and yaw may benefit from an increase in the bandwidth.

Not all alignment loops are shown here. Namely, the three pointing loops governing the power-recycling cavity pointing, and both arm cavity pointings are unmeasured. These loops are controlled via the camera servos and are of millihertz bandwidth.

Additionally, there are four loops controlling the AC spot motion in the arm cavities: CSOFT and DSOFY pitch and yaw. These loops are the most marginal loops, and therefore the most challenging to measure. All measurement results of these loops are very poor due to low coherence and not presented here.

4.5 The Benefit of POP Wavefront Sensing

Portions of this section are taken from a aLIGO technical report written by the author, see [101] with LIGO technical report number LIGO-T2400054.

This section investigates the benefit of placing POP wavefront sensors in vacuum. Previous aLIGO technical documents have already explored this necessity, such as the technical

document T1600580 [102]. The benefit of a POP WFS has been mainly considered for the ability to control the power recycling cavity alignment, especially in the presence of critical coupling, observed at LHO in O3 and LLO in O4. The in-air sensor POPX has been used previously to perform this control, but has the drawback that it is sensitive to backscatter from ISCT1 [103].

There may be an additional benefit to the new POP WFS, which is that with multiple demodulations (36 and 45 MHz), it could be used for MICH alignment control in addition to PRC alignment control. This could alleviate the AS port WFS that currently support three demodulations.

4.5.1 Current Alignment Control Configuration

Currently, all interferometric alignment signals at the Hanford Observatory are distributed between the in-vacuum WFS at the REFL and AS ports. CHARD (test masses) uses a combination of REFL 9 and 45 MHz signals from both A and B sensors, INP1 (IM4) a combination of REFL 45 MHz signals, and PRC2 (PR2) a combination of REFL 9 MHz signals. PRC2 was previously controlled via the POPX WFS, which currently demodulates at 36 MHz in-air. Part of the O4 commissioning effort involved moving the PRC2 signal from the POPX in air to REFL WFS in vacuum to avoid backscatter [104]. No measurement has been made showing any effect from keeping the POP beam diverter open at LHO, but it has not been tested extensively in the current LHO low-noise state.

At the AS port, all sensing is performed on AS WFS A. MICH (BS) uses the 36 MHz sideband, DHARD (test masses) the 45 MHz, and SRC1 (SRM) uses 72 MHz (beat of the 118 and 45 MHz). Tests have shown that no MICH signal appears in the 72 MHz in full lock [105], although previous work has closed BS control on the 72 MHz WFS in DRMI lock [106].

Previously, the Livingston Observatory operated with the PRC alignment at the REFL port to avoid any noise from scatter from ISCT1 from using the POPX sensor. However, due to loss of alignment signals from the increase in operating power, PRC alignment was moved back to POPX on ISCT1 [107]. Some mitigation has been put in place to reduce noise from scattered light from ISCT1 [108].

LHO sensing matrix measurements of the POPX sensor in air are shown in Figure 41. This measurement shows the response of PRC2 (PR2), CHARD (test masses), and INP1 (IM4) in pitch and yaw, driven at 8.125 Hz. The sensor is currently set to demodulate at 36

MHz. In O3, the POPX sensor at LHO was used to control PRC2, while REFL WFS were used for CHARD and INP1.




Figure 41: Sensing matrix measurements from LHO of the POPX 36 MHz sensor at 60W for pitch (left) and yaw (right), measured at 8.125 Hz. Units are counts/ μ rad. This corresponds to μ rad of IM4 motion for INP1 (single suspension), μ rad of PR2 M3 motion for PRC2, and μ rad of test mass L3 motion for CHARD.

4.5.2 Finesse Model

These simulations make use of an interferometer simulation software, `finesse3` [109], using the LIGO Hanford model from `finesse-ligo`. In this model, the arm and auxiliary cavities of the interferometer are locked in length with an input power of 60 W, but no alignment loops are closed.

The degrees of freedom measured here are as follows:

- CHARD: common hard motion of the arms, diagonalized ideally according to the g-factors of the mirrors in the interferometer cold state
- PRC2: motion of PR2
- PRC3: motion of PR3

- MICH: motion of the beamsplitter
- SRC1: motion of the SRM.




Figure 42: Finesse3 simulation of POP sensing matrix in pitch for 36 and 45 MHz demodulations. Units are in W/rad. Values are at DC with no feedback control engaged.

This particular model adds two new wavefront sensors to the LIGO Finesse3 model, on transmission of PR2 (POP port). Since there is currently no information on the real distances to the proposed in-vacuum POP wavefront sensors, the detectors are placed immediately “behind” PR2, and the Gouy phases are set by hand to be 0° (POP A WFS) and 90° (POP B WFS). There is a 50/50 beamsplitter between these two sensors, and the model measures the results at three radio frequency demodulations: 9, 36, and 45 MHz. The phase is optimized for each detector around the PRC2 degree of freedom.

The pitch response of the listed degrees of freedom are represented for the 36 and 45 MHz demodulations in Figure 42. These results demonstrate that PRC angular motion (PR2 and PR3) as well as CHARD motion show up strongly in the 36 and 45 MHz signal at similar phase and magnitude. This result is similar to the PRC2 and CHARD results from the

POPX measurements taken at LHO. The POP36 wavefront sensor will be a good sensor for controlling the PRC cavity alignment from either the PR2 or PR3 angle. While CHARD shows up with similar strength, the CHARD alignment control is performed at a bandwidth of 3-4 Hz [110] while PRC alignment is controlled at 0.1 Hz [111]. We can rely on the fact that the CHARD signal will be driven to zero in the POP detector during lock.

Also, these results indicate the possibility of controlling the MICH alignment from the POP port. Especially in the 45 MHz signal, the MICH response is phased closer to the Q quadrature, similar to the MICH length signal which is also controlled at POP. The wavefront sensors at the AS port currently support three demodulations: 36, 45, and 72 MHz. If the AS port WFS needed to be reduced to two demodulations (45 MHz to control DHARD and 72 MHz to control SRC1), then the MICH alignment control could be moved to the POP 45 WFS. Furthermore, the SRC1 response at POP is weak, so it is best to maintain SRC alignment control at the AS port. Measurements at LHO have demonstrated that there is no MICH alignment signal present in the AS72 signal [105].




Figure 43: Finesse3 simulation of POP sensing matrix in yaw for 36 and 45 MHz demodulations. Units are in W/rad. Values are at DC with no feedback control engaged.

The yaw response of the listed degrees of freedom are represented for the 36 and 45 MHz demodulations in Figure 43, with similar results to the pitch results. This result is also similar to the PRC2 and CHARD results from the POPAIR measurements taken at LHO.

The 9 MHz results for pitch and yaw are shown in Figure 44, and demonstrate that it is unlikely that decent signals for these alignment degrees of freedom can be derived from a 9 MHz POP sensor. Therefore, POP WFS with 36 and 45 MHz demodulations are likely the best choice.




Figure 44: Finesse3 simulation of POP sensing matrices in pitch and yaw for a 9 MHz demodulation. Units are in W/rad. Values are at DC with no feedback control engaged.

Other alignment degrees of freedom such as the input pointing from IM4 appear weakly in POP (not shown). IM4 alignment is still best controlled from the reflection port.

4.5.3 Benefits of POP Alignment Control

There are multiple benefits to alignment control using a signal from the POP port. One benefit has been previously summarized in T1600580 [102], which is that when facing critical coupling, the alignment signals derived from the REFL port WFS degrade. This was observed in O3 at LHO. A similar issue has been seen at LLO in O4. The CHARD signal

still shows up strongly at REFL, but the PRC2 and INP1 signals have been completely lost. IM4 is now controlled on dither, and PRC2 must be controlled using POPX.

The PRG at LHO is currently very high, and at 60W operating power the POPX sensor is not needed as the PRC2 signal in REFL is stable. However, the O5 upgrade will include replacement of all four test masses at both sites. Currently, we cannot predict what effects the new test masses will have on the PRG. The replacement of ITMY at LHO was beneficial for several reasons. In this particular instance it solved the issue with critical coupling at LHO. However, the replacement of test masses at LLO had the opposite effect of reducing the PRG and likely contributing to the loss of alignment signals at REFL. Hopefully, the O5 test mass replacements will be only helpful, but having in-vacuum POP alignment sensors will allow flexibility to manage any alignment issues related to critical coupling.

There is also evidence from the LHO power up attempts during O4 commissioning that POP alignment control is much more stable during high power operation. The tests of an 80W lock at LHO showed that regular locklosses occurred about 30 minutes into the lock due to the PRC alignment signal changing significantly at the REFL port [112]. This issue was avoided by moving the alignment signal back to POPX [113]. This is something to keep in mind for any further power increases in O5 commissioning.

Overall, the POP port provides a much more stable signal for PRC alignment than the REFL port in various interferometer configurations. If a POP alignment signal is required for high power operation at both interferometers, the best scenario from a noise perspective is to ensure the sensor is placed in vacuum.

One additional benefit that has not yet been explored is the ability to control other degrees of freedom at POP. The **Finesse3** simulations show that a MICH alignment signal can be derived at POP. If MICH is moved from the AS port to the POP port, this can improve other alignment signals. Currently, the signal-to-noise on the AS72 WFS is poor, partially due to the fact the AS WFS currently support three demodulations on a detector designed for two. With MICH alignment control on POP, the AS port WFS can be redesigned to manage only the 45 and 72 MHz signals, which can greatly benefit the interferometric control of the SRC alignment.

Two POP WFS will probably be the most beneficial to achieve these goals. With two detectors approximately 90° in Gouy phase apart, better diagonalization of the alignment signals is possible. In particular, the CHARD signal is phased almost 180° different between the two POP45 sensors. If there is any concern about cross coupling between PRC and

arm alignment, two sensors allow more flexibility to subtract the arm signal from the PRC alignment signal. Conversely, there is very little difference between the two POP36 signals, so there is not much benefit to two different Gouy phases of the POP36 WFS. However, these simulations do indicate the benefit of both the POP36 and POP45 demodulations, so even with one WFS, the upgrade should prioritize a WFS that is capable of the 36 and 45 MHz demods.

4.6 Conclusions

This chapter has provided an overview of alignment sensing and control in a complex interferometer topology. With some background in beam propagation and cavity stability requirements, the basics of wavefront sensing are presented. Several interferometer alignment degrees of freedom are described, along with the best methods for sensing and controlled these DOFs. An overview of the current Advanced LIGO sensing and control topology for alignment is summarized, with descriptions of the sensing, input and drive matrices. A method for measuring open loop gains of alignment control loops is described. Recent open loop gain measurements of almost all alignment degrees of freedom are shown, along with an analysis of their performance. Finally, a proposal for an improved alignment sensing method is described for future detector upgrades, along with both measurements and models to motivate the change.

With a full picture of the alignment topology of the Advanced LIGO interferometer, we can now move into discussion of the noises limiting these controls. Then, we can take a deeper look at the particular challenges of the Fabry-Pérot arm cavity alignment control, especially in the context of high operating power.

Chapter 5

Noise Characterization of Alignment Sensing and Control

As briefly mentioned in Chapter 2, noise from alignment controls is a significant limiting noise source in gravitational-wave detectors at low frequency. Reducing noise from alignment controls is one of the most important avenues for investigation in commissioning advanced detectors, as this is a region that gravitational wave detectors traditionally fall short of design goals. In observing run three (O3), the LIGO Hanford Observatory was nearly completely limited by noise from alignment controls below 30 Hz, see Figure 45.

With the previous chapter's overview of alignment sensing and control, and a better understanding of the control topology and challenges, we can begin an investigation into what alignment controls are the most limiting to gravitational wave strain, what sets the noise in the alignment control loops, and how we can manage the coupling of alignment controls to the differential arm length.

This chapters details efforts to characterize the noises that limit the alignment sensors and controls, and when possible, reduce them. Section 5.1 describes measurements to determine sensing noises that limit the reflection port wavefront sensors. Sections 5.2 and 5.3 detail noise limitations in the wavefront sensors to auxiliary mirror motion and optical table motion. Section 5.4 describes cross-coupling between auxiliary controls and global alignment loops. Section 5.5 explores alignment coupling to differential arm length, and Section 5.6 provides a summary of the current noise performance of each alignment loop given the mitigations discussed in this chapter.




Figure 45: The noise budget of the LIGO Hanford Observatory during O3. Alignment control noise almost completely limits the sensitivity below 30 Hz. Figure from [44].

5.1 Noise Characterization of Alignment Sensors

The alignment sensors, either WFS or QPDs, are a set of four photodiodes that are used for either RF or audio-band detection. As such, they will experience shot noise due to the incident power on the detector and dark noise due to the electronics. Dark noise can be further divided into noise from the electronics, such as Johnson noise, and noise from the analog-to-digital conversion of the signal, called ADC noise.

5.1.1 Shot Noise

Shot noise on a photodiode results from Poisson fluctuations of photon arrival time at the detector. A DC power is detected by measuring a finite number of photons that arrive randomly and independently of each other. This process results in a white noise that depends only on the incident power on the detector, P_{DC} , and the light quanta arriving on the detector, characterized by their wavelength, ω . The amplitude spectral density of shot noise is therefore

$$\sqrt{S_{shot}(f)} = \sqrt{2\hbar\omega P_{DC}} \left[\frac{W}{\sqrt{Hz}} \right]. \quad (5.120)$$

By measuring the incident power on each segment of the alignment sensor, the total shot noise of the detector can be calculated. At the reflection port, the wavefront sensors have approximately 6 mW of power incident in full lock with 50 W of input power [114]. At the antisymmetric port, the wavefront sensors have approximately 4 mW of the power incident in full lock with 60 W of input power [115].

However, this formulation of the shot noise assumes that the incident power on the detector is dominated by carrier power. If the incident power is instead dominated by RF sideband power, then we must account for cyclo-stationary shot noise, where the power incident in each quadrature varies due to the demodulation of the signal [116]. As a result, the shot noise ASD acquires an additional sinusoidal variance that depends on the demodulation phase. As demonstrated in [116], the resulting shot noise PSD becomes

$$\sqrt{S_{shot}(f)} = \sqrt{2\hbar\omega \bar{P} \left(\frac{1}{2} - \frac{1}{4} \cos 2\theta \right)}, \quad (5.121)$$

where \bar{P} represents the average power incident on the photodiode.

At the REFL port, the light incident on the WFS is dominated by carrier power, so Eq. 5.120 holds. However, the power incident at the AS port contains significant 45 MHz sideband power, so Eq. 5.121 must be considered.

5.1.2 Dark Noise

The electronics in each photodiode segment of the WFS have some inherent noise that does not depend on the incident power. This noise is often referred to as “dark noise”, as it is present even with no incident light on the photodiodes. This type of noise mostly arises from the thermal noise in the electronics, and therefore only depends on the temperature and the electronics used. This noise also appears as a white noise.

5.1.3 ADC Noise

Most of the Advanced LIGO controls are performed via the digital system, and so all analog signals must be converted to digital. The conversion of an analog signal to a digital signal acquires additional noise. This noise is also considered a dark noise, since it is present independent of the incident power. However, it is distinct from the noise described above, as it is acquired within the signal chain.

A full diagram of the signal chain is shown in Figure 46. First, an incident power, P , on the photodiode generates a photocurrent I . This photocurrent is converted to a voltage via a transimpedance amplifier. Within these steps, the signal acquires a dark noise, n_{dark} , from the photodiode electronics. This RF signal voltage is then demodulated to create an audioband voltage. The resulting analog signal is then converted from voltage to digital counts using the analog-to-digital converter, where it is whitened before passing through the converter. The digital signal is then dewatered to undo the whitening process. In this way, n_{dark} is both whitened and dewatered, while only n_{ADC} experiences dewatering. This allows the whitening/dewatering process to suppress ADC noise while having no effect on the rest of the signal.

To quantify ADC noise and dark noise, two measurements must be made. First, a measurement of the noise on the WFS with no incident light, and nominal whitening is made (A). Then, the electronics whitening is disengaged, and the noise is remeasured (B). The resulting noises measured in each are as follows

$$\text{noise}_A = \sqrt{[n_{\text{dark}}]^2 + [n_{\text{ADC}} \times DW]^2} \quad (5.122a)$$

$$\text{noise}_B = \sqrt{[n_{\text{dark}}]^2 + [n_{\text{ADC}}]^2} \quad (5.122b)$$

where DW represents the dewhitening filters, usually multiple stages of low-pass filters. These two measurements can be subtracted in such a way to calculate the dark noise and ADC noise present in the dark signal. These equations also demonstrate that these noises add in quadrature.




Figure 46: A diagram of the RF signal chain. First, the incident power on the photodiode is converted to photocurrent, with some photodiode responsivity, R . Then, the signal is converted to a voltage via the transimpedance amplifier with transimpedance, T . The RF signal is demodulated into an audio signal with demodulation Ω and demodulation gain D . Finally, the signal is whitened, passed through analog-to-digital conversion via the ADC and then dewhitened. The resulting digital signal is now measured in counts.

The signal chain depicted in Figure 46 also demonstrates how a measured digital signal is calibrated into Watts of incident power. The incident power, P , is converted to photocurrent, I , with some responsivity, R , that depends on

$$R = \eta \frac{e\lambda}{hc}, \quad (5.123)$$

with e representing the electronic charge, λ the laser wavelength, h Planck's constant, and c the speed of light. Because the conversion from an incident photon to excited electron in the photodiode is not perfect, the quantum efficiency, η , captures the conversion efficiency [117]. The aLIGO WFS have a photodiode responsivity of about 0.76 Amperes of

photocurrent per Watt of incident power [118]. The electronics design includes a transimpedance of 1 kOhm [118]. The electronics chain includes an additional 22 ± 1 dB of demodulation gain [119]. Finally, the signal chain includes a ± 20 V, 16-bit ADC [118]. The AS WFS also have an additional 20 dB whitening gain engaged [120].

The calibrated measurement of these noises in the WFS is shown in Figure 47. The figures show a calibrated measurement of the noises from two REFL port WFS at the 9 and 45 MHz demodulations. Both of these measurements show that the noise above 10 Hz in each sensor is mainly dominated by the dark noise of the sensor itself. Both ports make use of multiple stages of whitening/dewhitening which suppresses the the ADC noise. The shot noise in these plots is calculated according to Equations 5.120.

The noise budgets for the REFL WFS demonstrate that in order to raise the shot noise level to match the level of electronics noise, the shot noise would need to increase by at least a factor of three. This corresponds to approximately a factor of ten increase in power at the REFL port. However, this would likely saturate the electronics, as the design of the aLIGO WFS has a maximum incident power tolerance of about 50 mW due to the 16-bit ADC.

5.2 Contribution from Optic Motion

Reinjected local control noise from auxiliary optics couples into the wavefront sensor signals at the REFL port. These auxiliary optics are not under interferometric control, and only controlled via local damping loops sensed by local shadow sensors. The optics with the strongest contribution to noise in the REFL WFS are the mirrors along the input telescope to the interferometer (IMs 1, 2 and 3), as well as power-recycling mirror 3 (PR3). The wavefront sensors also detect other optic motion, namely the arm cavity mirrors and other power-recycling cavity and input mirrors, however those particular mirrors are under interferometric control, and estimating their contributions to the noise become more difficult. For out-of-loop optics, the contribution to the noise in each sensor can be directly measured via a coherence projection.

The coherence in power between two signals, $x(t)$ and $y(t)$, is defined as

$$\gamma_{xy}^2 = \frac{|S_{xy}(f)|^2}{S_x(f)S_y(f)}. \quad (5.124)$$

If $x(t)$ is the target channel and $y(t)$ is some witness channel of noise in $x(t)$, the coherence between them defines the fraction of noise in $x(t)$ that can be explained by noise witnessed




Figure 47: Noise budgets of ADC, dark, and shot noise for the REFL A WFS demodulated at 9 and 45 MHz. Most of the noise above 10 Hz can be explained by dark noise from the electronics. The measured noise and shot noise were measured in full lock with 50 W input power to the interferometer. This data was taken at the LIGO Hanford Observatory on September 4, 2022.

by $y(t)$. To avoid the bias of random noise, some minimum threshold for coherence must be set. Generally, the inverse of number of averages N of a measurement provides a reasonable minimum for this estimate.

For $\gamma_{xy}^2 > 1/N$, the amplitude spectral density of the noise in $x(t)$ as witnessed by $y(t)$ is

$$\sqrt{S_w(f)} = \sqrt{\gamma_{xy}^2 S_x(f)}. \quad (5.125)$$

This projection assumes that the witness channel is an independent witness of this motion. Since this method usually relies on a passive coherence between the two channels, it can be unreliable in estimating the total noise contribution from auxiliary channels. To better estimate the contribution of witness channel sensors, a direct measurement of the noise via a transfer function is more ideal. When multiple channels witness the same or similar noise in the target channel, a more complete method, such as Wiener filtering, should be applied. However, this method provides a good diagnostic tool for quick estimates.

One such projection is shown in Figure 48. Here, the contributions of PR3 motion, IM motion, and table Z motion are projected into two REFL wavefront sensor channels. These channels are combined to create the error signals for the common arm alignment loops. The figures show significant amounts of the noise in the sensors can be explained by coupled noise from optic motion and table motion. In particular, the reinjected sensor noise from the local sensors of the optics contributes heavily to the noise in the WFS, and the RMS of the signal.

5.2.1 Improving Noise from Local Controls

Examining the results in Figure 48 shows that above 10 Hz, excess noise in the sensors lies an order of magnitude or more above the fundamental noise in the sensors, which lies around $1 \times 10^{-10} \text{ W}/\sqrt{\text{Hz}}$. Much of this noise can be attributed to sensor noise from the local shadow sensors, OSEMS, that are used to control the suspensions [122]. The local controls used for the IMs and PR3 do not sufficiently suppress noise reinjection above 10 Hz.

To mitigate this noise, the local controls were recommissioned for these suspensions. The local control is designed to damp suspension motion, especially around the suspension resonances. As such, these are AC-coupled loops of low bandwidth.

The IM suspensions are single suspensions designed to have a combination of passive and active damping controls for the six degrees of freedom [123]. The active damping controls the length, pitch, and yaw degrees of freedom. The significant noise contribution demonstrated




Figure 48: Projections of optic motion into the REFL wavefront sensors at the LIGO Hanford Observatory. The noise contributions are estimated via coherence projections. Data taken July 6, 2022 [121].

in Figure 48 results from a projection of the damping of the pitch degree of freedom, but all three of these DOFs can contribute noise into the alignment sensors in pitch and yaw (length contribution is a result of length-to-pitch cross coupling).

The active damping loop design in use prioritized suppression of the ~ 1 Hz suspension resonances, but had very little noise suppression above 10 Hz. Figure 49 compares the pitch-to-pitch transfer function of the IM damping control with the old and new design. Above 10 Hz, the loop gain is suppressed by two orders of magnitude with little impact on the loop stability or performance around the suspension resonance of 1.1 Hz. A similar adjustment was made to the length and yaw local controls.




Figure 49: A model of the IM pitch-to-pitch transfer function comparing the loop design before and after improvements were made to reduce noise injection above 10 Hz. The loop maintains suppression around the suspension resonance at 1.1 Hz but reduces the gain by two orders of magnitude at 10 Hz.

The PR3 suspension is a large triple suspension [124] with active damping for all six degrees of freedom, using the BOSEMs [122]. To reduce reinjected BOSEM noise in the global alignment controls, the damping loops for all six degrees of freedom were recommissioned

to prioritize noise reduction above 10 Hz and reduction of the total RMS present on the REFL WFS. A comparison of the new and old design for the pitch-to-pitch transfer function is shown in Figure 50. Gain is reduced by a factor of ten around 10 Hz to improve noise reinjection. There is also some reduction of the gain around the resonances. This design allows some reduction of the noise that dominates the low frequency portion of the WFS signals, as seen in Figure 48, but still maintains desired suppression of the suspension motion. The coherence of the measurement is shown in the third subplot to be near one across the measurement band, demonstrating high accuracy of this measurement.

Both of these changes were beneficial to noise in the wavefront sensors, as well as the noise that couples into the differential arm length through the alignment control. Figure 51 shows a spectrum of wavefront sensors after these upgrades were implemented. No further coherence with the IM damping loops is observed. PR3 motion still dominates the spectrum, but there are some improvements in the noise level across the spectrum. Most importantly, the significant noise resulting from table motion is better resolved compared to Figure 48, allowing other mitigation techniques for the noise above 10 Hz.

5.2.2 Noise Improvements

The improvements in the IM and PR3 damping control also improved alignment noise coupling to the differential arm length. The CHARD signal is formed from a linear combination of all four REFL WFS signals, Table 6. Any improvement in the wavefront sensors is an improvement in the CHARD control noise. Due to the coupling of CHARD into the differential arm length, the noise contaminating gravitational-wave strain is also reduced.

Figure 52 compares the amplitude spectral density of the differential arm length at low frequency around the IM damping change. With the improvement in the IM damping control, the low frequency noise in DARM also improved between 20 and 30 Hz. To confirm this improvement is real and repeatable, a “chop” test is performed. The noise level in DARM is measured while alternating the control configuration from the old to the new configuration. The chop test demonstrated that this improvement was repeatable with the improved IM damping design.

Improvements to the PR3 damping control had a slightly different effect. No change in the noise within the gravitational-wave band was witnessed, but an improvement of the in-loop error signal RMS was measured. Figure 53 shows the DARM error signal, as measured by the OMC DCPDs in current amplitude spectral density. The improvement in the PR3




Figure 50: A measurement of the PR3 pitch-to-pitch transfer function before and after improvements were made to reduce noise injected into the alignment signals at REFL. The new design reduces some control around the suspension resonances, but suppressed noise at 10 Hz. The coherence of the measurement is shown. Data taken at the LIGO Hanford Observatory July 26, 2022.

local control coupled less BOSEM sensor noise to the alignment and length control loops that couple directly to DARM. The improvement resulted in a reduction of the RMS measured on the DCPDs to 0.01 mA. The DARM offset current is 20 mA in observing run four (O4).




Figure 51: Projected noise from both optic motion (PR3) and table motion to the reflection port wavefront sensors after the improvements to auxiliary mirror local controls. Data taken from the LIGO Hanford Observatory on September 4, 2022.




Figure 52: Improvement to the differential arm length after the reduction in IM damping noise injection into the alignment control loops.




Figure 53: A reduction of differential arm length error signal RMS by a factor of two due to the improvements in the PR3 damping loop.

5.3 Contribution from REFL Port Table Motion

The reduction in damping noise in the REFL WFS revealed significant coupling from the motion of the optical table where the sensors are mounted, located within HAM1, see Figure 5. When the reflected interferometer beam arrives to this table, it is steered onto the wavefront sensors with two tip-tilt mirrors, referred to as RM1 and RM2. The table itself is mounted on a passive isolation stack. This passive isolation is insufficient to suppress ground motion above 10 Hz. Any motion that this table experiences couples to the wavefront sensors via the imprint of the excess RM motion onto the light. Four table-top L4C seismometers mounted on this table are good witnesses for this motion.

While nominally six degrees of freedom should be accounted for in table motion (three rotational and three translational), only four sensors are mounted due to space constraints. The signals from the four sensors are diagonalized to account for rotational motion about the X and Y axes (RX and RY), as well X and Z translational motion. Due to poor diagonalization and placement, these four degrees of freedom are highly coupled. To achieve a coherent projection of the table motion to each wavefront sensor, the autocorrelation of the four witness sensors, as well as the cross-correlation of the witness and target sensors must be accounted for. This is achieved by applying a Wiener filter for each target channel, in this case the four REFL WFS channels.

There are two possible strategies to mitigate this excess noise. The first strategy involves implementing signal blending. Both the wavefront sensors at the REFL port and the transmission QPDs at the end stations witness the motion of CHARD, one of the loops most significantly impacted by this noise. As detailed by H. Yu [125] in Section 7.3.2, interferometric control of the arm cavity axis is essential at DC. The AC motion of the arm cavity axis is witnessed by the arm transmission QPDs, and is enhanced by the arm cavity gain, while motion of the recycling cavity is not. Therefore, an AC error signal of the arm cavity axis could be derived from the transmission QPDs and be distinguished from recycling cavity motion, as long as the DC portion of the signal is controlled interferometrically. A signal blend applies a filter that creates a CHARD error signal from wavefront sensors at DC, and QPDs at AC. The QPDs provide control above a few Hz, where the transmission QPD noise is flat.

This is a good strategy, and was applied to the LIGO Hanford Observatory control scheme in O3 [44]. There are two important caveats to this strategy. The first is that this strategy

can only be applied to CHARD, but all loops sensed at the REFL port see this excess noise. The second is that the transmission QPDs contain a significant amount of pitch/yaw cross coupling, demonstrated in Figure 54 [126]. When driving in CHARD pitch at 8.125 Hz, significant portions of signal appear in the four arm transmission QPD yaw signals. This cross coupling is not evident in other sensors. Therefore, application of this method would likely increase CHARD pitch/yaw cross coupling at high frequency, where noise coupling to DARM is significant. Furthermore, this strategy was applied when the table motion coupling was significantly weaker, due to the presence of other noise couplings, as shown by comparing Figure 48 and 51.




Figure 54: A transfer function of the transmission QPD signals to CHARD pitch motion while driving in a sine wave CHARD at 8.125 Hz. No signal blending was applied during this measurement. Notch filters are applied at the same frequency as the drive to avoid loop effects. Significant pitch/yaw cross coupling is measured in the sensors [126].

Given the improved resolution of this table noise in the sensors, the second method to reduce table motion coupling involves applying a feedforward cancellation, similar to the strategy described in 2.2.1. This strategy can be applied to any control loop impacted by this excess noise, and does not introduce any excess opto-mechanical coupling in its application. The application of feedforward relies on the accuracy of the measurement of the coupling—this strategy would not have been as successful before the improvements in the REFL WFS noise from other sources.

5.3.1 Applying HAM1 Feedforward

To apply feedforward of the HAM1 table motion, data was collected from all witness seismometer channels, and the desired control channels (CHARD, PRC2, and INP1). The coupling of these witness channels to the control channels were measured, and a feedforward filter was fit for each channel. The NonSENS (Non-Stationary Estimation of Noise Subtraction) algorithm was applied to fit these channels [127, 128, 129]. NonSENS is an optimal package to use for this noise coupling, as it accounts for noise modulation.

The specific application of this training involves using offline data to train for online noise subtraction. There are specific challenges related to this application [130]:

- The channels are often saved with a different sampling rate than the online channels in LIGO Digital System, e.g. the ASC model is a 2048 Hz model, while all ASC channels are saved offline at 256 Hz.
- The channels used for subtraction are in different models, so IPC delays occur, e.g. the seismometer witness channels are in a 4096 Hz seismic model, versus the 2048 Hz ASC model for the target channels.
- There may be additional filters in the frontend models that are not measured or accounted for in the offline data.
- The target channel is part of an in-loop signal.

A diagram that better elucidates these effects is Figure 55. This figure and work-through is all presented in LIGO technical report T2300064 [130]. Some channel x has a coupling to h defined by A . We would like to design a filter S that will counteract this coupling. Delays between models, or other real-time effects, are accounted for in F . The offline channels x_{DQ} and h_{DQ} are not necessarily the same as the live channels, because the data is often saved at a lower rate than the live channel due to the application of some anti-aliasing filter, Q . When using these offline data sets to train the online subtraction, T is generated, which is not the same as S , the desired coefficient.

We can write the coupling as

$$h = Ax \tag{5.126}$$

where our goal is to tune S to cancel out A

$$0 = (A + SF)x. \tag{5.127}$$




Figure 55: A diagram demonstrating the additional factors considered in offline subtraction using NonSENS. The offline subtraction is designed to implement some subtraction S that counteracts A . Model delay and real-time data effects must be taken into account with F . Q represents anti-aliasing filters used for both the data and strain channels when the data is saved is DQ . Therefore, a proper offline training must account for T , which contains all these extra elements, in order to properly train S . Presented in [130].

The data for offline training is obtained without any subtraction in place, so

$$h = -FSx \quad (5.128)$$

$$\frac{h_{DQ}}{Q_h} = -FS \frac{x_{DQ}}{Q_x} \quad (5.129)$$

meaning that without accounting for these effects, we would train

$$T = -\frac{Q_h}{Q_x} F \quad (5.130)$$

instead of the desired S .

We can refer to $\frac{Q_h}{Q_x} F$ as our “pre-shaping” filter, that we first apply to our offline data x_{DQ} , such that our training actually determines S , instead of T . This pre-shaping filter is applied to the signals before optimization by the NonSENS algorithm. These pre-shaping

filters can either be measured in the online system, or modeled based on the knowledge of the system.

If this subtraction is being applied to an in-loop signal, the effect of the loop also acts as a pre-shaping filter. By measuring the closed-loop effect, this can also be applied. In the case of HAM1 feedforward, the desired subtraction occurs above 10 Hz, while the loop bandwidths are 3 Hz or lower. Figure 56 shows a model of the CHARD pitch closed loop response. This closed-loop effect, above 10 Hz where our subtraction occurs, is negligible. Other control loops that require HAM1 feedforward have even lower bandwidth, and therefore are even less likely to have such an effect.




Figure 56: A model of the magnitude of the CHARD pitch closed loop response.

5.3.2 Noise Improvements

Application of HAM1 feedforward reduced the excess noise in the controls above 10 Hz by a factor of 2 or more. The improvements in the in-loop error signal of each control loop is shown in Figure 57. The target was improvement about 10 Hz, but the implementation reduced some noise below 10 Hz as well. Two of the three control loops, INP1 and CHARD, still show some excess noise that follows the shape of the HAM1 noise. It is possible that the feedforward is not fully optimized. However, PRC2 shows the full expected subtraction of the noise. PRC2 P is sensed only on the RF9 sideband (see Table 6), which INP1 is on the 45 MHz sideband, and CHARD on 9 and 45 MHz signals. It is possible that the coherence of

HAM1 motion is easier to measure on the 9 MHz signals, limiting the achievable subtraction on the 45 MHz signals.




Figure 57: In-loop error signals for three controls loops sensed at the REFL port comparing the implementation of feedforward. HAM1 feedforward subtracts noise due to excess table motion as witnessed by on-table seismometers.

Due to the coupling of the alignment controls to the differential arm length, this benefit extended beyond improvements to the alignment loops. Figure 58 shows a reduction in the differential arm length noise between 10 and 30 Hz with the application of this feedforward. To test and confirm the improvement, the HAM1 feedforward was chopped, while monitoring the differential arm length. This chop test proved that this improvement can be attributed to the reduction of the alignment sensor noise injected into DARM through the alignment controls.

Furthermore, it is imperative that the application of this feedforward be applied to all

REFL port controls, and not just the CHARD loop. While CHARD has the most contribution to the differential arm length noise due to both bandwidth and coupling strength, the lower bandwidth REFL loops have coupling to both DARM and to CHARD. The first implementation of HAM1 feedforward at the LIGO Hanford Observatory only targeted the CHARD loop, and was less effective. By applying this subtraction to all loops, any cross-coupling from INP1 and PRC2 that also contribute to CHARD or DARM noise is accounted for, and the benefit is evident.




Figure 58: The improvement in low frequency noise due to the implementation of HAM1 feedforward for all alignment signals at the REFL port.

5.4 Contribution from Centering Loops

Wavefront sensors require a centered beam to properly detect the first-order misalignment signal when a cavity is not well-aligned with the input axis. To ensure the proper centering of the beam throughout the lock, both the REFL and AS (anti-symmetric) port wavefront sensors have “DC centering” loops—low bandwidth loops that utilize the DC beam position between the four quadrants as an error signal and feed back to steering mirrors to center the beam. For the REFL WFS, DC1 and DC2 center on REFL WFS A and B respectively, and feed back to RM1 and RM2, see Figure 5. For the AS WFS, DC3 and DC4 center on AS WFS A and B, and feed back to OM1 and OM2, see Figure 5. These loop bandwidths do not need to be very high; they are designed to track slow beam spot drifts over the course of a lock.

In the course of investigating the noise coupling from HAM1, the REFL port optical table, it was found that there is high coherence between the CHARD pitch and yaw control loops, and DC1, DC2 and both RM steering mirrors. Each of these components have their own coherence to the HAM1 motion. Passive coherence alone is insufficient to determine if these small suspensions and centering loops couple noise into CHARD, or if all the controls witness the same table motion.

Therefore, the best method is to measure the coupling directly, through a noise injection. By injecting into the DC1 control loop, for example, the resulting transfer function from DC1 to CHARD can be determined. That transfer function can then be used to make a projection of the coupling to understand if noise from these auxiliary loops couple directly to the target (CHARD).

Figure 59 shows the results of such a measurement. A shaped broadband injection into the DC centering loop control (of DC1 and DC2) was driven approximately a factor of ten above ambient at 5 Hz and above. The coherence with the CHARD pitch and CHARD yaw control was measured. The resulting transfer function was used to calculate the coupling to CHARD, following Eq. 2.45.

This result indicates a significant amount of the noise in the CHARD control above 5 Hz was limited by reinjected sensing noise from these centering loops. Notably in this investigation was a lack of coherence with the DC3 and DC4 loops with the DHARD control that uses the AS WFS for its error signal. A significant difference here is the bandwidth of the DC control loop—DC1 and DC2 bandwidths were near 3 Hz, which DC3 and DC4




Figure 59: The measured DC centering loop coupling to CHARD pitch and yaw, calculated via a linear projection based on the measured CHARD/DC coupling. All DC loop points shown have a coherence with CHARD of at least 0.1 for a measurement of 30 averages.

bandwidths were below 1 Hz. Both loops feed back to similar tip-tilt suspensions, so there is no need for this level of control at the REFL port.

The DC1 and DC2 bandwidths were reduced to below 1 Hz, matching the AS loop design. This change had no impact on the locking procedure, which relies on the centering loops to stabilize the WFS signals. Similar to other improvements, this upgrade corresponded to improvement in the differential arm length noise spectrum. Again, this change was chopped multiple times to confirm the the repeatability and origin of the improvement. Figure 60 shows this improvement from 10-20 Hz, matching what we expect based on the transfer function coupling.

5.5 Alignment Coupling from Spot Positions

So far this chapter has focused heavily on the noise present in the loops and reducing it to achieve lower noise overall in the gravitational-wave strain. However, this is not the only method to reduce alignment controls noise. Investigating what actually drives the alignment coupling to differential arm length, regardless of the sensor noise, can also lead to a reduction in the overall noise.

Returning to Equation 2.54, the coupling of angle to length is directly related to the spot position on the mirror. In order to achieve the lowest angle-to-length coupling, the beam must be centered with respect to the actuation center of the mirror. This effect is




Figure 60: The differential arm length noise improvement from a reduction in DC centering control bandwidth.




Figure 61: A diagram of spot-miscentering leading to angle-to-length coupling.

demonstrated in the cartoon in Figure 61. Specifically, the relationship between the terms in this figure can be written as

$$\Delta L \simeq \Delta y \Delta \theta. \quad (5.131)$$

The larger the offset from center, the stronger the coupling of angular motion is to the length signal. Reduction of angular noise coupling can therefore be performed through a reduction of the residual angular motion, $\Delta\theta$, as described so far, as well as a reduction in the beam mis-centering, Δy . Note that this representation of the coupling assumes a singular frequency dependence of the coupling, which is likely inaccurate, as discussed in Barsotti et al [48].

While nominally this requires aligning the beam with the geometrical center of the mirror, there are a few points to consider about the best location for the beam on the mirror surface. First, the attachment point of the fibers on the mirror is above the center point of the mirror by just over one and a half millimeters. Aligning the beam to the geometrical center would then induce inherent pitch-to-length cross coupling. Furthermore, the actuation points on the mirror can be adjusted to set the mechanical center, or actuation node, away from geometric center. This can be used to, for example, compensate the pitch alignment for the slight offset in attachment point. However, the final point of consideration is mirror point defects, which are more important to avoid.

Mirror point defects, often referred to as “point absorbers”, are small defects on or in the optical coating that absorb more than the expected amount of power [131]. These point absorbers may be defects from the coating process, or other contaminants. The excess absorption of these points causes significant wavefront distortions in the cavity given the

large intracavity power in the arms [131, 44]. If point absorbers are present on a Fabry-Pérot cavity mirror, it is best to steer the beam away from the absorber and reduce optical losses from absorption, scatter, and mode matching loss.

To find the best beam position on the mirror, the beams are slowly rastered over the mirror surface by changing the angle-to-length actuation gains. This slowly changes the mechanical actuation point of the mirror, and therefore the mechanical center. Dithering servos run during this process to move the beam to the new center by minimizing angle-to-length coupling, following the coupling in Equation 5.131. During this process, cavity power buildup and sensitivity are monitored to find the favorable operating position, as cavity losses decrease when the beam is moved away from point defects. At the LIGO Hanford Observatory in O4, the beams were moved off-center in pitch and yaw on both ETMs by approximately 16 millimeters to avoid point absorbers. The beam was also mis-centered by a few millimeters on ITMX due to a point absorber. ITMY was replaced shortly before the O4 commissioning period began, and no significant point absorbers have been detected.

Three beam centering servos control the beam position on the test masses: two control the arm cavity beam position for X and Y, and a third controls the beam position on ITMX via actuating on the power-recycling mirror. Therefore, the beam position on ITMY is uncontrolled. If the actuation node of the mirror does not align with the uncontrolled beam position on ITMY, the alignment coupling to DARM increases.

To reduce this coupling, an excitation is driven into an alignment degree of freedom, such as the pitch or yaw of DHARD, CHARD, or ITMY itself. Then, the ITMY actuation gains are adjusted to reduce the measured DOF coupling to DARM by aligning the actuation node with the beam position on the mirror. The inherent assumption in this process is that this coupling is independent of frequency: the actuation gain provides a flat change in the mechanical coupling, and the injection is performed as a sine wave at a single frequency. By injecting a sine wave into the CHARD DOF at 8.125 Hz and adjusting the actuation gains, the pitch coupling was reduced by a factor of four, and the yaw coupling by a factor of 30 [132].

Even with this improvement, significant coupling of the DHARD yaw DOF persisted at LHO. First, a sine wave injections at 30 Hz determined the minimized coupling occurred at an yaw-to-length actuation gain of -1.65, while a sine wave injection at 15.1 Hz determined that an actuation gain of -2.4 best minimized the coupling [133]. To test the origin of this frequency dependence, a broadband excitation of DHARD yaw was injected between 10 and

60 Hz [134].




Figure 62: Broadband injections into the DHARD yaw degree of freedom measure a bilinear coupling to differential arm length. Adjusting the spot position on the mirror away from the nominal position changes the coupling.

Figure 62 shows the results of this test. Gains smaller than -1.65 (closer to mirror center), best minimize the coupling above 30 Hz, while worsening the coupling below 30 Hz. The exact opposite occurs for gains greater than -1.65 . A minimum is found between the actuation gains of -1.55 and -1.4 , where the phase of the transfer function flips sign. Due to the large DHARD yaw noise contribution to DARM below 30 Hz, this test determined that a gain of -1.65 was best to minimize this noise, while not worsening the noise above 30 Hz significantly. Finally, this test validates the presence of bilinear alignment coupling described in Barsotti et al [48].

5.6 Residual Noise in REFL Control Loops

This chapter has summarized in detail effort to understand noise in the REFL wavefront sensors and reduce that noise where possible. Considering the noise budget shown in Figure 47 and the noise projections to the WFS in Figure 51, a few conclusions about the state of the REFL alignment sensing, and the status of the control loops can be made.

First, the wavefront sensors are not shot-noise limited. The measured dark noise present is much higher than that of the shot noise, and it is not possible to increase the incident power to achieve a shot noise limit without saturating the electronics. This dark noise is dominated by inherent dark noise of the PDs or electronics, as measurements show the ADC noise is sufficiently suppressed.

Second, much of the noise, in particular the RMS, is dominated by motion from PR3. It is possible this is mostly reinjected sensor noise from the PR3 BOSEMs. An improvement in the local control of PR3 would likely improve the overall RMS of the REFL port alignment controls. Finally, the noise between 10-30 Hz is still dominated by the REFL table motion. While feedforward can subtract some of this noise, it cannot subtract all of it. Until the table motion is mitigated in some other way, differential arm length will still experience a noise limitation from the excess motion of that table.

With the projections of these noise sources to each wavefront sensor, and the input matrix values for each REFL port loop, summarized in Table 6, the in-loop error signal for CHARD, PRC2 and INP1 can be derived. Figure 33 shows sensing matrix results for each of these wavefront sensors in WFS counts/microradian of mirror motion.

Figure 63 shows this calibrated signal for CHARD pitch and yaw. CHARD is a linear combination of all four WFS signals at the REFL port. The total in-loop RMS of CHARD pitch is approximately 8 nanoradians, and yaw 20 nanoradians. This RMS is dominated by PR3 motion. The sensing noise limit is 7×10^{-14} rad/ $\sqrt{\text{Hz}}$, and coupled REFL port table motion above 10 Hz reaches 2×10^{-12} radians/ $\sqrt{\text{Hz}}$. The yaw signals are not contaminated with this table motion, but only reach a sensing noise value of 3×10^{-13} rad/ $\sqrt{\text{Hz}}$.

The projections for alignment controls noise in Barsotti et al [48] require 1 nanoradian RMS for the Fabry-Pérot arm cavity optics. This measurement indicates that the in-loop Advanced LIGO arm cavity motion due to common cavity axis motion is about factor of ten too high. Furthermore, it predicts that the CHARD sensing noise will reach approximately 2.3×10^{-14} radians/ $\sqrt{\text{Hz}}$ above 10 Hz. Above 30 Hz, the measured sensing noise is a factor




Figure 63: A calibrated measurement of the CHARD pitch and yaw in-loop error signal noise.

of 3 higher in pitch and a factor of 10 in yaw than this prediction, and up to 50 times higher than this prediction between 10 and 30 Hz in pitch.

Figure 64 shows the projected results of INP1, also using calibrations from Table 6 and Figure 33. INP1 pitch achieves 500 nanoradian RMS, and yaw 300 nanoradian, which is about an order of magnitude higher than the predicted 40 nanoradian RMS from [48]. The sensing noise level above 30 Hz is around 8×10^{-12} radians/ $\sqrt{\text{Hz}}$ in pitch and 5×10^{-12} radians/ $\sqrt{\text{Hz}}$ in yaw, about 20 times higher than required in [48]. The excess HAM1 noise increases this value in pitch to up to 2×10^{-10} radians/ $\sqrt{\text{Hz}}$ between 10 and 30 Hz.

Figure 65 shows the projected results of PRC2. In [48], the assumption was made that PR3 would be used for the control of the power-recycling cavity alignment. However, the results for PR2 line up well with this prediction. PRC2 achieves 8 nanoradian RMS in pitch and 1 nanoradian RMS in yaw, well within the requirement for auxiliary ASC [48], and a sensing noise level of 1×10^{-14} radians/ $\sqrt{\text{Hz}}$ in pitch and 7×10^{-15} radians/ $\sqrt{\text{Hz}}$ in yaw. PR3 motion also dominates the RMS, and the noise above 10 Hz is contaminated by REFL table motion up to 2×10^{-13} radians/ $\sqrt{\text{Hz}}$.

These calibrated signals indicate areas for improvement in the low frequency suppression




Figure 64: A calibrated measurement of the INP1 pitch and yaw in-loop error signal noise.

of alignment noise, especially for INP1. However, much of this low frequency noise is a result of BOSEM noise. An improvement in the local sensing noise for suspension control may help relax many of the controls requirements. Also, despite some improvement due to HAM1 feedforward, there is still considerable noise from HAM1 motion coupling into the controls above 10 Hz.

5.7 Conclusions

This chapter covers a variety of noise sources and noise couplings that effect the low frequency noise that limits gravitational wave strain. Due to this study, considerable improvement in the alignment controls was made ahead of the fourth observing run.

In O3, alignment controls noise limited LIGO Hanford Observatory at levels between 5 to 10×10^{-19} m/ $\sqrt{\text{Hz}}$ at 20 Hz, as shown in Figure 45. During first half of the fourth observing run, this alignment controls noise was measured at 1×10^{-19} m/ $\sqrt{\text{Hz}}$ at 20 Hz, and did not directly limit the gravitational wave sensitivity as it did in O3, see Figure 8. Many of the upgrades and studies listed in this section have enabled that reduction. A comparison between the low frequency displacement sensitivity along with the measured alignment noise




Figure 65: A calibrated measurement of the PRC2 pitch and yaw in-loop error signal noise.

projection for O3 and O4 is shown in Figure 66, showing that the work reported in this chapter led to an alignment noise reduction by a factor of ten below 30 Hz at the LIGO Hanford Observatory.

However, this analysis has pointed out several areas for improvement. First, the alignment controls are non shot noise limited, as desired. A deeper investigation into what drives this high level of electronics noise, and how to increase the incident power would be of great benefit to the alignment controls. This would lead to an improvement in signal-to-noise for the controls, and perhaps a reduction overall in the noise above 10 Hz.

Developing better local controls for suspensions, especially in the reduction of sensing noise, is another important part of improving the global alignment controls. Much of the RMS of these controls is driven not by real displacement, but reinjected sensing noise from local shadow sensors. A reduction in RMS would allow the control bandwidths to be reduced, and relax noise suppression requirements.

Despite the success of HAM1 feedforward in O4, much of the noise injected from reflection port loops results from this excess ground motion. An improvement in table isolation, or other mitigation strategies would not just benefit the alignment controls, but the overall noise level in the gravitational-wave strain.




Figure 66: A comparison of the O3 and O4 LIGO Hanford displacement sensitivities, along with the measured alignment controls noise contributions. Between 20 and 30 Hz, the alignment noise contribution has been reduced by a factor of ten. O3 alignment noise reported in [44]. O4 alignment noise measured in [42, 43]

Finally, more in-depth study is required into the behavior of alignment coupling in angle-to-length. A full comparison of the frequency dependence of each degree of freedom may provide more information about the bilinear nature of the coupling, and strategies for mitigation.

While Figure 66 shows exciting results for the improvement in the alignment noise contribution, this level of technical noise is still high, and could again become a limiting noise source as other low frequency noise sources are mitigated. Therefore, ongoing work to study and mitigate alignment controls noise is of the utmost importance to continue to improve detector sensitivity at low frequency.

Chapter 6

Alignment Controls at High Power

In the last two chapters, the alignment control of the arms has been described, and a significant analysis of the noise contributions to the sensors that help control those degrees of freedom has been performed. Now, we will move into a more in-depth look at how the arm alignment control must be designed, and many of the challenges therein, especially in the context of high-power operation.

First, the arm cavity degrees of freedom have been consistently referred to as “HARD” and “SOFT” modes. As promised, the reasoning for this nomenclature will now be described. Also, it has been alluded to that the most significant sources of noise from alignment controls arises from the control of the arms—the reason for this will become apparent in this chapter. To further highlight the importance of understanding arm cavity alignment, not only is the noise from these controls a limiting factor in the noise performance of a gravitational-wave detector, but the actual control performance of the loops limits the operational stability of the interferometer. Both the sensitivity *and* the duty cycle are important figures of merit for an interferometer that hopes to make astrophysical detections. As this chapter will describe, the arm cavity alignment controls are often a key part of that stability factor.

This chapter will include discussions on the opto-mechanical dynamics of a Fabry-Pérot cavity with high intracavity power. There will be a presentation of modeled controls with multiple modeling techniques, and a discussion of high power compensation methods. Then, in addition to outlining the alignment considerations for high power operation, there will be a description of alignment control challenges in the context of thermal effects which often go hand-in-hand with high operating power.

6.1 Radiation Pressure in Fabry-Pérot Cavities

In Chapter 4, the Fabry-Pérot cavity stability conditions were defined in terms of the g -factors, or geometric factors related to mirror curvature and cavity length, $g = 1 - \frac{L}{R}$. For mirrors with curvature R_1 and R_2 , it was stated that for light to stably resonate in the Fabry-Pérot cavity,

$$0 < g_1 g_2 < 1. \quad (6.132)$$

This lends to the result that $\text{sign}(g_1) = \text{sign}(g_2)$, such that their product is positive. Here, we follow the convention that a positive radius of curvature denotes a concave mirror.

The Advanced LIGO mirrors are suspended, and the high finesse of the Fabry-Pérot cavities, $\mathcal{F} = 440$, enables significant amounts of laser power to resonate within the two arm cavities. As a reference, the operating power goal for O4 was 400 kW intracavity power. Previously, the detectors operated around 200 kW intracavity power in O3. The presence of such a beam in the cavity means that the radiation pressure effect, literally the force of light on the suspended mirrors, cannot be ignored. The laser beam itself acts as an optomechanical coupling between the two mirrors in the cavity. When light resonates within a Fabry-Pérot cavity, the mirrors are no longer independent pendula—they operate as one coupled system.

This phenomenon was first outlined in the work by Sidles and Sigg [135], and further studied by Barsotti et al [48] in preparation for Advanced LIGO. The radiation pressure torque on the suspended mirrors as a result of a static misalignment of the input and end test masses ($\theta_{ITM,ETM}$) in a Fabry-Pérot cavity of length L and intracavity power P can be written as

$$\tau = \hat{\mathbf{K}}_{opt}(g_{ITM}, g_{ETM}, L, P) \begin{pmatrix} \theta_{ETM} \\ \theta_{ITM} \end{pmatrix} \quad (6.133)$$

where g_{ETM} and g_{ITM} are the g -factors of the end and input test mass respectively.

As demonstrated in [135] (and in our notation above), the torsional stiffness matrix, $\hat{\mathbf{K}}_{opt}$, only depends on the cavity geometry and intracavity power. This is not so for the radiation pressure force in length, which has a dependence on cavity detuning and other dynamics. One incredible result from this work by Sidles and Sigg is the dependence of radiation pressure torque—it is inherent to the system regardless of cavity design choice, cavity detuning, or other properties. Once a gravitational-wave detector cavity design is chosen, the only other dependency is power. With an ever-continuous desire to operation

gravitational-wave interferometers at high powers, understanding the dynamics of radiation pressure torque is essential.

Specifically, this torsional stiffness matrix has the form [48] (c representing the speed of light)

$$\hat{\mathbf{K}}_{opt} = \frac{2PL}{c(1 - g_{ITM}g_{ETM})} \begin{pmatrix} -g_{ITM} & 1 \\ 1 & -g_{ETM} \end{pmatrix}. \quad (6.134)$$

The relationship shown in this equation and 6.133 is for the case of the single cavity. The diagonalization of this matrix results in one positive and one negative eigenvalue, independent of the choice of cavity geometry or intracavity power. Emphasizing the significance of this result from [135], there will always be one statically stable eigenmode (positive eigenvalue) of the Fabry-Pérot cavity, and one statically unstable mode (negative eigenvalue).

However, this analysis leads to the understanding that there is one element here that can be chosen. The sign of the cavity g-factors will determine which cavity eigenmode is statically stable, and which is statically unstable. For positive g-factors, the major axis eigenmode is statically unstable, and the minor axis eigenmode is statically stable. This was the operating configuration for Initial LIGO, where intracavity power reached 15 kW [135]. However, the Advanced LIGO design was chosen such that the cavity g-factors are negative, which lends itself to a more favorable control design. The eigenvalues of the Advanced LIGO cavities therefore correspond the positive eigenvalue to increasing stability with increasing intracavity power, and the negative eigenvalue to decreasing stability with increasing intracavity power.

This leads to the naming of the two cavity eigenmodes as the “hard” and “soft” modes of the cavity; that is, the hard mode becomes more stiff with increasing power and the soft mode less stiff. A diagram of the two eigenmodes in a Fabry-Pérot are shown in Figure 67. The resulting hard mode corresponds to a tilt of the cavity axis, while the soft mode corresponds to a shift of the cavity axis. This diagram shows the relationship of mirror angle, $\theta_{ITM,ETM}$, with the resulting change in the cavity axis.

When no power resonates in the cavity, the suspension plant of each mirror is free, and has some mechanical frequency f_0 . With the addition of the torque from a resonating beam, the coupled ITM/ETM plant changes; the free pendulum frequency, f_0 , will be shifted from rest via the factor

$$\Delta f_{H,S}^2 = \gamma_{H,S} \frac{PL}{cI(1 - g_{ITM}g_{ETM})} \quad (6.135)$$




Figure 67: A diagram of the hard and soft modes of a single cavity. Diagram taken from Barsotti et al [48].

where I represents the mirror moment of inertia and

$$\gamma_{H,S} = (g_{ETM} + g_{ITM}) \mp \sqrt{(g_{ETM} - g_{ITM})^2 + 4}. \quad (6.136)$$

The hard eigenmode is always positive and the resulting frequency shift only increases this frequency with power. However, the soft eigenmode is negative and can be unstable if the frequency shift, Δf_S , exceeds the fundamental frequency of the pendulum, f_0 .

This is actually an oversimplified view of the system, given that Advanced LIGO test masses use a quadruple pendulum system that has four eigenmodes [37]. In fact, the frequency shift of the eigenmodes of the system becomes more complex.

The Advanced LIGO design includes a maximum operating power of 750 kW. Figure 68 demonstrates the effect of radiation pressure torque on the Advanced LIGO QUAD suspension over the range of operating powers. At zero power, each suspension is free and therefore has four yaw eigenfrequencies corresponding to the four suspension mechanical modes in yaw. Thought of slightly differently, the hard and soft modes of the radiation pressure coupling are degenerate with no cavity power.

With increasing intracavity power, the cavity mirrors become coupled opto-mechanically, breaking the hard/soft mode degeneracy, resulting in eight total mechanical modes. As the intracavity power is increased, the four hard mode frequencies increase with power and the four soft mode frequencies decrease in power. Since the opto-mechanical coupling occurs




Figure 68: Advanced LIGO quadruple suspension (QUAD) yaw eigenfrequencies as shifted with intracavity power according to the Sidles-Sigg effect. At zero power, the Fabry-Pérot cavity mirrors behave as free pendula with four eigenmodes corresponding to the four suspension mechanical modes. With increasing cavity power, the mirrors become one optomechanically coupled system with a splitting of the eigenmodes into eight total “hard” and “soft” modes. The hard modes, shown with solid lines, increase in frequency, while the soft modes, shown with dashed lines, decrease in frequency. For the full Advanced LIGO design power, the soft mode becomes completely unstable as the eigenfrequency becomes negative.

through the coupling of the final mirror stage of each cavity mirror, all mirror modes involving the final suspension stage must split [136].

Represented in Figure 68, the four mechanical modes of the Advanced LIGO QUAD are shown, each represented by a different color. With the presence of resonating cavity power, the degeneracy between the hard and soft modes is broken. The dashed lines represent the soft modes, each decreasing in frequency with increasing power, and the solid lines the hard modes, each increasing in frequency with increasing power.

At the full aLIGO operating power, the soft mode will become completely unstable, as the resulting frequency shift, Δf_S , is greater than the lowest frequency mechanical mode of the system. As shown in Figure 68, the lowest frequency soft mode crosses zero at approximately 700 kW of power. This corresponds to the soft mode becoming completely imaginary and unstable.

Also of note in Figure 68 is the behavior of the intermediate eigenfrequencies. While the highest frequency hard mode and lowest frequency soft mode continue to increase/decrease indefinitely, the intermediate frequencies do not cross, as expected in a stiffness matrix [136]. At some power, each intermediary eigenmode will appear to asymptote to a final value. Therefore, when considering the suspension stiffness behavior with power, it becomes important to understand how the eigenmodes of the stiffness matrix interact. For controls and noise considerations with multi-stage pendula, the most important eigenmodes of the system will be the lowest frequency mode and the highest frequency mode. An unstable soft mode will require a bandwidth of approximately three times the soft mode frequency to control, so the behavior of the lowest frequency mode can inform suspension design to avoid this instability. The highest frequency mode will determine the bandwidth of the hard mode control, as that mode will require damping if it is driven by an external disturbance.

Note: only the yaw behavior of the Advanced LIGO QUAD is shown, but the same effect exists for pitch. However, this becomes more complicated to demonstrate, as there is significant pitch cross-coupling with other degrees of freedom such as length, whereas yaw does not experience any cross-coupling in the design. Future suspension designs, such as the A# BSC Heavy Quadruple Suspension, reduce this coupling and enable a better visualization of the pitch modes. See LIGO Technical Document LIGO-T2300150 [136] for more details.

6.1.1 Beam Size and Radiation Pressure

For a gaussian beam of width w and laser wavelength λ resonant in the Fabry-Pérot cavities, the radiation pressure eigenvalues can be shown, to first order, to depend on the beam spot size in the cavity [135],

$$k_H = \mp \frac{LP w^4}{c w_0^4} \quad (6.137a)$$

$$k_S = \pm \frac{LP}{c} \quad (6.137b)$$

and $w_0 = \sqrt{\lambda L/\pi}$. The Advanced LIGO choice of negative g-factor results in a positive k_H and a negative k_S . The dependency on beam size here further motivates this choice. Advanced LIGO beam size was chosen to be as large as possible while still maintaining the required Fabry-Pérot stability requirements. This choice is to reduce the effect of coating thermal noise, which inversely scales with spot size. The resulting cavity g-factors are set to be

$$g_{ETM} = 1 - 3995 \text{ m}/2245 \text{ m} = -0.78 \quad (6.138a)$$

$$g_{ITM} = 1 - 3995 \text{ m}/1934 \text{ m} = -1.06. \quad (6.138b)$$

These values set the spot sizes to be about 5 cm on the input test masses and 6 cm on the end test masses. The statically stable mode depends strongly on beam size, meaning the stiffer degree of freedom will be more significant with power than the soft degree of freedom.

However, future gravitational wave detectors will exploit new technologies, such as lower thermal noise optical coatings, and longer Fabry-Pérot arms. In these cases, the beam size can be minimized out of preference or necessity. As will be detailed in further sections, reduction of beam spot size will have favorable effects on angular control loop bandwidths, corresponding to less technical noise.

For a 40 km detector, such as the type proposed for Cosmic Explorer [6, 31], beam diffraction must be limited over the distance of the arms. This sets a cavity g-factor requirement of about -0.33, or a 30 km radius of curvature for each test mass. The resulting beam spot size would be approximately 12 cm, twice the size of Advanced LIGO. Therefore, the Cosmic Explorer design will minimize coating thermal noise, which is spot-size dependent, while reducing the effect due to radiation pressure which is cavity-geometry dependent.

In the more near-term upgrade to the Advanced LIGO detectors, A#, there will be no change in the cavity geometry. However, the power requirements will increase, with a goal

of reaching 1.5 MW intracavity power [67]. To maintain cavity stability in these conditions, the test masses will be heavier (100 kg), more than twice the current Advanced LIGO mass (40 kg). This effectively increases the moment of inertia, which is inverse to the effect of radiation pressure torque (Eq. 6.135).

6.2 High Power in Advanced LIGO

In O3, the LIGO Hanford Observatory achieved about 200 kW arm power and LIGO Livingston achieved about 240 kW of arm power in full operation [44]. Arm cavity power is not easy to constrain in Advanced LIGO, as there is no sensor that is a direct witness of the arm power. However, one simple way to determine the arm cavity power is to use the knowledge of input power (P_{in}), power-recycling cavity gain (G_{PR}), and arm cavity gain (G_{arm}) to directly calculate, following

$$P_{arm} = \frac{1}{2} P_{in} G_{PR} G_{arm}. \quad (6.139)$$

This, of course, requires a well-calibrated understanding of the factors G_{PR} and G_{arm} , which we will soon find challenging.

At LIGO Hanford, the arm cavity gain was measured to be 262 and 276 for the X and Y arm respectively, mismatched due to asymmetries in the input test reflectivities and absorptions [44]. The power-recycling cavity gain, called “PRG” for short, was 44 ± 3 W/W [44]. Maximum input power was between 32 and 36 W [44]. In post-O3 commissioning and upgrades, the Y-arm input test mass was removed due to the presence of a point absorber that limited the operating power, and replaced with a test mass that more closely matches ITMX and had no point absorbers. After this replacement, the PRG was measured to be 55 W/W at 2W of input power [137].

During the pre-O4 commissioning period, LHO operated for significant time at three different input powers: 50 W, 60 W and 75 W from the pre-stabilized laser system (PSL). Assuming an arm cavity gain of 260, this corresponds to operation between 300 and 430 kW of arm power. The input power, power incident to the power-recycling mirror, PRG, and arm power are reported in Table 9. While not used in full detector operation mode, the values for 2 and 10 W input power are also included for reference later.

During the locking procedure, it is necessary to stabilize interferometric controls for several different operating powers. All locking sequence steps, from locking the arm cavities

PSL Input Power (W)	P_{PRM} (W)	G_{PR} (W/W)	Estimated P_{arm} (kW)
2	1.8	55	13
10	9.5	54	65
50	48	52	320
60	57	50	370
75	71	46	430

Table 9: Input powers, power-recycling gain, and arm cavity power to LHO during O4 commissioning.

with the ALS system through the CARM offset reduction must be performed at 2 W. This is mainly driven by the difficulty in holding a cavity locked when purposefully detuned, as the arms must be before bringing them to resonance on the main laser. The radiation pressure force in length, similar to what is described in described in Section 6.1, has a strong dependence on longitudinal cavity offsets and power, unlike the angular effects.

Once all cavities are locked and alignment controls are engaged at 2 W power, the locking procedure begins the process of increasing the input power of the interferometer. This is done in slow 5 W steps until the final operating power is reached to allow the slower alignment controls to converge with the changing power. Therefore, the radiation pressure torque dynamics described in Section 6.1 add an additional challenge to interferometer locking: the arm cavity alignment controls must be stable at the initial locking power, the final operating power, and all intermediate powers in between.

6.3 Controls Perspective of Radiation Pressure Torque

The radiation pressure torque effect can be thought of as a feedback loop: some initial torque on the test masses by the incident laser beam causes a small angular misalignment in the mirror. This misalignment changes the cavity axis and spot position, and the change in spot position creates optical torque on the test masses. The radiation pressure torque term, R , depends on the circulating power in the cavity, P_{arm} , and the angle-to-spot-position matrix eigenvalue, $\frac{dy}{d\theta}$

$$R = \frac{2P_{arm}}{c} \frac{dy}{d\theta}, \quad (6.140)$$

where $\frac{dy}{d\theta}$ depends on the cavity geometry. The angle-to-spot-position eigenvalues for the hard and soft modes in Advanced LIGO are

$$\left. \frac{dy}{d\theta} \right|_h = 4.5 \times 10^4 \frac{\text{m}}{\text{rad}} \quad (6.141a)$$

$$\left. \frac{dy}{d\theta} \right|_s = -2.1 \times 10^3 \frac{\text{m}}{\text{rad}}. \quad (6.141b)$$

This is essentially the same information given in Section 6.1, but we now view the system as a feedback loop, which will make our understanding of the controls effects more evident. To visualize this, we can draw a simple block diagram describing this situation, Figure 69. Our control system is two coupled quadruple pendula, which have some free pendulum torque-to-angle transfer function S_0 . The radiation pressure term, R , becomes an additional feedback block in our system.




Figure 69: A diagram of radiation pressure torque as a feedback loop. S_0 represents the pendulum without the presence of any radiation pressure torque. R represents the radiation pressure torque term which depends on cavity geometry and power.

The resulting torque-to-angle suspension transfer function becomes

$$\Theta = \frac{S_0}{1 - RS_0} \tau \quad (6.142)$$

$$\Theta = S_R \tau.$$

This estimation works well for a single suspension, or a multi-stage suspension where the angular actuation is performed at the final stage. For the Advanced LIGO QUAD, alignment actuation is performed from the penultimate mass stage (PUM), and at DC from the top

mass. Since the majority of the alignment control is performed from the penultimate stage, we can rewrite this model slightly to account for that.

$$\Theta_3 = \frac{S_2}{1 - RS_3} \tau_2 \quad (6.143)$$

Here, we have denoted the transfer function of the test mass torque to test mass angle as S_3 and PUM torque to test mass as S_2 . This distinction is important: the radiation pressure term in the feedback loop only applies to the mirror stage, as the laser only directly interacts with that stage.

To see how this manifests in the transfer function behavior, we can use aLIGO QUAD model transfer functions and apply radiation pressure as above in Equation 6.143. Figure 70 shows the plotted HARD and SOFT pitch-to-pitch transfer functions for a series of cavity powers where alignment actuation would be performed from the PUM. Only the pitch-to-pitch model is shown, but this effect is present in both pitch and yaw.

For the HARD transfer function, the poles of the suspension plant shift up in frequency, while for the SOFT they shift down in frequency. The behavior of the intermediate eigenfrequencies shown in Figure 68 are evident in this representation as well—for the HARD mode the low frequency poles appear to stop shifting up in frequency above a certain power, while the SOFT mode higher frequency poles appear to stop shifting down in frequency. The behavior of the low frequency pole of the SOFT mode is also apparent. Above 600 kW of power, it becomes unstable. For the HARD plant, it is stable with increasing power, but the highest frequency pole shifts up in frequency, and therefore any control bandwidth must shift with it. An interferometer that locks at a low power, such as 2 W, and increases the power steadily to the final operating point must account for each of these intermediary suspension plants in the controls design.

6.4 Compensating for Radiation Pressure

Gravitational-wave detectors must compensate for radiation pressure as a part of achieving high power stability. There are two main methods for radiation pressure compensation that have been applied in Advanced LIGO operation. Both methods involve accounting for the changing suspension plant with the digital controls during the power up sequence, although in slightly different ways.

The first method is relatively simple in concept, but more complex in implementation. Multiple plant compensator filters can be designed for different operating powers, such as




Figure 70: Modeled pitch-to-pitch suspension transfer functions for the HARD and SOFT modes at different cavity powers using the Advanced LIGO QUAD model. Modeled following Equation 6.143.




Figure 71: A control diagram with radiation pressure and control block K .

2 W input power, some intermediary input power, and the final operating power. Each of these compensators must be stable enough to work over different input power ranges, such as from 2 W to 25 W, 25 W to 50 W, etc. Therefore, the shifting suspension poles are accounted for, and every step in the power-up process is stabilized. These plants do not need to be optimized for observation mode until the final power is reached, so they may be fairly noisy in terms of reinjection above 10 Hz. Simply put, this control method involves one digital control path, denoted by the K block in Figure 71. The K block is designed to compensate for whatever power is present in R that may change the opto-mechanical plant in addition to meeting alignment controls requirements. It may be adjusted as necessary when the power changes. Therefore, the alignment control transfer function simply becomes

$$\Theta_3 = K \frac{S_2}{1 - RS_3} \tau_2. \quad (6.144)$$

Shifting between multiple suspension plant compensators is tricky. The loop must remain closed above 2 W and the transition between two different compensation filters must be continuous. The LIGO digital control system becomes beneficial in implementation. Two different loop inputs (A and B), can be designed, each with their own set of compensation filters and matching error signals. The output of A will be set on while B is off. The transition occurs when the output of A is slowly ramped to a digital gain of zero while B is slowly ramped to a digital gain of one. The digital system can applying this gain ramp smoothly and continuously without destabilizing the system.

A second method involves Radiation Pressure Compensation (RPC), as described in




Figure 72: A control diagram using Radiation Pressure Compensation block F , in addition to the control K .

dissertation work by H. Yu [125]. In this method, only one K filter design is required. Additionally, another parallel digital control path, denoted by F in Figure 72 is added to the digital control system. Ideally, $F = -R$, that is, F completely compensates for the effect of radiation pressure, and only one K filter needs to be designed that targets only the controls requirements and is designed for one particular opto-mechanical plant.

For example, K can be designed specifically for the 10 W input power radiation plant, for when

$$R^{10W} = \frac{2 \times 65\text{kW}}{c} \frac{dy}{d\theta} \quad (6.145)$$

Then, as power is increased beyond 10 W input, F is adjusted to compensate for the intra-cavity power above 65 kW. For example, at 60 W input power,

$$F^{10 \rightarrow 60W} = R^{60W} - R^{10W}, \quad (6.146)$$

with full loop transfer function

$$\Theta_3^{60W} = K^{10W} \frac{S_2}{1 - R^{60W} S_3 - F^{10 \rightarrow 60W} S_3} \tau_2. \quad (6.147)$$

Note that this requires that the block element F is directly applied to the test mass stage to counteract the radiation pressure which also directly acts at the test mass stage. This

means that F need only be some constant gain factor depending on the difference in power. Furthermore, this has the benefit of needing only one suspension plant compensator design for K , no matter the input power.

However, all alignment actuation must be performed from the PUM stage due to bandwidth and noise considerations, so Equation 6.147 will need to be adjusted, as the control F will be applied to the PUM. F will need to include a factor that accounts for the transfer function from the PUM to the test mass,

$$F = -R \times \frac{S_3}{S_2} \quad (6.148)$$

or in other words, a plant inversion of the S_2/S_3 transfer function.

Furthermore, the alignment controls requirements include stringent noise suppression above 10 Hz once the detector reaches observation mode. This usually involves applying some low pass filter to the control block K . However, the low pass filter compromises the response of the plant inversion and therefore the effectiveness of the RPC if not carefully designed.

$$F = -R \times \frac{S_3}{S_2} \times LP \quad (6.149)$$

The full form of the transfer function becomes

$$\Theta_3^{60W} = K^{10W} \frac{S_2}{1 - R^{60W} S_3 - F^{10 \rightarrow 60W} \times \frac{S_3}{S_2} \times S_2 \times LP} \tau_2. \quad (6.150)$$

This form assumes that K also contains some appropriate low pass filter for observation mode.

Comparing the nominal 10 W controller for a hard loop (K), with the required radiation pressure compensation for control at 60 W (F), as in Figure 73, two concerning aspects are apparent. First, above 10 Hz, the low pass for the radiation pressure control is not well-matched to the low pas required for low noise performance in K . However, the low pass in K comes with significant loss of phase around a few Hz, which will compromise the stability of F if applied on both digital control paths. The more troubling issue is the consideration of what each of these parallel control paths do. The main compensator path, or K , is designed to produce a closed-loop response of a few Hz bandwidth to appropriately control the HARD loop. However, the RPC control path, F , applies essentially a digital soft mode that acts against the stiffening of the suspension dynamics at high power. These two effects occur with a similar response but act oppositely in phase, creating right-half plane poles.




Figure 73: A comparison of the digital controllers for the main loop compensation path K , and the radiation pressure compensation path F , matching the block diagram in Figure 72.

This is more clearly evident in the modeled open loop gain response for the HARD loop, Figure 74. The goal of radiation pressure compensation is match some lower-power plant with a well designed controller. However, measuring the total loop gain, the destabilized zeros around a few Hz are evident, as well as a loss of DC gain.

Adjustments can be made to account for this, such as relaxing the loop compensation in the main path, given that the RPC path requires a strict inversion of the opto-mechanical plant. Or, the compensation can be designed for a higher intermediary power, such at 25 W, where the RPC required to match the 60 W plant to the 25 W plant is not as significant. There is little to be done about the noise impacts, as a needs of RPC are in direct conflict with the effect of a low pass—the phase evolution at a few Hz where radiation pressure has impacted the suspension poles the most must be well matched by the RPC.

LIGO Hanford successfully operated with this radiation pressure compensation method in O3, where the final operating power of 35 W required much less RPC to return the HARD loop control to the 10 W plant. During O4 commissioning, when the input power was doubled to 60 W, the challenges of RPC implementation became evident. To ensure full stability of the alignment controls at the operating power, and to meet noise goals once in observation

mode, LHO instead uses a broadly stable plant compensator during the power up sequence, and then transitions to a tuned controller for the final power, where loop shaping can be better designed to meet the noise goals above 10 Hz.

This method of radiation pressure compensation has so far only been applied to the HARD control. It may be a better strategy to compensate the effect of radiation pressure on the SOFT mode, as the operating power approaches the point of instability for the SOFT modes. Namely, unlike the HARD mode control challenges, RPC for the SOFT mode would require applying a digital HARD mode to stiffen the opto-mechanical spring against the softening effect of the power. That fact, coupled with the relatively low bandwidth of the SOFT loop design would overcome many of the issues with the application of RPC to the HARD loops. At the current powers, the SOFT mode is stable, but this method of compensation can be reconsidered as the detector operating power approaches 700 kW.

6.5 Thermal Effects in Advanced LIGO

One unfortunate consequence of increasing the operating power is the thermal impact the additional operating power has on the mode-matching state of the interferometer. As previously discussed, the cavity geometry and beam parameters are carefully chosen so that the carrier has some Gaussian beam profile that propagates in the cavity, while higher order modes are suppressed. However, the optical coatings and test masses have approximately 0.5 ppm of absorption, meaning that some resonating power will be absorbed by the mirrors. The absorption of the power within the high reflective surface causes a thermo-elastic effect which changes the radius of curvature of the mirror. There is an additional thermo-refractive effect within the bulk of the test mass that the beam experiences as it transits the input test masses of each arm. Overall, this has the effect of changing the overall cavity geometry and causing wavefront distortions as the optimal cavity mode becomes mismatched from the nominal state.

Methods of compensating these thermal effects involve applying other thermal actuation to counteract the test mass heating. The Thermal Compensation System, or TCS, involves actuators like ring heaters around the test masses, and CO₂ lasers that heat the ITM compensation plate. The heating of the test mass causes an increase in the radius of curvature, while ring heaters, for example, provide a negative radius of curvature effect to counterbalance.

Despite the simplicity in description, implementation of these mode-matching corrections often add to the complexity of the system. Thermal actuators are slow, and mode matching must be performed to match the interferometer in its “hot state” when power is resonating within the Fabry-Pérot arm cavities. However, the mode must still be accurate in the “cold state” before power resonates and heats the mirrors. This often means finding a setting that provides decent mode control for the hot state, but does not strongly impact the cold state and hinder the locking procedure. As such, LIGO Hanford operates with imperfect compensation of the thermal effects, which will have an effect on all the controls.

6.6 Modeling Alignment Controls with High Power

Measured arm cavity open loop gains were previously presented in Chapter 4, but now will be studied more in-depth. All measurements presented in this section were taken using unbiased measurement method, which has been outlined in Section 4.4.

Some alignment controls modeling has already been presented in this chapter, as a part of understanding the effects of radiation pressure compensation. More broadly, the benefits of modeling the alignment controls of the arms are increased understanding of the stability of the controls, and how different commissioning changes might effect them. Specifically, the LHO input power was increased from 50 W to 60 W and then to 75 W during the commissioning period. These power increases are challenging due to the effects already outlined in this chapter. Accurate models of the alignment controls enable a better understanding of how to adapt the controls during these changes.

The ultimate goal is to create a reliable model of the arm controls so that appropriate compensation can be developed ahead of time, and instabilities can be understood in the context of the known alignment control status. For example, when increasing the interferometer power from 50 to 60 W, growing 0.5 Hz instabilities were seen in all alignment controls. Several control loops have the possibility of instability at this frequency. However, they cannot all be measured when the interferometer is at risk of lock loss. Modeling helps narrow the source of stability issues.

6.6.1 Modeling the HARD Loops

Equation 6.143 provides a starting point for creating an open loop gain model. We have already established the radiation pressure torque action on the test mass and the digital

control which applied via the PUM stage. Next, the model must account for the split actuation between the PUM and top mass stages, as well as the digital control filters that create that crossover. Finally, a calibration of the suspension response and loop optical gain is required.

The full open loop gain response is described by

$$G_{ol} = G_{opt} \times K \times \left(D_2 \frac{S_2}{1 - RS_3} + D_0 \frac{S_0}{1 - RS_3} \right) \quad (6.151)$$

where the subscripts represent the suspension stage: 0 representing top mass, 2 representing PUM, and 3 representing test mass. K represents the control loop filter, similar to what is shown in Figure 71, and D_n represents the digital locking filters that set the suspension actuation crossover. Finally, G_{opt} represents the calibration factors, which are derived from a combination of the measured optical gain, such as from measurements described in Section 4.3, and the known suspension calibration factors, which are catalogued in [138].




Figure 75: DHARD P model varying the ITM radius of curvature within 1% and the ETM radius of curvature 0.5% for a given power of 360 kW. The nominal, or cold state value, is demonstrated with the black dotted line. Modeled using **Finesse3**.

The free suspension transfer functions are well-modeled, and the models can be further confirmed with offline suspension measurements. The main free parameter in the modeling process is the radiation pressure term, R , which has a dependency on power and cavity geometry. In principle, all that should vary is the power, except that at high power thermal effects impact the cavity geometry.




Figure 76: Model and Measurement of the CHARD P open loop gain. CHARD P measurement taken using the unbiased open loop gain method detailed in Section 4.4. Modeled analytically using the methods described in this section.

At LHO, estimations of the test mass radius of curvature change can be made using the Hartmann Wavefront Sensors, which capture the thermal state of the test mass. A finite element analysis model using the Hartmann images during LHO power ups showed that the ITM radii of curvature can increase by almost 1% due to the thermalization of the interferometer. For the ETMs, this value is just under 0.5%. Figure 75 demonstrates the effect of small radius of curvature changes on the test masses on the DHARD P transfer function. For a given operating power, the radii of curvature are varied within 1% for the ITMs and 0.5% for the ETMs. The largest effect is on the high frequency pole near 2.6 Hz. For a longer radius of curvature, the pole will shift down in frequency, and the phase response around the pole will change. This transfer function was modeled using interferometer simulation software, *Finesse3* [109].

If this effect is not accurately accounted for and free modeling parameter is only the circulating power, the model will misattribute a lower intracavity power to a given open loop measurement. Taking this factor into account, Figure 76 and 77 model the DHARD and CHARD P open loop gain with 410 kW of circulating power, estimated from the 76 W input power. A comparison with an unbiased open loop gain measurement and calculated measurement error at that operating power is shown. The radius of curvature change due to the uncompensated thermal lens is also included in the model, shifting the ITM RoC to




Figure 77: Model and Measurement of the DHARD P open loop gain. DHARD P measurement taken using the unbiased open loop gain method detailed in Section 4.4. Modeled analytically using the methods described in this section.

1957 m and the ETM RoC to 2255 m. The full open loop gain is calculated according to Equation 6.151. It accounts for many features of the measured open loop gain, with some discrepancies in the phase near the 2.7 Hz pole. This suggests that the full thermal effect has not been sufficiently modeled to capture the phase evolution.

Additionally, these thermal effects may have some effect on the cavity matrix diagonalization. The cavity basis of the HARD and SOFT diagonalization is determined according to the expected cavity geometry, and further adjusted to account for suspension actuation strengths. If the thermal effect is significant enough at high power, this could degrade the effectiveness of the cavity alignment drive matrix, so that the HARD and SOFT modes become increasingly more coupled in the hot state. Figure 78 models the effect of a 10% drive mismatch on the DHARD Y suspension plant for a given cavity power. The largest effect occurs between 0.4 and 0.5 Hz.

Measurements of the CHARD Y and DHARD Y open loop gain indicate some mismatch occurring between 0.4 and 0.5 Hz. Figures 79 and 80, taken at 60 W input power, show a model of the open loop gain using an estimated 360 kW intracavity power. The drive mismatch is not included in the models on the plot. However, the discrepancy near 0.5 Hz could possibly be explained by some drive variation. This is also a potential source of the 0.5 Hz instabilities that were experienced during the power up process. CHARD Y in particular




Figure 78: The effect of drive mismatch on the HARD yaw suspension plant, modeled using *Finesse3*. This model shows that alignment drive mismatch could explain some of the anomalous features in the DHARD and CHARD Y OLG measurements.




Figure 79: Model and Measurement of the CHARD Y open loop gain. CHARD Y measurement taken using the unbiased open loop gain method detailed in Section 4.4. Modeled analytically using the methods described in this section.




Figure 80: Model and Measurement of the DHARD Y open loop gain. DHARD Y measurement taken using the unbiased open loop gain method detailed in Section 4.4. Modeled analytically using the methods described in this section.

shows this rapid gain evolution right around unity gain. Note that these models do also account for the test mass radius of curvature change due to the higher operating power.

All models shown in this section use an analytical method, calculating the expected open loop gain using suspension models, known radiation pressure dynamics, and some adjustments in the parameters to match the phenomenological criteria. Overall, this modeling work has been fairly successful at predicting HARD loop dynamics.

6.6.2 Modeling the SOFT Loops

Unfortunately, the same success in HARD loop modeling is not found in the SOFT loops. The SOFT loops are the most marginal loops controlling the arm cavities. The SOFT loop sensors experience significant cross-coupling, as discussed in Chapter 5. This makes these loops difficult to reliably measure, and therefore also difficult to model.

Figure 81 compares one such SOFT loop model with a measurement. First, notably, is the significant measurement uncertainty around a few tenths of a Hz. This is also exactly where the model of the CSOFT pitch loop diverges significantly with the expected mode. CSOFT pitch is set to be a low bandwidth AC coupled loop that provides additional suppression of the low frequency suspension resonances near 0.5 Hz. Similarly, the control is offloaded to the top mass, although the AC coupling of the loop means that the majority of the control

lies within the PUM itself.




Figure 81: Model and Measurement of the CSOFT P open loop gain. CSOFT P measurement taken using the unbiased open loop gain method detailed in Section 4.4. Modeled analytically using the methods described in this section. This model misses several important details, such as the low Q resonances and right-half-plane pole.

The measurement further indicates the presence of a right-half-plane pole around 0.5 Hz, which cannot be recreated in a model using drive mismatch. Figure 82 shows a model similarly varying the arm cavity drive matrix within 10%. No effect is observed at 0.5 Hz, where the model discrepancy lies. A similar variance of test mass radius of curvature also shows little effect on the transfer function.

Similar unstable dynamics are present in all SOFT loop transfer functions. These effects so far cannot be attributed to radiation pressure or thermal effects. Due to the inherent instability of the loops, it is unlikely measurements of the open loop gain will provide more answers, as any significant drive around 0.5 Hz to resolve the loop features is destabilizing to the control.

6.7 Conclusions

Increasing interferometer operating power is a key design feature for current and future gravitational-wave interferometers. This section presents a work-through of the alignment controls challenges at high power, and some methods for combating them. Effective modeling of alignment controls is a valuable tool during high power commissioning as a way to




Figure 82: Modeled drive mismatch on the CSOFT pitch suspension plant. Modeled using `Finesse3`. This modeling attempt shows that drive mismatch cannot explain the odd features seen in the CSOFT P open loop gain measurement.

understand the strange controls behavior that often arises at these operating powers. This chapter describes the efforts to model the alignment controls at LIGO Hanford during the high power commissioning to 400 kW. Modeling the HARD loops proved mostly successful, which has enabled controls upgrades to promote better detector stability and improve the overall noise. However, much work still needs to be done to understand the dynamics of the SOFT loops, which still may be impacting the stability of the interferometer.

Chapter 7

The Performance of the Laser Interferometer Gravitational-Wave Observatory During the Fourth Observing Run

This chapter contains excerpts from a paper in preparation by Capote, Jia, Aritomi, Nakano, and Xu.

7.1 Introduction

On May 24, 2023, the LIGO Hanford and LIGO Livingston Observatories began the fourth gravitational-wave observing run. The first half of this run, dubbed O4a, ran until January 16th, 2024. Both detectors underwent significant upgrades in the intervening period between observing run three and observing run four to achieve unprecedented sensitivity.

7.2 O4 Overview

The Advanced LIGO detectors achieved unprecedented sensitivity in the fourth observing run, O4. This section details the performance and operation of the detectors in O4, including astrophysical range, detection rates, duty cycle, and lock acquisition.

7.2.1 Astrophysical Range

During O4a, 81 non-retracted, high-significance public alerts were released [139]. Figure 83 shows the cumulative event rate in O4a compared with the event rate in previous observing runs [139, 140]. The increase in sensitivity of both detectors corresponds with an increased event detection rate relative to previous observing runs.




Figure 83: Cumulative events up to January 16, 2024, including the entirety of O4a. Candidates in O3 and O4a are the set of preliminary candidates identified in online real-time triggers [139]. The events represented in O1 and O2 are all the identified candidates from that observing period announced in GWTC-1 [140]. The increase in event detections in O4a results from the increase in sensitivity of the detector network from O3.

Figure 84 shows the binary neutron star (BNS) range of two LIGO detectors during O4a. The BNS range provides a standard metric of sensitivity for gravitational-wave detectors, using the distance to which a $1.4 M_{\odot}$ - $1.4 M_{\odot}$ BNS merger can be observed with a signal-to-noise ratio of 8. This range estimate accounts for the antennae pattern of the two LIGO detectors, as described in [52].

In O4a, the detectors achieve a BNS range of 130-160 Mpc for LHO and 145-164 Mpc for LLO. Detector commissioning work continued during the O4 run, increasing the detector sensitivity at both sites. Figure 84 shows that the O4a median range for the Livingston detector is well-distributed around 156 Mpc, while the Hanford detector range has a bimodal distribution around a median range of 148 Mpc. The lower peak of the range is around 140 Mpc while the upper peak is around 155 Mpc, similar to LLO. The significant increases in LHO detector sensitivity come from the reduction of input power shortly into the run that enabled low frequency noise improvements (Sec. 7.5.2), and the reduction of squeezer crystal losses that enabled higher squeezing levels.

7.2.2 Duty cycle

The duty cycle of the detectors in O4a is 67.5% for LHO and 69.0% for LLO. The duty cycle of coincident observing of two detectors is 53.4%. The duty cycle in O4a is lower than that in O3 by 7-8% for both detectors [44]. Several factors impacted the decreased duty cycle relative to O3. The LIGO Hanford duty cycle was impacted by lock stability from operating at higher power at the start of the run. Further details on this effect will be discussed in Section 7.5.2. Both detectors have a longer locking process relative to O3 due to higher power operation and more complex subsystems. Other factors that impacted duty cycle during O4 include microseismic activity, storms, wind, earthquakes, and even logging near the LIGO Livingston detector.

7.2.3 Lock Acquisition

In order to detect gravitational wave signals, the interferometers must be in the “locked” state: all mirror lengths and alignments are controlled within a stable linear regime. The procedure to bring the inteferometer from an uncontrolled state to this stable state is a lengthy process called “lock acquisition”. This lock acquisition process is described in detail in Buikema et al. [44].

One significant difference in the lock acquisition process relative to previous observing runs is the addition of the filter cavity for the injection of frequency dependent squeezing. Once the main interferometer locking procedures have been carried out, the filter cavity is locked and frequency dependent squeezing is injected. Once squeezing is injected and all other locking procedures are complete, the interferometer enters observation mode.




Figure 84: Range trends for the LIGO Hanford (red) and LIGO Livingston (blue) observatories during the first half of the fourth observing run, O4a. The left plot shows the median hourly binary neutron star (BNS) range of the LIGO Hanford and LIGO Livingston detectors over O4a. The BNS range from Observing runs 1-3 are shown for comparison. In O4, both detectors increased sensitivity over the course of the run, especially the Hanford detector. At Hanford, significant range improvements result from lowering the input laser power (mid-June 2023), and the reduction of squeezer crystal losses (mid-October 2023). The right plot shows a range histogram for O4a. Both detectors achieved maximum BNS ranges of over 160 Mpc, with the Livingston detector achieving close to 170 Mpc.

The squeezer locking is performed simultaneously with the interferometer locking, and is been described in detail in [21]. Squeezer locking stabilizes the relative phase between the squeezed vacuum and interferometer beams at the readout photodetectors. This relative phase is often called the “squeeze angle,” and it determines the optical field quadrature (e.g. phase, amplitude) in which squeezing reduces quantum noise. The squeezer laser is independent of the interferometer laser, and so the two lasers must be phase-locked. For this, an RF sideband of the squeezer laser called the Coherent Locking Field (CLF) is generated and injected into the squeezer along with the pump light [141]. The CLF copropagates with the squeezed vacuum beam to the readout photodetectors. The photodetector signal is demodulated at the CLF frequency to sense the relative phase between the squeezed beam and interferometer beam. This error signal is used to by actuate on the frequency of the squeeze laser, and thus stabilizes the relative phase (squeeze angle) between the squeezer and interferometer lasers at the output photodetectors.

In addition to locking the squeeze angle, the filter cavity length and detuning must be stabilized. To lock the filter cavity, another RF sideband called the Resonant Locking Field (RLF) is injected alongside the CLF. The RLF sideband is designed to resonate in the filter cavity, unlike the CLF sideband. Comparing the resonant RLF to the off-resonant CLF, the beat signal between the CLF and RLF is demodulated at their difference frequency to obtain an error signal for filter cavity length and alignment controls. The RLF sideband frequency is nearly an integer multiple of the filter cavity’s free spectral range (FSR), such that when the RLF resonates, the squeezed vacuum beam is nearly simultaneously resonant in the filter cavity. By changing the RLF frequency, the filter cavity resonance is detuned with respect to the squeezed vacuum frequency. Thus, changing the RLF frequency changes the filter cavity detuning, and allows for optimization of the frequency-dependent squeeze angle rotation [142].

The full locking process is automated, and managed by a state machine called *Guardian* [143]. Locking time can vary depending on the initial state of the mirrors and the environmental conditions. In good alignment and low environmental disturbance, locking can take 30 - 45 minutes. This process can extend beyond one hour if the environmental conditions are poor. The lock acquisition steps that take the most time are related to the convergence time of slow alignment loops, the careful hand-off of different sensors and actuators, and thermalization effects as the optics absorb incident power.

If a disturbance causes a controlled degree of freedom to move outside the linear regime, a

“lock loss” occurs. The locking process must then begin again. Lock losses can be caused by a variety of phenomena, such as earthquakes, storms, high winds, control instabilities, and instrumental glitches. Still, to-date the majority of locklosses occur for unknown reasons.

7.3 Instrument Noise

The LIGO Hanford and Livingston detectors are limited by various noise sources stemming from both fundamental limits to gravitational wave detectors and technical noises due to the operation of the detectors. The best understanding of the limiting noise sources in these detectors is shown in Figure 85. The plots demonstrate noises quantified in two distinct ways. Generally, fundamental noise sources can be derived from first principles, and are calculated from knowledge of the detector’s optical and mechanical properties. Technical noises in the interferometer are directly estimated via projections from auxiliary channels that serve as a witness.

The traces shown in dashed lines in Figure 85 represent estimated fundamental noise sources that provide an upper limit to the sensitivity of a gravitational wave detector. The quantum noise is estimated through both measured and modeled parameters related to the operating power of the detectors and operation of the squeezing subsystem. Thermal noise is estimated through external laboratory measurements of the mirror coatings, substrates and suspensions. Residual gas noise is estimated through the understanding of the achieved high vacuum levels at each detector.

The traces shown in dotted lines in Figure 85 represent directly measured technical or fundamental noise sources. These estimations are made via a measurement of a witness channel, w , and its contribution to the differential arm length channel, d . The amount of noise contributed from a witness channel to the strain readout is estimated via [44]

$$S_{w \rightarrow d}(f) = G(f)S_{w,0}(f), \quad (7.152)$$

where $G(f)$ is a coupling factor that is calculated from a direct injection. The excess noise in the witness channel and strain channel during the excitation is compared to the ambient noise in both channels [44]

$$G(f) = \frac{S_{d,\text{exc}}(f) - S_{d,0}(f)}{S_{w,\text{exc}}(f) - S_{w,0}(f)}. \quad (7.153)$$

This method allows the calculation of both linear and nonlinear couplings to be measured, and can be used to measure channels with very low coherence to the main channel.

Finally, the sum of the known noises is calculated and compared to the measured noise at each detector. Discrepancies between these two traces in Figure 85 indicate regions where unknown or nonstationary noise limits the gravitational-wave strain sensitivity.


At low frequencies below 50 Hz, the detector noise differences are most visible, largely due to the differences in scattered light control, controls noise, seismic noise, and input beam jitter. At intermediate frequencies between around 50 - 500 Hz, detector sensitivity is primarily limited by thermal noise, and secondarily by quantum noise. At kilohertz frequencies, the two detectors achieve comparable shot-noise-limited strain sensitivities, with slight differences in the circulating laser power and squeezing efficiency. Both detectors have some amount of unknown noise, especially below 50 Hz, where the total measured detector noise exceeds the sum of known noises.

Of the noises shown in Fig. 85, the estimated residual gas and seismic noises are the same as in O3 [44]. Here, we describe the limiting noise sources of the LIGO detectors in the fourth observing run, O4.


7.3.1 Quantum noise

Quantum noise of interferometer originates from vacuum fluctuations at the antisymmetric port [144] and limits detector sensitivity across most of the astrophysical band. Quantum noise manifests as both shot noise and quantum radiation pressure noise. Shot noise can be understood as fluctuations of the phase quadrature of the vacuum field entering the antisymmetric port, and it is the dominant noise source at high kilohertz frequencies. Quantum radiation pressure noise results from fluctuations in the amplitude quadrature and dominates at low frequencies. At the crossover frequency around 40 Hz, the two forms of quantum noise contribute equally.

Non-classical states of light known as “squeezed vacuum” can be injected at the antisymmetric port to reduce interferometer quantum noise. Quantum correlations between photon pairs reduce the noise variance in one optical quadrature of the squeezed state at the expense of increased noise in the conjugate quadrature, e.g. reducing fluctuations in phase at the cost of increasing fluctuations in amplitude. This associated increase in amplitude noise is a form of quantum back-action. The previous observing run, O3, marked the first



(a) LIGO Hanford Observatory (LHO) noise budget, Dec. 2023.



(b) LIGO Livingston Observatory (LLO) noise budget, Oct. 2023.

Figure 85: Noise budgets for the LIGO Hanford detector (LHO, upper) and LIGO Livingston detector (LLO, lower) in strain noise amplitude spectral density units. Solid lines show calculated noise contributions. Dots show measured and projected noise contributions. (Black) lines show the total expected detector noise (sum of budgeted noises), while the blue lines show the measured detector noise. Most discrepancies between expected and measured noise are at low frequencies, where scattered light noise is prominent. Individual noise terms are discussed in Section 7.3. Both sites face similar limiting noises: noise from angular and length controls and scattered light below 50 Hz, a mix of quantum noise and coating thermal noise between 50-500 Hz, and photon shot noise above 500 Hz.

astrophysical run of the LIGO detectors that used squeezed light injection to enhance detector sensitivity [145]. In O3, squeezing improved kilohertz sensitivity by 2.0 dB (LHO) and 2.7 dB (LLO), at the cost of increased amplitude noise that could degrade low-frequency performance.

In O4, broadband quantum noise reduction is achieved by injecting frequency-dependent squeezing [21] to reduce the limiting form of quantum noise at all frequencies. Using a new 300-m optical “filter cavity” installed at each detector, frequency-dependent squeezing injects amplitude-squeezed light at low frequencies where the mirrors are susceptible to quantum radiation pressure noise, and phase-squeezed light in the shot-noise-limited region above 500 Hz. By mitigating back-action noise, frequency-dependent squeezing further allows for higher injected squeezing levels. Section 7.4.3 summarizes the upgrades made to squeezer system for O4. Altogether, in O4, the LIGO detectors observed squeezing levels of up to 5.4 dB (LHO) and 5.8 dB (LLO) around 2 kHz.

The purple traces in Figure 85 show the calculated quantum noise models, given the best estimates of interferometer and squeezer parameters. The observed squeezing is limited by optical losses, fluctuations of the optical field quadrature being squeezed (phase noise), and mode mismatch amongst the various cavities in LIGO. The Gravitational Wave Interferometer Noise Calculator (GWINC) is used to calculate the full detector quantum noise models. Below 40 Hz, quantum noise is limited by quantum radiation pressure, and is thus sensitive to the circulating laser powers and filter cavity parameters. At intermediate frequencies between 100 - 450 Hz, where squeezed light begins to resonate in the signal recycling cavity, quantum noise is sensitive to many detector parameters such as the detuning of the signal recycling cavity, homodyne readout angle, circulating arm power, and squeezer mode-matching. Above the coupled-cavity pole at ~ 450 Hz, the shot-noise-limited strain sensitivity is determined by the circulating laser power and the squeezing efficiency, both of which are limited by losses and mode mismatch.

7.3.2 Thermal noise

Thermal motion in the mirror suspensions, mirror substrates, and optical coatings contributes noise in gravitational wave detectors. The most significant noise source of these three in Advanced LIGO is coating thermal noise (CTN), arising from Brownian motion in the optical mirror coatings, limiting the detector sensitivity around 100 Hz. The Advanced LIGO test masses are mirrors using Bragg reflectors to create high reflectivity surfaces. The

Bragg reflector consists of alternating layers of low- and high-index of refraction materials. In the Advanced LIGO case, the low index material is silica, SiO_2 , and the high index material is tantala doped with titania, $\text{TiO}_2\cdot\text{Ta}_5\text{O}_2$. Thermal noise arising from an incident Gaussian beam can be shown, via the fluctuation dissipation theorem, to depend on the loss angle of the material [29, 25]. Of those two materials, titania-doped tantala has a higher loss angle, and it is the loss from this material that sets the current CTN limit in aLIGO.

To quantify the CTN in the Advanced LIGO interferometers, smaller witness samples are produced during the test mass coating process. The loss angle and thermal noise of these samples are measured, and the expected thermal noise in the detector is calculated. In O3, the estimated CTN was 1.3×10^{-20} m/ $\sqrt{\text{Hz}}$ at 100 Hz, resulting from a direct measurement of the CTN [146]. This estimate has been updated to be 1.1×10^{-20} m/ $\sqrt{\text{Hz}}$ at 100 Hz [66] using the same measurement method. Based on these measurements, the form of coating thermal noise from the Advanced LIGO coatings follows

$$\sqrt{S_{\text{CTN}}} = 1.1 \times 10^{-20} \left(\frac{100 \text{ Hz}}{f} \right)^{0.45} \frac{\text{m}}{\sqrt{\text{Hz}}}. \quad (7.154)$$

The coating thermal noise trace in Figure 85 plots this measured trend.

New end test masses were installed in the Livingston Observatory ahead of O4 to remove point defects that limited the operation of the detector [131, 147]. Since the replacement of these test masses, LLO has observed a higher-than-expected level of noise around 100 Hz that appears similar to coating thermal noise. For further information about this test mass install, see Section 7.4.2. While witness samples produced during the coating of these new optics demonstrate thermal noise at the expected aLIGO level, the total detector noise observed at LLO between 100 - 200 Hz fits to a CTN amplitude of 1.79×10^{-20} m/ $\sqrt{\text{Hz}}$ at 100 Hz, with a slope of $1/f^{0.5}$. This noise source appears unaffected by changes in interferometer power or the injected squeezing. While this excess noise source appears similar to coating thermal noise, it is important to note that laboratory measurements suggest the expected frequency dependence of CTN from the aLIGO mirror coatings is $1/f^{0.45}$, as in Equation 7.154 [66]. Currently, there is no other explanation for the source of this excess noise. Correlated noise measurements at both observatories see this excess noise around 100 Hz, with higher CTN amplitudes than expected from witness sample measurements, and a slightly different frequency dependence ($1/f^{0.5}$) than expected ($1/f^{0.45}$).

A fit of this excess noise has been performed in the Livingston detector strain data. The result of this fit is shown in thermal noise sub-budget in Figure 86. Due to the similarity in

behavior to coating brownian noise, it is referred to as a “coating brownian curve fit”. With this fitted noise, the excess noise from 60 Hz to 200 Hz witnessed in the LLO noise budget in Figure 85 is fully explained. There is currently no other method to confirm the identity of this noise source, but work is underway to remeasure the CTN of the witness samples. Better understanding of the quantum squeezing parameters also enables better fits of this excess noise witnessed at both detectors.

The test mass substrates are made of high quality fused silica, and therefore the mechanical loss of the substrates is very low. The resulting mechanical resonances of the substrates have a very high quality factor (Q), which can impact the operation of the detectors. Therefore, Advanced LIGO also makes use of acoustic mode dampers to slightly reduce the Q of the test mass resonances. Acoustic mode dampers (AMDs) are piezoelectric dampers tuned to reduce the Q of the test mass mechanical resonances over specific frequency ranges [148]. Based on external measurements, the noise degradation due to thermal noise from AMDs is estimated to be less than 1%. Figure 86 demonstrates the thermal noise contributions from each of these sources mentioned here in the Livingston detector strain.




Figure 86: Thermal noise sub-budget for the LIGO Livingston detector. Both the fitted Coating Brownian noise and the witness-sample measured Coating Brownian is shown. Coating Brownian noise is potentially a limiting noise source in the current detectors.

7.3.3 Laser noise

Intensity noise, frequency noise, and spatial fluctuations of the laser (beam jitter) also couple with gravitational wave signals. Generally, intensity noise and frequency noise are described as sidebands on the main laser carrier frequency. Beam jitter can similarly be represented as sidebands of the Hermite-Gaussian first order mode. These noise sidebands enter the interferometer from its input ports, but ideally, in a perfectly symmetric interferometer, these sidebands are reflected back to the input ports and do not appear at the anti-symmetric port where gravitational wave signals are detected. This is known as common-mode rejection. However, in real interferometers, asymmetries exist due to macroscopic length difference in arms, differences in mirror reflectivity, etc., leading to a decrease in the efficiency of common-mode rejection, allowing noise sidebands to couple with gravitational wave signals. Therefore, it is necessary to stabilize the input laser frequency and intensity noise, and to mitigate the input beam jitter.

Laser frequency noise

The laser frequency is stabilized using three optical cavities as the frequency references. The first reference is a 10-cm linear reference cavity made of ultra low expansion glass, located on the in-air optical table. Another optical cavity employed as a frequency reference is the input mode cleaner (IMC), a 25-m triangular cavity comprised of three suspended mirrors. The ultimate frequency reference is the 4-km arm cavity length. While the motion of the differential arm length is utilized for gravitational wave detection, the motion of the common arm length (CARM) is employed as the frequency reference. Errors from resonant frequencies of these references are fed back to the laser frequency through the nested feedback loop. The CARM, IMC, and the reference cavity stabilize the laser frequency in different frequency bands: below 10 kHz, between 10 kHz and 50 kHz, and above 50 kHz to 300 kHz, respectively. Fig. 88 shows the frequency noise and its DARM projection at each site. The frequency noise is calibrated by the error signal of the stabilization loop by CARM, resulting in similar noise levels at both sites. However, the noise coupling differs between the two sites due to higher order mode coupling caused by differences in the thermal state of the interferometer. When the curvature of the test mass is changed by using the thermal compensation system, the frequency noise coupling also changes significantly. This not only affects the coupling level but also alters the response itself, resulting in variations in the




Figure 87: The laser noise subbudget for both detectors, show the contributions of frequency noise, intensity noise and beam jitter noise.

frequency and width of dips in the coupling function appearing in the LHO curve of Fig. 88. Adjustments are made to the thermal compensation system to ensure that the frequency noise does not limit the sensitivity.

Laser intensity noise

Laser intensity noise is stabilized by two feedback loops. The first loop involves a photodetector installed on an in-air optical table, which is fed back to an acousto-optic modulator also installed on the same in-air optical table. This loop operates within a bandwidth of 80 kHz. The second loop stabilizes the output power of the IMC. Here, eight PD arrays installed in vacuum are used to stabilize the power with larger power than the first loop,




Figure 88: Laser frequency noise for both detectors. The top plot shows the projected frequency noise compared to the total detector noise at LHO. The middle plot shows the calibrated laser frequency noise. The frequency noise is measured with the in-loop error signal of the CARM loop. LHO has an out-of-loop sensor for the CARM loop to obtain both the in-loop and out-of-loop frequency noise, while LLO has only the in-loop sensor. The bottom plot shows the coupling function from laser frequency noise to total detector strain noise.

effectively reducing shot noise, which limits the stability of noise. This loop has a bandwidth of 30 kHz. These loops stabilize the relative intensity noise down to a level of $4 \times 10^{-9} 1/\sqrt{\text{Hz}}$ within the bandwidth of 30-1 kHz.

Beam jitter noise

Input beam jitter noise is one of the most significant sources of noise masking the gravitational-wave signal in O4. Currently, LIGO does not have an active stabilization system for beam jitter noise similar to other laser noises, but a passive reduction of jitter is achieved through the IMC. As mentioned above, beam jitter can be described as the Hermite-Gaussian first order mode, and in a cavity resonating at the fundamental mode, higher-order modes are filtered out, leading to a reduction in the input beam jitter in the output beam of the cavity.

However, the input beam jitter is significant, limiting sensitivity in multiple bands between 100-1000 Hz at the LHO and around 500 Hz at the LLO. Jitter noise coupling, like frequency noise, is influenced by the thermal state of the interferometer. At LHO the input jitter coupling functions increased by a factor of about 10 as the power was raised from 46 W to 75 W and decreased by a factor of about 2 when the power was lowered to 60 W. One hypothesis is that the higher power produces greater thermal distortions of test mass surfaces around coating defects, making the arms less symmetrical and reducing common mode rejection of input noise. Scattering noise from the cryobaffles also increased with power for both sites, possibly also due to greater surface distortions, prompting warnings about increased vibration coupling with power for both input jitter and scattering noise.

In the frequency band above 100 Hz, it can be seen from the bottom plot of Fig. 89 that the coupling at LLO and LHO, excluding pitch coupling at LHO, are at similar levels. However, in the DARM projection, input beam jitter noise at LHO is about an order of magnitude larger than at LLO around 100-300 Hz. This is because the environmental noise in the laser room is larger in LHO compared to LLO, which causes the larger jitter as shown in the middle plot of Fig. 89. This jitter is known to be introduced by a periscope installed to correct the height difference between the optical system in the laser room and the interferometer plane. In LHO, the vibration of the periscope caused by environmental noise is larger than that in LLO.




Figure 89: DARM projection of the input beam jitter noise (top), input beam jitter noise (middle), and input beam jitter coupling function (bottom). The input beam jitter noise was estimated by quadrant photodiodes that witness beam motion in reflection of the input mode cleaner cavity.

7.3.4 Controls noise

The length and alignment control topology of the interferometer is divided into the control of the arms and control of the auxiliary lengths that make up the dual-recycled Michelson portion of the optical layout. Length control of the arms is divided into the common arm and differential arm length control, with common motion read out from the symmetric port,

or on reflection of the interferometer, and differential motion from the anti-symmetric port, or transmission of the interferometer. Auxiliary lengths include the Michelson cavity length (MICH), formed between the beamsplitter and the two input test masses, power-recycling cavity length (PRCL), formed between the power-recycling mirror (PRM) and the average distance of the input test masses, and signal-recycling cavity length (SRCL), formed between the signal-recycling mirror (SRM) and the average distance of the two input test masses.

These auxiliary length degrees of freedom are read out via a pick-off port of the power-recycling cavity, on transmission of power-recycling mirror 2 (PR2), where an RF photodiode reads out both the 9 MHz and 45 MHz demodulations. Their actuators are the beamsplitter, PRM and SRM, respectively.

Auxiliary alignment of the arms is divided into the common and differential cavity axis alignment, similarly sensed at the symmetric and anti-symmetric port, respectively, using an RF QPD known as a wavefront sensor (WFS), which is four RF photodiodes in a quadrant array. A pitch misalignment is sensed by subtracting the signal of the upper and lower quadrants, and a yaw misalignment is detected from the subtraction of the left and right quadrants. A misalignment of the arm cavity axis beats a first-order carrier beam mode with a zeroth order RF sideband mode to provide the alignment signal for control. The translation of the cavity axis is sensed via DC QPDs on transmission of the arms. The pitch and yaw signal is extracted in a similar way, except this is a purely DC signal with no sideband component. Finally, dither schemes and camera servos provide slow control of the spot positions when an absolute beam position reference is required. Camera servos were commissioned for the first time in for observing run four; more details on this upgrade can be found in Section 7.4.4. Dither schemes are used in lock acquisition, but then transitioned to camera servos to reduce noise from alignment controls in the gravitational wave strain channel.

The other alignment control involves the alignment of the input axis of the interferometer, controlled from a mirror on the input telescope and the power-recycling cavity alignment, controlled from PR2, both sensed at the symmetric port. Additionally, the signal-recycling cavity alignment, controlled from the SRM and the signal-recycling cavity pointing, controlled from SRM and SR2 are sensed at the anti-symmetric port. All cavity axis misalignments are controlled from wavefront sensors distributed at these two ports, while the signal-recycling cavity pointing degree of freedom is controlled from a DC QPD at the output of the signal-recycling cavity.

Auxiliary length control noise

While control of the auxiliary lengths is required to keep the interferometer operating in the sensitive linear regime, this control comes at a cost. The gravitational-wave readout is sensitive to motion of the Michelson interferometer, power-recycling cavity, and signal-recycling cavity, and therefore noise due to these controls couples into the gravitational-wave strain.

These control loops are designed to provide enough suppression of the large low frequency ground motion to hold the interferometer in the linear operating range, while reducing the re-injection of sensing noise in the gravitational-wave band. Despite careful controls design to minimize noise re-injection, these auxiliary length controls still imprint excess sensing noise in-band. To reduce the noise contribution of these loops to the gravitational-wave signal, a feedforward cancellation is applied.




Figure 90: A comparison of the differential arm length (DARM) and Michelson length (MICH) coupling at LHO. With no feedforward applied, coupled Michelson noise dominates the differential arm length at low frequency. With an initial feedforward tuning, this coupled noise is suppressed to no longer limit total detector noise. Iterative tuning of the feedforward provides improved signal-to-noise in the injection to suppress it further.

Feedforward cancellation is a method of real-time noise cancellation by mechanically driving the expected sensing noise to the test masses. It is made possible by the fact that

the residual noise of the auxiliary lengths is witnessed by independent sensors from the gravitational-wave readout. The transfer function from the auxiliary loop to gravitational wave signal are each measured, as well as the transfer function of the feedforward actuation path. The transfer functions are fit using fitting algorithms such as `vectfit` or `IIRrational` [149, 150]. The opposite sign of the control signal of this auxiliary loop is injected into the differential arm length using the measured transfer function to cancel out the noise that would appear. This feedforward scheme can reduce the contribution of auxiliary length noise in the differential arm length by a factor of 100 or more in the gravitational wave band. Figure 90 shows how the noise contribution of the Michelson length motion changes with feedforward control applied. The coupling of this noise source to differential arm length can be demonstrated by

$$\frac{|S_{wd}(f)|^2}{S_w(f)} = \gamma_{wd}^2(f)S_d(f) \quad (7.155)$$

where $S_w(f)$ and $S_d(f)$ are the power spectral densities of the auxiliary witness channel and the differential arm length, respectively, $\gamma_{wd}^2(f)$ is the coherence (in power) between these channels and $S_{wd}(f)$ is the cross spectral density between the two channels.

The error on feedforward control depends on the coherence of the measurement [100]

$$\sigma = \sqrt{\frac{1 - \gamma_{wd}^2(f)}{2N\gamma_{wd}^2(f)}} \quad (7.156)$$

where N is the number of averages taken during the measurement. Therefore, the ability to subtract noise via feedforward depends strongly on the precision of the measurement. To overcome this challenge, an iterative method of feedforward tuning was applied [151]. After initial injection, an appropriate feedforward filter is tuned to subtract the excess noise. Next, the same injection is applied, with the new feedforward engaged. This allows residual linear coupling to be measured and suppressed further with an improved filter. Figure 90 compares Michelson noise coupling to differential arm length in each of these scenarios.

Feedforward cancellation was applied to suppress the noise from the Michelson and signal-recycling cavity loops in O4. While the power-recycling cavity does couple as well, this coupling is usually indirect, and through improper sensor diagonalization or other optomechanical coupling to the other lengths. Therefore, it can be reduced through the main feedforward application of MICH and SRCL. With the iterative tuning method, a factor of one hundred or more suppression of Michelson length noise is possible, and a factor of thirty suppression of signal recycling cavity length noise is possible.

Alignment control noise

Alignment control is performed to suppress motion from the environment, improve optical gain, reduce angle-to-length coupling, and stabilize mirrors against radiation pressure effects. Again, this comes at the cost of noise injected into the gravitational-wave readout. In principle, alignment control of all arm cavity and auxiliary mirrors couples into the differential arm length, but the arm cavity control dominates the noise from alignment control. Any change in arm cavity mirror alignment is a direct change in arm length.

The most significant coupling results from the control of the arm cavity alignment control. Eight control loops govern the common and differential cavity axis alignment and cavity pointing in pitch and yaw. The cavity axis alignment is controlled with bandwidths of up to 5 Hz, and the cavity pointing alignment with lower bandwidths that are no more than 1 Hz. Figure 91 shows the four most limiting alignment control loops and their contribution to differential arm length.

During lock acquisition, the alignment control loops are kept in a high bandwidth, or high noise, state to better suppress disturbances in the alignment that result from procedures such as increasing the intra-cavity power. Once the final power is achieved, the control loops are transitioned into a low noise state that reduces the bandwidth of the loops and implements aggressive low pass filters that suppress the injection of alignment sensing noise above 10 Hz.

While most wavefront sensors and QPDs are sensing-noise-limited above 10 Hz, the WFS at the symmetric port are mounted on an in-vacuum table with poor isolation from the ground. As such, there is excess noise imprinted on the signal from 10-30 Hz. Feedforward subtraction using table accelerometers as witnesses is implemented to reduce the presence of this noise. Further details can be found in Section 7.4.4.

Both LIGO detectors were limited by noise from alignment controls during O3. Significant equipment and controls upgrades were implemented ahead of O4 to mitigate some of this noise and prepare for high power operation. These upgrades will be further detailed in Section 7.4.4.

Local suspension controls noise

Each suspension is controlled locally to suppress excess motion, namely around the suspension resonances. The local sensors used for these damping loops are optical shadow sensors, known as A and B OSEMs [122]. The local control of each suspension is an AC-coupled




Figure 91: A sub-budget showing the limiting alignment control loops measured at the LIGO Hanford Observatory. These four loops control the arm cavity alignments, and have bandwidths between 3 and 5 Hz.

control design at the top mirror stage only with frequency cutoffs below 10 Hz to avoid sensor noise contamination in the gravitational-wave band. However, even with these cutoffs, OSEM sensor noise still contaminates the gravitational-wave strain.

The contribution of these local damping controls is shown in the Figure 85 noise budget traces labeled “Suspension damping”. The LLO noise budget shows contributions from the quadruple suspensions (QUADs) that control the test masses, and the large and small triple suspensions that control auxiliary mirrors (triples). The LHO budget shows contributions from only the QUADs. Both detectors experience noise contamination from the QUADs and triples, but the direct contribution of the triples to the total gravitational wave strain noise has only been measured at the Livingston detector. A sub-budget of the contributions from the quadruple suspensions and triple suspensions at the Livingston detector is shown in Figure 92.

Efforts were made during commissioning to reduce the bandwidth of the local controls to avoid sensor noise contamination. However, even with the reduced bandwidth, the low frequency RMS of auxiliary length and alignment controls is still dominated by OSEM sensor




Figure 92: Local damping noise contributions to gravitational wave strain noise measured at the LIGO Livingston Observatory. After local control upgrades, contribution from auxiliary optic triple suspensions is comparable to the contribution from test mass quadruple suspensions.

noise, in addition to the limits in the strain sensitivity.

Actuator noise

The digital-to-analog converters (DAC) used for the test mass control signals also contribute noise to the gravitational-wave strain. This digitization noise results from nonlinear behavior in the conversion of a discrete digital signal to a continuous analog signal.

Digitization noise can be reduced through a technique called dithering, which involves injecting high-frequency sinusoidal signals through the DACs. A large, high-frequency signal dominates the RMS of the DAC, such that the digitization noise at low frequency, where the key portion of the control signal is generated, is reduced. At LLO, a single dither line at 7 kHz was injected and improved digitization noise by a factor of 3, suppressing DAC noise sufficiently far below the total detector noise.

While a similar digitization noise is observed at LHO, the dithering method did not change the total noise level. Therefore, LHO operates without DAC dithering. At this time,

DAC noise is not a limiting noise source for LHO.

7.3.5 Photodetector dark noise

For O4, the output photodetector electronics improvements summarized in Section 7.4.8 reduced dark noise by almost 3-fold over the previous electronics. Dark noise is now measured to be approximately a factor of 10 in amplitude below the total strain noise.

7.3.6 Scattered light noise

Scattered light produces noise when light scattered from the main interferometer beam reflects off a vibrating source and re-interferes with main interferometer light. The peak frequency of this noise is proportional to the velocity of the vibrating source, meaning that even if the frequency of the vibrating source is low, scatter can still produce noise within the observation band of gravitational waves depending on the vibration amplitude and the optical power in the stray beams.

With the increased laser powers for O4, scattered light couplings appear to have increased across both detectors. From the noise budgets in Fig. 85, excess low-frequency noise is visible below 50 Hz, where both sites face a discrepancy between the total measured noise (red, LHO; blue, LLO) and the sum of known noises (black). Some of the unknown noise is likely from scattered light. Section 7.4.6 describes efforts to mitigate known sources of scattered light for O4, such as the removal of an in-vacuum window at the output port through which the main laser beam transmitted (significantly reducing acoustic coupling at the output), and the damping of various in-vacuum baffles. Still, vibration measurements suggest other sources of scattered light remain to be identified and mitigated.

The squeezer system presents a potential path for back-scattered light. Interferometer light that is rejected by the output faraday isolator is directed into the squeezer system, via the exact optical path by which squeezing is injected into the interferometer. New for O4, the filter cavity was a potential new source of back-scattered light from the squeezer: leakage light from the interferometer can be modulated by the filter cavity length control signal, and reflected back into the interferometer. This filter cavity backscatter can be estimated by exciting the filter cavity length control signal, and measuring the corresponding excess noise in the detectors. The filter cavity length control signal includes both the filter cavity length motion and photodetector dark noise. Therefore, the filter cavity backscatter projections

in Fig. 85 can be considered an upper limit. Due to the complexity and number of sources of scattered light in the interferometer, it is difficult to project all forms of scattered light noise. Therefore, only filter cavity backscatter noise is shown in Fig. 85.

7.4 Instrument improvements

Major detector upgrades were implemented between O3 and O4 to reduce quantum noise, controls noise, electronics noise, and noise from scattered light. Other upgrades include efforts to improve lock stability through the mitigation of parametric instabilities and improved controls. In this section, the instrument improvements between O3 and O4 are presented.

For the main interferometer, in O4, additional laser amplifiers were installed to increase the available input laser power, and the test mass mirrors were replaced to reduce the point defects that limited circulating powers previously [131]. Both LIGO detectors operated with input laser powers of 60-75 W in O4, up from the 30-40 W previously used in O3 [44]. With cleaner test mass mirrors, the higher input powers translated into higher circulating arm powers of 260 kW - 410 kW in O4, compared to the 200 kW - 240 kW used in O3. Upgrades to the alignment controls enabled stable operation at these higher arm powers, along with reduced noise at low frequency from controls. Together, these upgrades increased the unsqueezed shot-noise-limited strain sensitivity by around 40% in LHO, and 16% at LLO.

At the output port of the interferometer, the O4 squeezer system was upgraded to enable higher levels of frequency-dependent squeezing, achieving broadband quantum noise reduction with the highest squeezing levels observed to-date in the LIGO detectors [21]. By modifying the squeezed light source to generate higher squeezing levels, implementing high-efficiency Faraday isolators to reduce losses, adding new active optics to optimize mode-matching, and installing a high-finesse optical filter cavity for frequency-dependent squeezing, the LIGO interferometers achieved broadband quantum noise reduction of up to 5 dB (LHO) and 6 dB (LLO) in O4.

7.4.1 Pre-Stabilized Laser upgrade

The Pre-Stabilized Laser (PSL) was upgraded at both observatories to supply high power to vacuum. The PSL begins with a 2 W NPRO 1064 nm seed laser. The seed beam is passed through 2 neoVAN-4S-HP amplifiers, the first amplifying the beam to 70 W and

the second amplifying to 140 W of power. The amplified laser beam passes through an optical cavity called the pre-mode cleaner for spatial mode cleaning and filtering of intensity noise in the radio frequency (RF) band. The laser frequency is actively stabilized by using a rigid optical cavity as the frequency reference. The intensity is also actively stabilized using the photodetector and an acousto-optic modulator. The phase modulation for the main interferometer control scheme is applied to the PSL by using electro-optic modulator. The PSL beam path includes a rotation stage waveplate to control the laser power injected into vacuum. The maximum power that can be supplied to vacuum 110 W. In O4, both observatories use up to 60-75 W of PSL input power during observing.

7.4.2 Test mass replacements

Test mass mirrors were replaced at both observatories due to the presence of point defects in the mirrors, which contributed to control instabilities and excess noise from scattered light and beam jitter [131]. Control issues related to excess loss from point defects limit the achievable operating power of the detector, while excess scatter and jitter noise can degrade the detector sensitivity proportional to the input power.

At the Hanford Observatory, the input test mass mirror to the Y-arm (ITMY) was replaced for O4. The presence of a point defect on ITMY was noted during the commissioning for O3, and nonuniform absorption from this point defect compromised the operation of the detector, especially when increasing the operating power from 25 W to 40 W [44]. After replacing ITMY, the Hanford Observatory achieved 57 W of input power with a corresponding power-recycling gain of 50 W/W, compared to the 34 W input power and power-recycling gain of 44 W/W in O3. No point absorbers have been observed on the ITMY mirror since its installation. However, point absorbers remain on the input test mass mirror to the X-arm (ITMX) and both end test masses.

At the Livingston Observatory, the two end test mass mirrors (ITMY and ITMX) were replaced due to the observation of point absorbers on both end mirrors. After replacement, no point absorbers have been observed on either end mirrors according to Hartmann Wavefront Sensor images of the mirror surfaces during high-power operation. However, the power-recycling gain in the Livingston detector decreased with the installation of these new mirrors, dropping from 47 W/W with an input power of 38 W in O3, to 38 W/W with an input power of 63 W in O4.

7.4.3 Squeezer upgrades

Previously in O3, the measured squeezing levels of up to 2.7 dB in LLO and 2.0 dB in LHO were limited by loss, phase noise, and mode-mismatch [145, 152]. Use of higher squeezing levels was further limited by excess back-action noise in the form of quantum radiation pressure [153], which could degrade low frequency detector sensitivity.

These limitations were addressed for O4: the squeezed light source was modified to generate higher squeezing levels (from 7-9 dB to 17 dB); losses were reduced with three new high-efficiency Faraday isolators (from single-pass losses of 3-4% in O3, down to 0.5-1% in O4); and a 300-m optical “filter cavity” at each detector to prepare frequency-dependent squeezing that mitigates back-action noise to enable broadband quantum noise reduction [21, 154].

With these upgrades, the LIGO detectors achieved broadband quantum enhancement with up to 4.5(5) dB (LHO) and 5.5(5) dB (LLO) of quantum noise reduction visible in O4.

More generated squeezing. The “squeezer” is a sub-threshold optical parametric oscillator (OPO) that consists of a dually-resonant bowtie cavity surrounding a nonlinear PPKTP crystal. For O4, the squeezer cavity mirrors were changed to increase the finesse from 37 (17.6) to 85 (110) for 1064 nm squeezed (532 nm pump) light, allowing lower 532 nm pump powers to generate higher levels of squeezed light at 1064 nm. Issues with crystal degradation persisted in O4. In October 2023, both sites translated their squeezer crystal to recover green losses mid-run, with additional motivations at LHO to find a crystal position with reduced 1064 nm squeezing losses. Compared to the generated squeezing levels of 7 – 9 dB in O3, both detectors generated around 17 dB of squeezing in O4.

Lower losses. High-efficiency, in-vacuum Faraday isolators are necessary to meet LIGO’s strict requirements on stray light isolation. In O3, the squeezed beam passed 3 times through 2 Faraday isolators, each with 3 - 4% single-pass loss [145]. In O4, the squeezed beam passes 5 times through 3 new high-efficiency Faraday isolators, each with single-pass losses of around 0.5 – 1%, reducing total loss and improving stray light isolation.

Improved mode-matching. Squeezer light is sensitive to the mode-matching, or the spatial overlap, between the squeezed laser beam and the resonant modes of the various optical cavities. To actively change the mode-matching, four new deformable mirrors with a controllable radius of curvature, known as Suspended Active Matching Stages (SAMS), were installed as suspended relay optics to optimize mode-matching and throughput along the squeezing

injection and GW readout paths.

Along the squeezing injection path, three SAMS actuated by piezoelectric actuators (“P-SAMS” [155]) were installed: one PSAMS to optimize mode-matching from the squeezer cavity into the filter cavity, and two PSAMS to optimize squeezer mode-matching into the interferometer and output mode cleaner cavities. Along the interferometer readout path, a thermally-actuated SAMS (“T-SAMS” [156, 157]) was installed to control mode-matching from the outgoing interferometer beam through the output mode cleaner cavity.

Frequency-dependent squeezing. At each detector, a 300-m filter cavity was installed to enable frequency-dependent squeezing [21]. Frequency-independent phase-squeezed light is generated by the OPO, reflected from this detuned filter cavity to become frequency-dependent squeezing, and then injected into the interferometer. Reflection from the filter cavity imparts a frequency-dependent phase shift for frequencies within resonance, rotating phase squeezing at high frequencies into amplitude squeezing at low frequencies. Ideally, this frequency-dependent squeezing rotation will fully compensate the opto-mechanical squeezing rotation within the interferometer arm cavities. This optimum is achieved when the filter cavity half-width-half-maximum linewidth γ_{FC} and detuning δ_{FC} are both equal to the crossover frequency $\Omega_{\text{SQL}}/\sqrt{2}$ ($\gamma_{\text{FC}} = \delta_{\text{FC}} = \Omega_{\text{SQL}}/\sqrt{2}$), where quantum radiation pressure noise and quantum shot noise are equal [142]. This frequency Ω_{SQL} depends on the circulating laser power in the arm cavities.

The O4 filter cavities are designed to meet $\Omega_{\text{SQL}} = 2\pi \times 42$ Hz (by choosing an input mirror power transmissivity of $T_{\text{in}} \sim 900\text{-}1000$ ppm with assumed losses of 60 ppm) to compensate a circulating laser power of 500 kW [142, 158]. However, in O4 the detectors operate with arm powers of 300-400 kW, which is less than the 500 kW that the filter cavity is designed to compensate. Higher circulating powers require more broadband filter cavities to fully compensate back-action noise, as higher powers couple in quantum radiation pressure noise over a broader bandwidth. Without increasing the arm power, an optimal filter cavity would have a narrower linewidth, for example achieved by choosing a less transmissive input coupler. As the arm powers increase to 500 kW, the low-frequency squeezing rotation is expected to reach the designed performance.

7.4.4 Auxiliary controls upgrades

Beam spot control with camera servo

In O3, an alignment dither system was used to control the beam spot positions on the mirrors in the arm cavities. A sine wave was injected into the angular motions of the mirrors in the arm cavities, and then a feedback signal was derived by demodulating the differential arm length signal at the dither frequency. This dither system was used to control the spot positions on the end test masses and one input test mass by feeding back to the angular motions of test masses and power recycling mirror. However, these dither lines were injected around 20 Hz, causing excess noise in the gravitational wave band.

To avoid the use of dither lines in the gravitational-wave band, in O4 the spot position control was moved to a camera servo system. Cameras monitor the beam spots on the mirror, and the centroid of each image is calculated to monitor the position of the beam over time. At both observatories, the beam spot positions on the end test masses and the beamsplitter are monitored by cameras and fed back to control the angular motions of the test masses and the power-recycling mirror. The alignment dither system is still used for lock acquisition at the Hanford observatory, but the control is transitioned to camera servos before entering science mode.

At the Hanford Observatory, it is necessary to mis-center the beam spots on several arm cavity optics to avoid the remaining point absorbers on both end test masses (ETMX, ETMY) and on the input test mass to the X-arm (ITMX). The point absorbers are significant enough that the beams are miscentered by several millimeters. The Livingston Observatory no longer has significant point defects, so the beams are kept near the center of the mirrors.

Local damping upgrades

The gravitational wave strain was found to be contaminated by re-injected sensor noise from the local controls of the triple suspensions that control auxiliary mirrors, such as the mirrors of the power- and signal- recycling cavities. Noise from the suspension sensors coupled into the strain sensitivity through the auxiliary length and alignment control loops. The re-injected sensor noise from the triple suspensions was found to be even more significant than from the quadruple suspensions controlling the test masses. Other singly-suspended mirrors, such as mirrors in the input telescope, also contributed damping noise to the auxiliary controls. For O4, local damping control of auxiliary mirrors and test masses was recommissioned

to suppress re-injected sensor noise within the gravitational wave band.

At the Hanford Observatory, the local damping noise for the mirrors on the input telescopes coupled into the sensors used to control the common alignment degrees of freedom. Due to the alignment control coupling to gravitational wave strain, this impacted the strain noise from 20-30 Hz. To improve this noise coupling, the local controls for the input mirror length, pitch and yaw degrees of freedom were reduced to bandwidths 1 Hz. This corresponded to a reduction in noise re-injection of an order of magnitude or more above 10 Hz in the alignment controls, and a direct reduction of the gravitational strain noise by a factor of 2 between 20 and 30 Hz.

Local damping noise from the large triple suspensions that control the power-recycling and signal-recycling mirror 3 (PR3 and SR3) also contaminated the gravitational wave strain. Adjustments to the controls maintained the same suppression of the suspension resonances while achieving an order of magnitude reduction in the noise injection above a few Hz for all degrees of freedom. These improvements reduced noise in the alignment sensors, and also contributed to a reduction in the differential arm length control RMS.

The small triple suspensions controlling the power- and signal- recycling mirrors and power- and signal- recycling mirror 2 (PRM, SRM, PR2, SR2) were also recommissioned at LHO. These damping controls had the least effect on the gravitational wave strain noise, but contributed significantly to the noise in the auxiliary length control of the power-recycling, signal-recycling, and Michelson cavities. By reducing the damping control above a few Hz, both the RMS of the control loops and re-injected noise above 10 Hz were reduced.

The Livingston Observatory performed a similar recommissioning of the local controls for the large and small triple suspensions, beamsplitter suspension, and quadruple suspensions. Noise from these controls limited the gravitational wave strain noise from 10-20 Hz. By reducing loop gains, improving cutoffs, and reshaping the loops to suppress noise above 10 Hz, noise from the local controls is five times below the gravitational wave strain noise above 15 Hz. However, below 15 Hz, there is still significant noise contribution from these controls. A sub-budget of the local control noise contributions at LLO is shown in Figure 92.

The Hanford Observatory does not observe significant contribution from the quadruple suspensions or beamsplitter suspension to the gravitational wave strain noise. Even though both observatories are not directly limited by local damping noise within the gravitational wave band, there is significant contribution to the low frequency noise and RMS of the differential arm length control, as well as auxiliary alignment and length controls from local

damping noise.

Symmetric port feedforward

Wavefront sensors mounted on an in-vacuum table at the symmetric port of both LLO and LHO witness excess noise from 10 - 30 Hz due to ground motion coupling. Feedforward subtraction is applied to reduce this noise coupling using accelerometers mounted on the table that witness the excess motion. The feedforward is applied to subtract noise from all control loops sensed at the symmetric port, including the control of the common hard mode of the arm cavities, power-recycling cavity alignment, and the input pointing. The feedforward can suppress the excess noise in alignment controls by up to a factor of ten from 10 - 30 Hz.

7.4.5 Online noise subtraction

In O4, up to three calibrated strain channels from each detector are available: calibrated strain, calibrated strain with line subtraction, and calibrated strain with line subtraction and online cleaning¹. Figure 84 shows the online cleaned range for both detectors.

Both detectors provide the calibrated strain with low-latency line subtraction, where the sinusoidal calibration lines and the 60 Hz power lines are removed from the online data [159]. The line subtraction provided over two Mpc of additional BNS range for both detectors.

The LIGO Hanford Observatory, in addition to low-latency line subtraction, implemented online cleaning of the strain data to subtract residual noise couplings during O4. This data cleaning is performed using the NonStationary Estimation and Noise Subtraction code (NonSENS) [127, 128, 129]. NonSENS has previously been demonstrated to successfully subtract both linear stationary couplings and nonlinear modulated couplings in the offline data from O3 [160]. For O4, front-end infrastructure was developed to perform cleaning of the strain noise data in low latency using measured noise couplings.

The greatest benefit of this online cleaning is in subtracting laser noises such as beam jitter, described in Section 7.3.3, and frequency noise, described in Section 7.3.3. The noise reduction with online cleaning could yield up to 3 Mpc of additional BNS range for LHO when running, due to its mitigation of nonstationary laser noise couplings. The LIGO

¹These channels are named `GDS-CALIB-STRAIN`, `GDS-CALIB-STRAIN_NOLINES`, and `GDS-CALIB-STRAIN_CLEAN`, respectively

Livingston Observatory did not implement online cleaning for O4, and so the cleaned and line-subtracted strain data are equivalent.

7.4.6 Scattered light mitigation

In the commissioning period before O4, vacuum incursions were made to mitigate known sources of scattered light at both observatories. As scatter sources are removed, or the motion of scattering surfaces is damped, in each case the acoustic and vibrational couplings in each area are measured to decrease, sometimes dramatically.

Septum window removal

At both interferometers, a septum window between two vacuum chambers at the output port was removed. This window separated the volumes housing signal-recycling mirrors and the output mode cleaner, and the outgoing interferometer light transmits through this window to reach the output mode cleaner.

Scattered light from this window coupled excess acoustic noise into the DARM spectrum. Removal of the septum window at both observatories significantly reduced vibration noise coupling at the interferometer output, in some frequency bands by up to two orders of magnitude. The reduced vibration noise coupling decreased the total detector noise from 30 - 100 Hz at LHO and 15 - 25 Hz at LLO. This window removal increased the astrophysical range at both detectors by a few Mpc.

Baffle damping

At both sites in O3, increased vibration levels, such as from wind and anthropogenic noise, produced harmonic series of noise peaks in DARM that are characteristic of scattered light noise. The scattering noise sources were identified using targeted vibration injections and movies of scattered light modulation in the test mass chambers. The sources were baffles, the cryopump manifold baffles and arm cavity baffles, that were designed to meet scattered light requirements at low motion amplitudes, but whose motion was amplified by high quality-factor resonances. This excess motion produced harmonic series of the resonance frequencies that reached into the sensitive frequency band of DARM.

For O4, most of these baffles were damped or their suspensions improved in order to reduce the velocity of reflecting surfaces, though not the amplitude of the reflected light.

7.4.7 Parametric instability mitigation

Parametric instabilities (PI) are caused by the optomechanical interaction between the circulating laser powers and the bulk mechanical resonance of test mass mirrors [161]. In particular, the radiation pressure of higher-order transverse optical modes in arm cavities can excite the mechanical modes of test masses, which can further excite the higher-order optical modes in the arm cavities. This interaction can become a positive feedback loop and cause an instability in high-power interferometer operation. Such PI modes have been observed at both LIGO interferometers [162].

In the previous runs, both active and passive damping have been used to mitigate such instabilities. For active damping, the test mass electrostatic drive has been used [163]. For passive damping, acoustic mode dampers (AMDs) were installed on the test masses before O3 [148]. The target frequencies of the AMDs is 15-80 kHz. However, during the arm power increase in the commissioning for O4, a PI mode at 10.4 kHz was observed at LHO, and a PI at 80 kHz was observed at both LHO and LLO. Both of these PI modes are outside of the AMDs target frequencies.

At LHO, the 10.4 kHz PI mode is driven by circulating laser power scattered into the second higher-order transverse mode due to misalignment, and is observed at twice the arm cavity transverse mode spacing (~ 5.2 kHz [44]), i.e. 2×5.2 kHz = 10.4 kHz. At LHO, this PI appeared on ETMY and is caused by the intentional beam miscentering on ETMY. To mitigate this instability, active damping using the electrostatic drive to ETMY was used at LHO during O4.

At both detectors, the 80 kHz PI is likely driven by the circulating laser power in a first higher-order transverse mode that 2 arm cavity free spectral ranges (FSR ~ 37.5 kHz) away from the main carrier frequency, i.e. 2×37.5 kHz + 5.2 kHz ~ 80 kHz. Thanks to the fast 512 kHz analog-to-digital converters installed for O4, the fast arm power modulations produced by the 80 kHz PI can be digitally sensed in transmission of the output mode cleaner. A control signal was generated and sent to coil drivers for the problematic test mass, to attempt electrostatic damping of the 80 kHz PI. In LHO, the 80.3 kHz PI appeared on ETMX at 57 W input laser power, and despite attempts at electrostatic damping, was ultimately mitigated by tuning the ETMX ring heater to offset the optical mode from the mechanical mode. In LLO, the 80.4 kHz PI appeared on ITMY at 50 W input laser power, and was correspondingly mitigated by tuning the ITMY and ETMY ring heaters.

7.4.8 Output photodetector electronics upgrades

Since transimpedance amplifiers are limited in range, the maximum signal to noise ratio of the photodetectors used to sense the gravitational wave signal is set by the voltage across the sensing resistor. The addition of an inductor to the transimpedance network creates a frequency dependent conversion from photocurrent to voltage, allowing for better signal to noise ratio at audio band frequencies and higher dynamic range, as was demonstrated at GEO600 [164]. The transimpedance network was redesigned for O4 to include a 2.5 H inductor with mu metal shielding and two capacitors to tune the frequency response, such that the main sensing resistor only sees 1% of the photocurrent. The speed of the ADC was also increased, and now reads out at 512 kHz.

7.4.9 Grounding improvements

During the commissioning period, spectra of the variation in current flowing from building electronics ground to neutral earth were observed to be correlated with noise in DARM. A newly developed electronics ground injection system showed that large spectral features in DARM could be produced by injecting 100 milliamp-scale currents onto the building electronics ground. The coupling is thought to be produced by fluctuations in the potential of the electronics ground system due to the variations in current flows across the finite resistance between electronics ground and the neutral earth around the building (measured to be about 2 ohms at LHO). Forces on the charged test mass may fluctuate with the potentials of nearby electronic systems that are referenced to the fluctuating electronics ground, such as the electrostatic drive and ring heaters.

The noise from electronics ground potential fluctuations was reduced in two ways. First, the resistance between certain electronics chassis and the building electronics ground were reduced in order to reduce the total resistance to neutral earth for those electronics. Second, the biases of the electrostatic drives were swept and set to values that minimized the coupling to DARM of current injections onto the electronics ground. It is hypothesized that, at the coupling minimum, the force on the test mass due to its charge is partially cancelled by the force due to the bias-induced polarization. These mitigations resulted in range improvements of a few Mpc. Further mitigation could likely be obtained by shielding electronics inside the chamber from the test mass with shields connected to the chamber.

7.5 Summary of High Power Operation

7.5.1 Arm power characterization




Figure 93: Comparison of the total detector noise in O3 and O4 without squeezing. The unsqueezed shot-noise-limited sensitivities are primarily a product of the circulating laser power and the detector’s readout losses. At LHO, around 900 Hz, the unsqueezed strain sensitivity increased by about 40% in O4. At LLO, around 900 Hz, the unsqueezed strain sensitivity increased by about 16%.

Quantifying the circulating power in the Fabry-Pérot arm cavities is important to understand the interferometer response to gravitational waves. As described in [44], this quantity is not trivial to constrain due to large uncertainties in estimates of the incident power on the beamsplitter and arm cavity gain. This section describes several methods that have been used to quantify arm power, and the associated challenges with each method.

One of the simplest ways estimate arm power uses the input power and power recycling cavity finesse. Specifically, this considers the input power incident on the power-recycling cavity (P_{in}), the power gain from the finesse of the power-recycling cavity (i.e. power-recycling gain, or PRG, G_{PR}), and the arm cavity gain (G_{arm}) given input and end cavity mirror transmissivities along with round-trip optical losses. The circulating arm powers can

be then calculated as

$$P_{\text{arm}} = \frac{1}{2} P_{\text{in}} G_{\text{PR}} G_{\text{arm}}. \quad (7.157)$$

This method was previously applied to estimate the arm cavity power in the third observing run [44]. The main source of uncertainty in this estimate stems from calibration of the photodetectors used to estimate the input power, power in the PRC, and arm cavity gain. This estimate is made more accurate by further estimating how much of the input power, P_{in} is carrier power. For both LHO and LLO, a few percent of the input power is measured to be higher order mode and RF sideband power. This value is constrained by changing the modulation depth of the RF sidebands and measuring the changing in power at the output port.

Table 10: Summary of the operating power parameters of the LIGO interferometers during O4. Circulating powers estimated using Eq. 7.157.

Parameter	Symbol	LHO	LHO start	LLO	Units
Input laser power	-	57.5	71.8	64.5	W
Input carrier power	P_{in}	55.9	68.6	61.9	W
Arm cavity power	P_{arm}	364	407	327	kW
Power-recycling gain	G_{PR}	50.1	45.6	39.85	W/W

LHO Circulating arm power

Previously in O3, the arm power at LHO was estimated to be around 195 ± 14 kW, as reported in Table II of Ref. [44]. The replacement of the Y-arm input test mass at LHO (see Section 7.4.2) corresponded with a significant improvement in the arm losses, and a corresponding increase in the recycling gain at high power. The resulting power-recycling gain at input power 57.5 W is measured at 50.1 W/W. At this operating power, 97.2 % of the input laser is power is carrier power, and the rest is higher order mode and RF sideband power. Applying the measured arm cavity gain of 260, the resulting arm power estimate is $364 \text{ kW} \pm 18 \text{ kW}$ following Equation 7.157.

With an increase to 71.8 W input power, optical losses increased slightly due to increased thermal distortion contributing to higher mode mismatch within the interferometer. This mode mismatch reduces the estimated power-recycling gain to 45.6 W/W. An estimated 4.4% of the input power is higher order mode and RF sideband due to this excess mode

mismatch. Using the same arm cavity gain, the resulting estimated power is $407 \text{ kW} \pm 20 \text{ kW}$.

LLO Circulating arm power

Previously in O3, the arm power at LLO was estimated to be around $240 \pm 18 \text{ kW}$, as reported in Table II of Ref. [44]. The replacement of the end test masses at LLO (see Section 7.4.2) corresponded to a small decrease in the power-recycling gain due to an increase in the arm cavity losses from higher scatter and absorption. The resulting power-recycling gain at input power 64.5 W is calibrated to be 39.85 W/W . Within the input power, 4% is estimated to be higher order mode and RF sideband power. Applying the measured arm cavity gain of 265, the resulting arm power estimate is $327 \text{ kW} \pm 6.5 \text{ kW}$ following Equation 7.157.

7.5.2 High power challenges

Both the Hanford and Livingston detectors underwent an increase in operating power ahead of the fourth observing run, as a part of improving the shot-noise limited in sensitivity. However, the power increase at both observatories was met with several challenges related to optical loss, excess low frequency noise, and controls challenges. In particular, the Hanford Observatory began the fourth observing run operating at a higher power that was later reduced to improve both the duty cycle and sensitivity. Table 10 summarizes the relevant operating power parameters for both observatories. Both the current LHO parameters (LHO) and the LHO parameters for the first month of the run are presented (LHO start).

This increase in operating power was enabled by pre-stabilized laser upgrades detailed in Section 7.4.1. The test mass replacements at both sites, detailed in Section 7.4.2, also enabled power increase, as operating power was limited in O3 by the presence of point absorbers [44].

During the power-up procedure at LHO, saturations of the power recycling cavity pick-off detector (POP) prevented increasing the operating power above 55 W . This photodetector is the main sensor for the auxiliary length degrees of freedom: PRCL, MICH and SRCL. A vacuum incursion was performed to attenuate the beam along this pick-off path to avoid saturation of the photodetector. This procedure was performed at both observatories to allow power up efforts to proceed. This attenuated light along both the POP path and well as the reflection port path used to control the common arm length and the common

alignment controls.

Thermal compensation for high power

The thermal compensation system (TCS) is designed to sense and correct the thermal lenses and aberrations on the core optics that occur due to absorption [165]. At both detectors, the estimated uniform test mass absorption is around 0.5 ppm, dominated by the optical absorption of the mirror coatings. For the O4 operating powers, the mirrors absorb hundreds of milliWatts of power, resulting in an induced thermal lens that distorts the finely tuned cavity dynamics. This estimated uniform absorption is further impacted by any point defects that absorb further power. The effect of these point absorbers was discussed at length in Buikema et al. [44].

Two main actuators provide thermal compensation control to core optics: ring heaters placed around the barrel of each optic and CO₂ lasers incident on a compensation plate behind each input test mass. The ring heaters are designed to create a negative thermal lens on the test masses that correct aberrations on the high-reflective surface of the mirror. The CO₂ lasers can induce a positive or negative thermal lens on the compensation plate. Both LLO and LHO make use of the ring heaters on the input and end test masses. Only LHO makes use of the CO₂ lasers at this time.

Hanford Detector TCS Tuning After LHO achieved 57 W input power, end test mass ring heaters were adjusted in both common and differential steps to provide better arm mode matching and reduce contrast defect light measured at the output port. This had the further benefit of reducing frequency noise coupling present at LHO. By increasing the end test mass ring heater power in a common step, the frequency noise coupling was reduced by a factor of three and the contrast defect light by a factor of five. Due to this change, a parametric instability around 10.4 kHz is driven during the lock. Due to the success of this tuning, active damping of this PI was implemented. This is further detailed in Section 7.4.7.

Meanwhile, slight differential tuning of each end test mass ring heater improved contrast defect further, but drove a different parametric instability near 80.3 kHz. At the time, no active damping solution was found, so the ring heater setting was left in a common mode for the end test masses. Ring heater changes were attempted on the input test masses to further improve arm mode matching and improve detuning in the signal recycling cavity. However, neither common nor differential changes of the ITM ring heaters were successful. Some changes completely prevented locking of the interferometer.

After the input power increase to 71 W at LHO, similar tests were performed to adjust the end test mass ring heaters in both common and differential steps. However, all tests of the ITM ring heaters were unsuccessful at improving mode matching and severely affected the ability to perform locking procedures. The contrast defect increased slightly as well as the frequency noise coupling due to arm mismatch. When the power was reduced mid-run to 57 W, the ETM ring heaters were reverted to the previous settings that produce low frequency noise coupling.

The CO₂ heaters were also adjusted to reduce frequency noise and jitter noise present in the interferometer. An annular beam mask was used to closely mimic the effect of the ring heaters. Differential tuning was found to worsen noise coupling, while common heating further improved the noise coupling.

Livingston Detector TCS Tuning At LLO, both the ITM and ETM ring heaters are used to improve the thermal lens from the increase in operating power. Differential ITM ring heater actuation is used to reduce laser noise coupling. ETM differential changes are used to adjust the higher order mode spacing to avoid parametric instabilities.

Increasing the input power at the Livingston Observatory did not result in a proportional increase in sensitivity due to loss of mode matching from the increased power. Increasing the ring heater power of both the input and end test mass ring heaters recovered some of the optical gain lost with each power increase, but not all. Furthermore, increasing ITM ring heater power also impacted locking procedures. These challenges limited the power increase at LLO further beyond 63 W input.

Thermalization

Both detectors experienced challenges related to *thermalization*: transients in the operating point of the detector due to the evolution of the induced thermal lens from the heating of the interferometer beam. Previous demonstrations of dynamic thermal compensation demonstrate that this thermal evolution can last for several hours at the beginning of each lock as the test masses absorb hundreds of mW of power [166]. The thermal compensation system is designed to counteract induced thermal lensing, but is less effective at dynamically counteracting the changing thermal lens that occurs at the beginning of each lock sequence.

During the thermalization process, optical losses of the main beam and RF sidebands change dramatically, higher order mode frequencies shift, and noise couplings change as the mirrors absorb power. The effects of thermalization can limit the operating power of both

detectors; it is challenging to compensate the thermal lens generally, and more challenging to compensate the early thermal evolution each lock.

Operating the Hanford Observatory with 71.8 W input power during the early part of O4a came with significant controls challenges due to thermalization. In particular, while the carrier beam power increased during the thermalization period, the sideband gain decreased dramatically. The 9 MHz sideband used to control the power-recycling cavity length lost up to 75% of its optical gain during the first few hours of every lock. This dramatic effect on the sideband but not the carrier is exacerbated by the induced ITM lens. This effect is doubled in the sidebands due to anti-resonance in the arm cavities. A servo to maintain a stable unity gain frequency was commissioned to manage this loss, but the gain evolution was unpredictable in each lock, as the overall thermal state of the interferometer varies significantly from lock to lock depending on how quickly the interferometer can relock after lockloss and maintain the high power thermal state. This corresponded with regular fast locklosses within the first few hours of observing that impacted the duty cycle.

Furthermore, the significant evolution of the sidebands impacted the ability control noise couplings consistently during each lock. With significant evolution in the sideband gain used to control the auxiliary length degrees of freedom, tuning length feedforward, as described in Section 7.3.4, became a challenge. A feedforward scheme measured at the start of a lock would lose efficacy during the thermalization process, increasing noise at low frequency during observing periods. LHO also witnessed a significant increase in vibrational noise at higher power. The effect of high power operation on the low frequency sensitivity is demonstrated in Figure 94. The increase in operating power corresponds to an increase in the broadband noise from 20-50 Hz. Some of this noise can be explained by excess noise in the length control loops that is not well-mitigated by length feedforward. However, even with a subtraction of the noise associated with the length controls, there is an approximate doubling of the broadband noise at low frequency. Most of this noise can be attributed to vibrational noise causing excess scattered light and beam jitter.

While the most significant impacts to the sensitivity of LHO at 71 W result from an increase in broadband low frequency noise, thermalization also impacted squeezing losses. An increase in the detuning of the signal-recycling cavity, combined with a reduction in mode matching of the interferometer to the output port The signal-recycling cavity would become increasingly detuned during the lock period, impacting the optical spring in the differential arm length and compromising squeezing.

After the first month of observing, the input power of LHO was decreased from 71 W to 57 W. This corresponded to an improvement in both duty cycle and detector sensitivity. While operation at the lower power still includes a thermal transient within the first few hours of the lock, the effects are significantly reduced, enabling improvement of the low frequency noise and reducing losses that impacted the injected squeezing. The decrease in power also corresponded to an increase in the stability of the auxiliary controls, enabling a significant improvement in the duty cycle. The relatively low duty cycle of LHO during O4a can mostly be attributed to the impacts of high power operation early in the run. Also, the steady increase in the LHO sensitivity, as shown in Figure 84, is a result of this change in power.

7.6 Future work

Between O4a and O4b, a 3-month commissioning break from January 16, 2024 to April 10, 2024 allowed the observatories to install several upgrades. At LLO, both X- and Y- end test mass mirrors (ETMX, ETMY) were cleaned with First Contact *in-situ* while suspended. At LHO, the output mode cleaner cavity was replaced to reduce optical losses; baffles near the input mode cleaner were further damped to reduce scattered light; and the squeezer's piezo-deformable optics (PSAMS) were adjusted to allow further optimization of squeezer mode-matching throughout the interferometer. At both detectors, a squeezer pre-mode cleaner cavity was installed to reduce high frequency (\gtrsim MHz) squeezer laser intensity noise.

After O4, a more extensive two year-long commissioning break will precede the fifth observing run (O5). Major planned detector upgrades for O5 include the test mass mirrors with reduced coating thermal noise and an upgraded thermal compensation system to achieve high power.




Figure 94: A trend of the differential arm length noise level from 20-50 Hz, normalized to data taken on March 30, 2023. The noise in this band increased with the increase in power. Even with coherent subtraction of limiting controls noise, the overall noise increased. This excess noise is attributed to an increase in coupled vibrational noise from scattered light and beam jitter.

Chapter 8

Technical Noise and Data Quality Considerations for Future Detectors

8.1 Introduction

This chapter covers work from a paper submitted to *Classical and Quantum Gravity* by Capote, Dartez, and Davis [167].

Advances in gravitational-wave detector technology have pushed the limits of possibility for the detection of new and exciting gravitational wave events [2, 168, 140, 169, 170, 171]. Despite many advances, ground-based gravitational-wave detectors such as Advanced LIGO [9], Advanced Virgo [14], and KAGRA [15] continue to be limited in sensitivity and operation due to both fundamental limits of gravitational-wave detectors and technical challenges in their operation [44, 172, 173, 174].

Several next-generation (XG) gravitational-wave detectors have been proposed, such as Cosmic Explorer [6], planned to be located in the United States, and the Einstein Telescope [30], planned to be located in Europe. Designs for these new detectors are based on current understanding of the best technologies for gravitational-wave detector performance. To enhance these new detectors, several key upgrades are proposed that push the limits of current technology to further improve detector sensitivity.

Noise sources that limit these types of gravitational-wave detectors are usually divided into a few main categories. Fundamental noise sources are derived from first principles and set an upper limit to any detector sensitivity. Baseline designs for XG detectors tackle many of these fundamental sources first by lengthening the Fabry-Pérot cavity arms and increasing

the mass and size of the test masses. Additionally, these designs include enhanced seismic isolation requirements through longer and larger suspensions, improved thermal noise via lower loss optical coatings, and reduction in quantum noise through high optical power and high levels of quantum squeezing [6, 30, 59]. Meanwhile, technical noises result from the operation of a detector and can usually be reduced if studied and understood. While baseline design choices focus on fundamental noise reduction, some design factors also tackle technical noise challenges. Some work has already been done to consider the design of the Cosmic Explorer detector with knowledge of fundamental, technical, and environmental limiting noises at low frequency, as in Ref. [31].

Technical noises are typically considered to be stationary but can also be non-stationary over long time periods or even create short bursts of transient noise called “glitches” [175, 176, 177, 172, 174]. Technical noise effects that are non-stationary or transient are often considered separately and are referred to as “data quality issues.” While the origin of these time-varying technical noises is similar to other sources of technical noise, they can create additional challenges for data analysis pipelines compared to stationary noise [178, 179, 180, 177, 181, 182, 183, 184, 185, 186, 187]. For this reason, understanding and mitigating data quality issues is an active area of research and vitally important to improving the robustness of astrophysical analyses with gravitational waves [176, 188, 179].

Addressing these challenges will require active research that builds upon current work being done to design XG detectors. In this work, we focus on Cosmic Explorer as an example of an XG gravitational-wave detector to explore the impacts of technical noise, data quality, and calibration on XG gravitational-wave science. We also compare and contrast how these issues may affect the Einstein Telescope when appropriate. Cosmic Explorer is planned to be a dual-recycled Michelson interferometer with 40 km or 20 km Fabry-Pérot arms [6]. The test masses serving as the arm cavity mirrors are planned to be 320 kg [6], eight times the mass of the current Advanced LIGO test masses [9]. These and other design choices, such as improved seismic isolation, will allow Cosmic Explorer to achieve unprecedented sensitivity to gravitational waves. Due to the focus on Cosmic Explorer, the projections included in this work use previous results from the Advanced LIGO (aLIGO) detectors, as these are the current-generation detectors most similar to the proposed Cosmic Explorer design [6].

First, we consider technical noises that create broadband, stationary noise in the detector, and possible mitigation strategies (Section 8.2). We then discuss the challenges of narrow-band technical noises (Section 8.3). Finally, we comment on the outlook of addressing these

challenges to best realize XG science goals (Section 8.4).

8.2 Broadband Technical Noises

To begin the discussion of technical noises and their impacts, we focus on technical noises that impact the detector over a broad range of frequencies. In current generation detectors, one of the most significant sources of technical noise, especially at low frequency, comes from auxiliary controls such as length and alignment controls [189, 44, 172, 190]. Control servos are a necessary part of sensitive interferometer operation to achieve stable resonance conditions in optical cavities, ensure good optical gain, and maintain operation in the linear control regime. However, controls can also inject excess noise into the gravitational wave band through various means. Part of the operation of a gravitational-wave interferometer involves finding a delicate balance between an optimal control point and minimal noise as these are two situations that can be at odds [44]. In this discussion we will enumerate controls requirements as they are in current generation detectors, and what considerations are required for next generation science goals.

8.2.1 Auxiliary Length Control Noise

Auxiliary cavity length control is an essential part of stable interferometer operation. Assuming Cosmic Explorer continues to follow a similar topology as current advanced gravitational-wave detectors, these auxiliary lengths are the power recycling cavity (PRC) length, the Michelson cavity (MICH) length, and signal extraction cavity (SEC) length [31]. Auxiliary lengths couple marginally to the differential arm length (DARM) and, therefore, noise related to controlling these cavities couples into the DARM spectrum from which the gravitational-wave readout is derived [44]. The significant noise sources in aLIGO are from the MICH length and the SEC¹ length [44]. Advanced LIGO does not experience any significant noise contribution from the PRC length, and this is so far expected to be the same for Cosmic Explorer. In principle, no noise will couple from the SEC in Cosmic Explorer due to the balanced homodyne readout [31], however, in practice effects like contrast between each arm cavity mode could result in SEC noise coupling. For the purpose of this discussion, we assume that the main auxiliary noise coupling results from MICH.

¹In many aLIGO papers this is referred to as the “signal recycling cavity”, or “SRC”.

Michelson cavity motion creates differential phase sidebands that become amplitude sidebands at the output of the interferometer. However, this noise is suppressed in the interferometer via the arm cavity finesse, as the DARM signal is amplified by the Fabry-Pérot arm gain relative to the MICH length signal. Current design considerations of Cosmic Explorer assume the same finesse as aLIGO [6]. Therefore, the Michelson coupling will have approximately the same factor, $\pi/2\mathcal{F} = 1/280$. However, studies are underway to consider the benefits of changing the Cosmic Explorer finesse; see Ref. [191] for a more in-depth discussion.

To understand the effects of noise coupled from Michelson length control, it is beneficial to understand what noises limit the Michelson loop itself. In Advanced LIGO, the Michelson loop was shown to be limited by sensing noise roughly above 10 Hz, at about 10^{-16} m/ $\sqrt{\text{Hz}}$. This sensing noise is shown in Fig. 9a of Ref. [189] to result from the quantum (shot) noise on the photodetector used to detect the MICH length signal. Below 10 Hz, MICH is limited by suspension damping noise from the local control of the beamsplitter, the mirror used to control Michelson length [189]. Besides suspension damping noise, there is also seismic noise, roughly following a $1/f^6$ slope and reaches a level of 1×10^{-17} m/ $\sqrt{\text{Hz}}$ at 10 Hz. The slope of this noise results from the triple suspension used to suspend the beamsplitter. All of these noises couple to DARM and limit the gravitational-wave channel, though for the aLIGO band which extends only to 10 Hz, the most significant noise is MICH sensing noise.

To estimate of the level of MICH noise in a gravitational-wave detector, we can take the current level of MICH sensing noise above 10 Hz, 10^{-16} m/ $\sqrt{\text{Hz}}$, and multiply it with the assumed coupling factor, $1/280$. The resulting displacement noise is approximately 3×10^{-19} m/ $\sqrt{\text{Hz}}$. This noise level is significant, and completely limits Advanced LIGO if left unmitigated. However, a feedforward control is applied that suppresses this noise out of band. Feedforward is a form of real-time noise cancellation based on a known transfer function of MICH to DARM coupling [9, 192]. Currently, it is possible to achieve at least a factor of one hundred suppression of this sensing noise in the gravitational-wave band. Such a suppression factor is dependent on the measurement—a measurement precision of 1% is required.

One drawback of feedforward is that it can be tuned to suppress sensing noise or displacement noise, but not both. Given the high level of MICH sensing noise that can couple to the gravitational-wave band, feedforward subtraction must be tuned to reduce sensing noise, but will also, by the nature of feedforward, amplify any displacement noise present. In

Advanced LIGO, the bandwidth of the MICH loop is set to around 10 Hz. There are several considerations for the setpoint of the MICH bandwidth, but the one most important for this discussion is the suppression of displacement noise. In order to maintain the linear control regime required for detector operation, the residual displacement of MICH length must be around a few picometers [189, 44]. This bandwidth is high enough to achieve the required displacement noise residual for aLIGO and any amplified displacement noise that occurs is below the observation band for gravitational waves.

However, the Cosmic Explorer band will extend down to 5 Hz, so the effect of MICH displacement noise must also be considered. Part of the design for Cosmic Explorer involves better local control for suspensions, namely, making use of local sensors with much lower noise [31, 6]. The proposed Homodyne Quadrature Interferometer (HoQI), is one such sensor that has the required performance for a detector like Cosmic Explorer [193]. The benefit of using the HoQI would be much less suspension damping noise, such that the low frequency displacement noise that limits MICH would likely be limited by seismic noise. To further improve upon displacement noise, Cosmic Explorer design also considers improved seismic isolation via improved suspension design, such that Cosmic Explorer could achieve a factor of ten reduction in seismic noise at 1 Hz [31]. These improvements allow for the ability to relax the controls design for MICH. Feedforward will still be required for Cosmic Explorer: in a 40 km detector, a MICH loop with sensing noise of 10^{-16} m/ $\sqrt{\text{Hz}}$ and an arm finesse of 450 will result in a strain noise of about 6×10^{-24} 1/ $\sqrt{\text{Hz}}$ [31]. If a one hundred times suppression of the MICH sensing noise is achieved, this noise would be at least five times below the level of the Cosmic Explorer fundamental noise curve, which reaches 3×10^{-25} 1/ $\sqrt{\text{Hz}}$ at 20 Hz [31, 44]. However, any residual displacement noise below 10 Hz will be amplified by this feedforward, which could completely contaminate the detection band from 5-10 Hz.

By suppressing MICH displacement noise with better local controls and seismic isolation, the Cosmic Explorer MICH bandwidth could be set below 5 Hz. With less MICH loop gain above 5 Hz, less sensing noise will be coupled in the band, and less displacement noise will be amplified by the feedforward. This has two beneficial impacts: more sensitivity in the region from 5-10 Hz, and less stringent requirements on the feedforward suppression. To completely avoid any risk of displacement noise contaminating the gravitational-wave strain at 5 Hz, a MICH bandwidth that is closer to 3 Hz would be the most beneficial. Loop bandwidth cannot be lowered indefinitely, both for displacement noise considerations, and for control

of any suspension resonances. It will be key to keep the beamsplitter suspension length resonances low enough that such a loop design is possible for Cosmic Explorer.

Assuming no change in the sensing noise level, that means the seismic noise needs to be less than $10^{-16} \text{ m}/\sqrt{\text{Hz}}$ below 5 Hz. Seismic noise reaches approximately $5 \times 10^{-15} \text{ m}/\sqrt{\text{Hz}}$ at 5 Hz [189] in aLIGO. Switching to an improved suspension design (see [31]), and changing the beamsplitter suspension from triple in aLIGO to a quadruple suspension to gain an additional additional $1/f^2$ suppression, can reduce the seismic noise at 5 Hz below the sensing noise as required. This is one path forward to keep seismic noise low enough that any amplification of displacement noise from the feedforward would be kept out of the Cosmic Explorer band. Overall, even with little change in the MICH sensing noise, design strategies related to reducing displacement noise via improved local damping and improved seismic isolation will have a significant effect on suppressing any length controls noise that can contaminate above 5 Hz.

These projections are designed to be only a first look at potential auxiliary length coupling in XG detectors. As the designs for Cosmic Explorer and other XG detectors develop, simulation can further investigate how this noise may appear in an XG detector. Elements such as seismic isolation, suspension design, and optical layout could affect the appearance of Michelson noise or even other auxiliary lengths. Notably missing here is a full corner optical layout for the Cosmic Explorer interferometer, showing the planned lengths of the auxiliary cavities, and the mirrors and telescopes required. Development of such a layout is still underway as a part of the Cosmic Explorer preliminary design work, and better modeling of auxiliary length noise can be undertaken once this layout is proposed. Furthermore, any changes to the Fabry-Pérot cavity finesse will adjust the MICH coupling; increasing the finesse will be beneficial in reducing MICH noise while lowering the finesse benefits SEC losses that limit high-frequency sensitivity [191]. Even in current detectors, there is still much to be studied about auxiliary length coupling. For example, the projections here only consider a linear coupling of MICH to DARM, but some evidence from O3 indicates that auxiliary length coupling could have a nonlinear component [44]. Nonetheless, this initial work highlights the importance of studying this type of technical noise as a part of maximizing XG detector sensitivity.

8.2.2 Alignment Sensing and Control Noise

Alignment sensing and control, or ASC, is arguably one of the most topologically complex controls system in ground-based interferometers and is often a limiting noise at low frequency [44]. The alignment of the Fabry-Pérot arm cavities and the auxiliary cavities must be controlled together: one nanoradian residual angular motion is required for the arms, while tens of nanoradians of residual motion is acceptable for auxiliary cavities [48]. These requirements are set to stabilize resonant conditions, suppress higher order spatial modes that can cause optical loss, and reduce cross coupling between degrees of freedom.

In aLIGO, alignment degrees of freedom are sensed via quadrant photodiodes (QPDs) and wavefront sensors (WFS) [48]. QPDs provide pointing control for while WFS are interferometric sensors using radio frequency quadrant photodetectors to sense the relative position of the carrier and sidebands within the cavities [44]. A third type of alignment control, called dither control, will be briefly discussed in Section 8.3. Beams are picked off at different ports of the interferometer to provide signals for these sensors. It is likely similar alignment sensing schemes and sensor layouts will be applied in Cosmic Explorer. Given the residual angular motion requirements, the alignment controls that provide the most significant noise impacts are the arm controls—achieving one nanoradian RMS requires loop bandwidths around a few Hz. Control of auxiliary cavity alignments, such as the PRC, MICH, and SEC alignments, are important, but the bandwidth required for these loops in Advanced LIGO is less than 1 Hz to achieve the RMS requirement. Therefore, this paper will focus mainly on the control of the arms when considering the technical noise.

The Fabry-Pérot arm control is performed in the cavity basis, known as “hard” and “soft” modes (the reason for these names will be explained shortly). A hard mode results from a shift in the cavity axis alignment, while the soft mode results in a translation of the beam position within the cavity. The hard mode is typically detected on a wavefront sensor (WFS), which detects the first-order misalignment mode from a cavity axis shift, while the soft mode is sensed at a quadrant photodiode (QPD) after the beam reduction telescope. In Advanced LIGO, the highest bandwidth control in alignment sensing is from the hard loops. Above 10 Hz, the limiting noise source for both WFS and QPDs is the sensing noise of the detector, about 10^{-14} rad/ $\sqrt{\text{Hz}}$ [48]. However, in aLIGO, some of these sensors are contaminated with excess noise from 10-30 Hz, due to spurious vertical and rotational coupling of the motion of the in-vacuum table where the sensors are mounted [44]. This

particular excess noise contamination is the result of poor seismic isolation of the table itself and can elevate the sensor noise level in the WFS, magnifying the injected noise from ASC into the DARM spectrum. While there have been some successful mitigation strategies using feedforward subtraction or signal blending in aLIGO [44], next-generation detector design should consider table isolation of similar importance to mirror isolation when it comes to reducing technical noise from alignment controls.

The coupling of angular mirror motion to cavity length is best understood via a geometric argument as

$$\Delta L = d_{spot} \times \theta_{mirror}, \quad (8.158)$$

where d_{spot} represents the beam spot motion on the mirror, and θ_{mirror} represents the mirror's residual angular motion [48]. This is a bilinear process, as both beam spot motion and mirror angular motion can have a separate frequency dependence. By achieving the nanoradian RMS requirement in the arm cavities, for example, the effect of the beam spot motion that is convolved with the residual angular motion is suppressed and has minimal noise coupling. In this case, the remaining noise coupling is linear, and results from residual angular motion coupling to a static beam offset [48]. The coupling of the arm cavity mirrors to the differential arm length, DARM, is the strongest, as any change in the angle of an arm cavity mirror is a direct change in DARM. The auxiliary cavity mirrors also experience a coupling to DARM, usually with a coupling strength of two or more orders of magnitude below the coupling strength of the arm cavity mirrors [48]. This is further motivation to consider only the noise coupling of the arm degrees of freedom as the dominant noise source from ASC.

Assuming a similar ASC sensing noise level as aLIGO, 10^{-14} rad/ $\sqrt{\text{Hz}}$, and a similar residual beam spot motion of 0.1 mm [48], the estimated strain noise from ASC at 20 Hz in Cosmic Explorer would be about 2.5×10^{-23} 1/ $\sqrt{\text{Hz}}$. In order to suppress this noise in the gravitational-wave band, low pass filtering can be applied to the alignment control loop. Considering the arm alignment loop bandwidth in aLIGO [44], two orders of magnitude loop suppression is achievable by 20 Hz, which would place the resulting strain noise from ASC just below the Cosmic Explorer design sensitivity at 20 Hz. However, this is not sufficient to achieve the desired Cosmic Explorer sensitivity at 5 Hz. If the alignment control loop bandwidth can be reduced, the sensing noise can be suppressed more in the detection band. This requires a reduction of the angular displacement noise, such that less control is required to meet the same residual motion requirements. Loop bandwidth reduction will also depend on the suspension design, as the suspension dynamics play a significant part of the design

of the arm alignment control. In particular, the optical torque from the high power in the arms will affect the angular suspension dynamics.

The high circulating power in the arm cavities creates a radiation pressure which exerts a torque to the suspended mirrors. This torque provides an additional mechanical term to the suspension system, which “hardens” (mirrors rotate with the same sign) or “softens” (mirrors rotate with opposite sign) the suspension modes [135]. Given the high expected operating power of Cosmic Explorer (1.5 MW) [6], this radiation pressure effect will be significant for determining the success of the interferometer alignment stability and resulting noise couplings.

The hard and soft suspension modes will be shifted in frequency from the free suspension resonance, f_0 , via the factor [31, 48, 135]

$$\Delta f_{h,s}^2 = \frac{\gamma_{h,s} P_{cav} L_{arm}}{c I_m}, \quad (8.159)$$

where $\gamma_{h,s}$ are the geometric factors of the hard/soft mode, P_{cav} and L_{arm} represent the intra-cavity power and arm length, respectively, and I_m and c represent the mirror moment of inertia and speed of light. The factor $\gamma_{h,s}$ depends only on the cavity geometry (length of the arm and radii of curvature of the test masses); $\gamma_h > 0$, as the hard mode increases in frequency (stiffens) with increasing cavity power, and $\gamma_s < 0$, the soft mode reduces in frequency (softens) with increasing cavity power. If the soft mode frequency shift, Δf_s^2 , exceeds the fundamental frequency of the suspension, f_0 , the soft mode becomes completely unstable. Although the hard mode frequency shift is always positive and always stable, the suspension mode must be damped [31]. As such, the hard mode frequency shift, Δf_h^2 , sets a lower limit for the bandwidth of any alignment control loop applied to the arm alignment control.

With these dynamics in mind, much of the consideration for the technical noise impact of alignment controls noise comes directly from the chosen design of the suspension. Current Cosmic Explorer includes a quadruple pendulum system similar to aLIGO’s, scaled up to support the 320 kg test masses [6]. A suspension design with low-frequency mechanical resonances will be required to keep the hard mode frequency low enough to design a low bandwidth control loop. Conversely, it would be beneficial to maintain a high enough suspension frequency that the soft mode does not become unstable. With the intra-cavity power and arm length for Cosmic Explorer set to be 1.5 MW and 40 km respectively, adjusting the cavity geometry also allows some control over the impact of the radiation pressure effect. It

is beneficial for Cosmic Explorer design to limit diffraction over 40 km, so the design keeps the beam size as small as possible. This design component has an added benefit for alignment control considerations, as smaller beams have smaller geometric factors, shifting the suspension resonances less [135]. Furthermore, it is beneficial to increase the mirror moment of inertia, I_m , as this would also reduce the shift in frequency from high power. If the mass of the mirror is fixed, this involves adjusting the aspect ratio of the mirror to maximize I_m .

For the current parameters in the Cosmic Explorer design, the resulting frequency shift of the hard mode is likely to be around $+(1.1 \text{ Hz})^2$ and for the soft mode $-(0.6 \text{ Hz})^2$, as found in Ref. [31]. Both the hard and soft modes should be considered in suspension design; minimizing the instability of the soft mode is beneficial, as a control bandwidth for an unstable soft mode will need to be around $3f_s$ [31]. The bandwidth of control for the hard mode will also determine how much the injected angular noise can be suppressed in the gravitational-wave band—the lower the bandwidth, the more suppression at 5 Hz is possible. The quadruple suspension design will involve four main eigenmodes of angular motion, so the placement of both the lowest and highest frequency eigenmodes will be key design parameters.

If the lowest frequency eigenmode of the quadruple suspension is designed to be around 0.5 Hz, that means the soft mode frequency, f_s , will be about -0.1 Hz . The required control bandwidth for this unstable mode would be less than 1 Hz, enabling any sensing noise contribution from the soft mode to be suppressed by multiple orders of magnitude by 5 Hz. Assuming that the highest frequency suspension eigenmode is around 1 Hz, the hard mode frequency, f_h , would be around 1.5 Hz, and the required control bandwidth for this loop would be around 2 Hz. It is more difficult to achieve the required loop suppression for the hard loop without completely destabilizing the loop at the unity gain frequency. A suspension design that reduces the eigenmode frequencies further would be beneficial for the hard mode, but would risk increasing the bandwidth required for the soft mode control.

Furthermore, the aLIGO ASC is currently only sensing noise limited above 10 Hz [194]. Similar to the Michelson noise requirements, the displacement noise below 10 Hz will need to be reduced to reach the Cosmic Explorer design goals at 5 Hz. Again, improved local control from design aspects like the HoQI will reduce damping noise, and the proposed reductions of the seismic noise in Ref. [31] to the seismic noise will help further. However, the overall contribution of ASC noise will, in the end, depend heavily on the suspension design and resulting hard and soft mode frequencies. One way to improve ASC noise regardless of the

resulting suspension design is to consider the sensing noise level. A factor of three reduction in sensing noise would reduce the injected ASC noise to the level of the MICH noise at 20 Hz, as calculated in Section 8.2.1. This improvement would require increasing the power on the WFS to improve the shot-noise limit. Any reduction in the ASC sensing noise would help reduce requirements on the suspension design to achieve noise suppression at 5 Hz.

The coupling of alignment noise in Cosmic Explorer must be studied more closely. The calculations in this paper assume the majority of the angular coupling from Eq. 8.158 would result from residual angular motion coupling to a static beam offset. Simulations will better reveal how Cosmic Explorer upgrades can improve the coupling of beam spot motion with residual angular motion, especially in regard to improved sensors and seismic isolation, as well as the significant suspension upgrades already proposed in works like Ref. [31]. Given the potential limitation to the gravitational-wave science from alignment control noise, it will be important to research its effect and various strategies for mitigation. For example, recent work projecting the current Virgo ASC design onto the Einstein Telescope design sensitivity curve demonstrates that significant improvements in ASC noise levels are required for XG detector goals [195]. At this time, improving the sensing noise could be the most effective method for reducing ASC noise, in concert with the already-proposed design choices for Cosmic Explorer. Since much of the coupling will also be related to the resulting shifts in the suspension modes due to radiation pressure, the effects of the alignment controls should be kept in mind during any suspension design work.

8.2.3 Science impacts of low-frequency noise

The technical noises highlighted in this section will have the most impact on detector sensitivity at low frequencies. Any loss in sensitivity at low frequencies compared to the expectation from fundamental noise sources will impact XG science. This section will detail a few impacts we foresee from any degradation in low-frequency sensitivity.

To demonstrate the importance of low-frequency sensitivity, Figure 95 shows how the total sensitivity of Cosmic Explorer is impacted by the lowest sensitivity considered. The upper panel shows the Cosmic Explorer design curve [196] with the multiple low-frequency cut-offs we consider. The two lower panels show the redshift horizon² of Cosmic Explorer with respect to the total source frame mass for each of these low-frequency cutoffs. The benefit of

²We calculate the redshift horizon [52] using the `inspiral-range` package [197], the IMRPhenomX-PHM [198] waveform, and the Planck18 [199] cosmology.

a higher detector bandwidth is most extreme for high-mass signals, but it also impacts the sensitivity of Cosmic Explorer to all signals. In the most extreme scenario, a low-frequency cut-off of 20 Hz, there is over a 25% loss in sensitivity to all masses and a complete loss in sensitivity to masses above $2000 M_{\odot}$. For these high-mass systems, the higher frequency cut-off means that the merger frequency is lower than this cut-off in the detector frame due to redshift effects; this effect is demonstrated in the top panel of Figure 95.

In addition to the detectability of signals, we highlight how the low-frequency sensitivity of Cosmic Explorer will impact the estimation of source parameters from compact binary coalescences (CBCs). A given detector's power spectral density, $S(f)$, is related to the signal-to-noise ratio (SNR) of a CBC signal (and hence the horizon distance of a detector) via [178]

$$\text{SNR}^2 \propto \int \frac{f^{-7/3}}{S(f)} df. \quad (8.160)$$

However, the measurability of some CBC parameters is much more dominated by low frequencies than the SNR. For example, the measurability of chirp mass, $\Delta\mathcal{M}$, is given by [200]

$$\Delta\mathcal{M} \propto \int \frac{f^{-17/3}}{S(f)} df. \quad (8.161)$$

Due to the $f^{-17/3}$ frequency dependence, the loss in low-frequency sensitivity will have even more substantial impacts on parameter estimation than shown for the redshift horizon in previous sections.

The loss in low-frequency sensitivity will also impact multi-messenger science. One of the exciting science goals that is unlocked with the increased bandwidth of XG detectors is the potential for warning of future signals well before their merger time. XG detector networks have the ability to confidently and accurately predict the sky location of impending merger over 10 hours in advance [201, 202, 203]. However, as the time to coalescence, and hence the early warning timescale, is related to the lower frequency cutoff of the detector by [204]

$$t_{\text{warning}} \propto f^{-8/3}, \quad (8.162)$$

the lowest frequency that XG detectors are sensitive to will strongly influence the potential early warning timescale.

To demonstrate how different levels of broadband technical noise could impact the SNR and the early warning timescale for binary neutron star (BNS) signals, we estimate these quantities for a putative BNS signal with different low-frequency cut-offs. These results can




Figure 95: Top: The proposed [196] Cosmic Explorer (CE) fundamental strain sensitivity with different low-frequency cut-offs overlaid. Example compact binary coalescence (CBC) waveforms are also shown to demonstrate how different redshifts will affect the appearance of signals in the Cosmic Explorer detector. Signals from high redshifts will not be detectable by Cosmic Explorer as the signals will merge (in the detector frame) at frequencies lower than the detector’s bandwidth.

Bottom left: The furthest distance a compact binary merger with a given mass can be observed with a signal-to-noise ratio of 8 for each low-frequency cut-off shown in the top panel. Source frame mass is quoted, as the measured mass at the detector will be redshifted. Masses in the compact binary are assumed to be equal.

Bottom right: Relative redshift horizon for each low-frequency cut-off compared to the full-bandwidth Cosmic Explorer fundamental curve. Higher low-frequency cut-offs most impact the sensitivity to high-mass signals.

Detector configuration	SNR ₁₀₀₀ Horizon [Mpc]	Early warning [s]
CE fundamental	113	1340
CE, 7 Hz cut	113	931
CE, 10 Hz cut	110	536
CE, 15 Hz cut	100	214
CE, 20 Hz cut	88	102
LIGO Livingston, O3	2.4	1.0

Table 11: For the detector bandwidth cut-offs considered in Figure 95, we estimate the horizon distance at which a $1.4 M_{\odot}$ binary neutron star merger could be observed with a signal-to-noise ratio (SNR) of 1000 and the relative timescale before merger that a binary neutron star system located at a distance of 100 Mpc could be well localized. The early warning timescale is when each detector configuration could localize such a signal with the same precision as one second of early warning for LIGO Livingston in O3. In both cases the signal is assumed to be directly overhead the plane of the detector. All Cosmic Explorer values have been rounded to the nearest integer.

be seen in Table 11. We also include the same quantities for the LIGO Livingston detector in O3 for reference. The horizon distance in this table is the maximum distance at which a $1.4 M_{\odot}$ BNS signal could be observed at $\text{SNR} = 1000$ in each detector configuration, and the early warning timescale is the duration before merger that a similar system located at 100 Mpc could be well-localized. We define “well-localized” as similar in localization to 1 s of early warning for the LIGO Livingston detector in O3.

Other low-frequency science targets that will be heavily impacted include gravitational-wave displacement and spin memory effects [205], imprints of eccentricity on compact binary waveforms [206, 207], and the detection of persistent signals from rotating neutron stars (as will be discussed in Section 8.3). In general, the lowest frequencies accessible by XG detectors will be some of the most scientifically fruitful frequency regions. Maximizing the low-frequency sensitivity of XG detectors by addressing technical noise will be key to unlocking these science targets.

8.3 Narrow Noise Features

The presence of persistent narrow noise features in XG detectors will be one of the main limitations to the discovery of new sources of persistent gravitational waves and hamper the analysis of known signals. Often referred to as “lines,” these near-sinusoidal instrumental artifacts are present in all current gravitational-wave detectors at frequencies across the entire band of interest. There are a wide variety of different sources of lines, all of which have the potential to hamper astrophysical analyses [208, 209, 177, 210, 211].

While these types of instrumental artifacts may not impact analyses of transient gravitational-wave sources, searches for persistent gravitational-wave signals can be severely impaired by data containing narrow-band features [209, 208, 211]. The most well-known source of persistent gravitational waves is rapidly spinning neutron stars [212]. As shown in the left panel of Figure 96, the density per unit frequency of known spinning neutron stars (detected as pulsars) is inversely correlated with the gravitational wave frequency. Unfortunately, the same is true of instrumental lines. The frequency width impacted by instrumental lines comprises two effects: the physical width of the instrumental line and the Doppler broadening [213] of the astrophysical signal. As it is common to veto any candidate where the frequency of the signal overlaps an instrumental line after Doppler broadening is concerned, we can consider the entire frequency region lost [211].

At low frequencies, the physical width of the line dominates the bandwidth that is impacted. This can be seen in the right panel of Figure 96; as $f \rightarrow 0$, the maximum density of lines that is allowable for a 10% loss in frequency space plateaus based on the assumed width of lines. At high frequencies, the effect of Doppler broadening dominates, leading to stricter requirements for a fixed loss of bandwidth per frequency. Even though the density of lines does not grow at low frequencies as quickly as the density of pulsars, the fixed maximum density of lines means that there is a risk of significant loss in discovery potential unless the density of lines is reduced below current levels. The presence of narrow-band features is already a problem for current gravitational-wave detectors, and the lower noise floor of XG detectors will exacerbate this issue in the absence of mitigation.

Sources of lines can be roughly divided into three categories: purposefully injected lines, such as calibration lines or dither lines; well-understood lines related to the detector design, such as mechanical resonances; and spurious lines resulting from technical and environmental noise couplings. Many lines only impact a narrow frequency band, but some lines can be




Figure 96: A comparison of the density of gravitational-wave signals from pulsars and instrumental lines in the LIGO Hanford (LHO) detector during O3. Left: the relative density of signals from pulsars and line artifacts per unit frequency. No lines were recorded below 10 Hz as data from this frequency region is not considered usable. Note that the density of both pulsars and lines grows at lower frequencies. Right: a comparison in the number of lines per unit frequency and the maximum number of lines allowable per unit frequency for 10% of pulsars to be undetectable due to lines. Both the physical width of lines and the Doppler broadening of the signal are considered. In O3, over 10% of pulsars were already not observable at some frequencies. Pulsar frequencies are derived using data from the Australia Telescope National Facility Pulsar Catalog [214, 215]. LHO instrumental lines based on data from the Gravitational Wave Open Science Center [216, 217].

found in so-called “combs,” where groups of lines are present at evenly spaced frequencies and have a shared physical source [208, 172]. Spurious lines without evidence for an instrumental origin are particularly problematic; such lines are challenging to monitor and can be confused with an astrophysical signal more easily than those with a well-understood source. Ongoing monitoring and investigation of lines will be essential to address spurious lines that will appear during the operation of XG detectors.

Any lines that do not have explicit evidence of instrumental origin will be especially problematic for searches of persistent signals as lines with such evidence can be safely discarded as candidates for an astrophysical signal. Discarding candidates related to unknown sources may unintentionally veto a true signal, resulting in a trade off between polluting a search with instrumental lines that do not have evidence of their origin versus risking missing an astrophysical signal.

In addition to the challenges presented by individual lines, mixing (or “modulation”) of different instrumental lines has been observed in current detectors. When two lines with frequencies f_1 and f_2 mix, they can create additional lines at frequencies $|f_1 + f_2|$ and $|f_1 - f_2|$. Injected lines like dithers, mechanical resonances, calibration lines and spurious lines have been observed to modulate with each other and spoil additional frequency bands [44, 208, 177]. Mixing can also occur between narrow lines in the gravitational-wave frequency band and low-frequency sources, such as seismic noise [177]. Even lines that are well above the gravitational-wave frequency band, such as dither lines injected at around 2 and 4 kHz, have been known to mix with calibration lines or mechanical resonances to create artifacts at frequencies of astrophysical interest [44, 208, 177]. Line mixing is particularly problematic as it can arise unexpectedly, allowing well understood lines to create additional spurious lines. As mentioned, mixing is also one way that injected lines or mechanical resonances that are intentionally placed outside a specific frequency band can spoil frequencies of interest.

Noise from purposefully injected lines can be thought of as a technical noise similar to broadband controls noise in that they are required for the successful operation of the detector. In O3, both LIGO detectors used Alignment Dither System (ADS) lines [44]. ADS lines maintain the beam spot position on the test masses, a requirement for maintaining stable alignment controls. A dither line from each test mass is injected into the DARM signal and then demodulated to obtain an angle-to-length coupling signal that is minimized when the beam position is co-aligned with the optic rotation point [44]. Both detectors injected these lines below 20 Hz, which has the potential to impact low-frequency persistent

gravitational-wave searches. Nothing in the initial Cosmic Explorer design [6] nominally requires dither schemes over other alignment or locking schemes, but they may become necessary for successful detector operation. If dither schemes are required for XG detectors, careful consideration of the location and amplitude of dither lines should be taken to avoid contaminating searches as much as possible and careful monitoring for modulation effects should be planned.

Calibration lines are another set of purposefully-injected lines. These lines are used to track slowly occurring changes in the response of the detector. While these time-varying changes are small enough that they do not impact the operation of the detector, they may cause systematic biases in the calibrated gravitational-wave strain data if uncompensated [218, 219]. The presence of uncompensated systematic biases in the calibrated strain data would negatively impact astrophysical parameter estimation of gravitational-wave sources [220, 221]. Calibration lines are loud, known sinusoidal excitations whose detector response amplitudes and phases are tracked over time and used to correct the static calibration model to account for time dependence. The resulting correction factors are called “time-dependent correction factors” (TDCFs) [218, 219]. Both LIGO detectors use TDCFs calculated from calibration line injections for calibration purposes: to track the actuation and sensing components of the detector response function. While their frequencies may differ between LIGO Hanford and LIGO Livingston, the lines serve the same purpose at both facilities [222, 219]. The calibration lines used to track the detector’s actuation function are all placed between 10 and 20 Hz, where the actuators in the lowest three stages in the quadruple pendulum are comparable in magnitude. The detector’s optical gain and coupled-cavity pole frequency are tracked using a line near the center of aLIGO’s most sensitive frequency band; this is where their effects are most significant [219]. Also tracked are additional sensing function parameters to characterize the opto-mechanical effect of operating a detuned signal extraction cavity (~ 8 Hz) [221, 222]. Line height is determined by a desired SNR for calibration precision, the limitations of the injection system, and consideration for mitigating often unpredictable contaminations, such as those caused by linear and non-linear frequency mixing. While calibration lines will be a necessary part of detector operation, their placement, quantity, and amplitude will require consideration for the next generation of detectors due to the promise of detection of persistent gravitational waves.

Mechanical resonances of optic suspensions also appear as lines in DARM. In Advanced LIGO, the suspension chains are designed to have very low loss, so any mechanical resonance

of that suspension will have a high quality factor (Q) [44]. The same is expected for Cosmic Explorer [31, 6]. For example, resonances of the fused silica suspension fibers used to suspend the test masses and beamsplitter (often referred to as “violin modes”) appear in the DARM spectrum as well as their harmonics [44]. In the LIGO detectors, these suspension resonances can create severe line contamination when they are rung up by seismic activity [223]. Other mechanical resonances, such as the test masses’ vertical and roll mode resonances, can appear in DARM. While the most significant mechanical resonance lines in aLIGO have resulted from the suspensions of the test masses, in principle, mechanical resonances from auxiliary optics could also contaminate the spectrum. With the increased sensitivity of Cosmic Explorer, these mechanical resonances can become prominent and contaminate other portions of the spectrum. Unique to Cosmic Explorer design is the potential to also see the presence of high-Q test mass acoustic modes within the detector bandwidth due to the larger size of the test masses relative to aLIGO.

Mechanical resonances can be passively or actively damped [224], but that is usually done to mitigate the effect these lines have on locking or control loops, as these resonances can become excited in amplitude and saturate interferometric sensors [44]. Mechanical damping of these resonances lowers the quality factor of the resonance, which could broaden the peak in the DARM spectrum. Acoustic modes are also damped passively to avoid parametric instabilities that cause locklosses [44, 148]. Like mechanical damping, acoustic mode damping also lowers the quality factor of the mode [44]. Depending on the location of the resonance and the relative broadening of the peak, it may not be favorable to persistent gravitational-wave searches to add these types of damping mechanisms. However, damping may be required as a part of detector operation. Therefore, if possible, consideration of the locations of mechanical resonances should occur during the design stage of a new detector like Cosmic Explorer, as no mitigation technique applied during detector operation will be able to eliminate them entirely. Better witnesses of this mechanical motion could allow subtraction of the lines, but that may not account for line modulation unless non-linear effects, such as those introduced from digital-to-analog conversion, are considered [225, 226, 227, 159, 228].

Finally, current generation detectors have also witnessed contamination from “spurious” lines that result from unintended noise couplings [208, 177, 172]. The sources of these lines are often traced to weak coupling between electronics operating at the observatory or another periodic source of noise. For example, aLIGO has experienced significant impacts from synchronized 1 Hz blinking of LED lights on various electronics, magnetic fields produced

by vacuum sensors, and influences from ethernet cable adaptors [208, 177]. Due to the low amplitude of these instrumental artifacts, identifying them requires extensive investigations over multiple months. It also requires long time periods to evaluate if any technical changes have mitigated the source of noise [208]. To ensure that these types of lines have minimal impact on analyses, constant monitoring of the data is required, along with careful record keeping; since the effects of instrumental changes may not be noticed for months, it is essential that it is possible to precisely correlate (sometimes down to the minute) the time of appearance of an artifact with on-site activities. Similar to mechanical resonances, much of the mitigation of these lines is ideally done at the design stage, as better isolation (electronic, magnetic, and mechanical) can prevent these lines from ever appearing. However, these examples demonstrate that not everything can be reasonably anticipated in the design of an observatory; for this reason, constant monitoring of line sources and coordination between data analysts and on-site investigators will be required throughout the era of XG to catalog or mitigate new lines as they appear.

Due to the significant discovery risk that narrow-band features present to searches for persistent gravitational waves, additional research in understanding the causes of instrumental lines and how to address them in data analysis pipelines is prudent. In particular, methods of how to mitigate spurious lines through instrumental design are not well represented in the literature and will be important to understand more clearly in the near future. Furthermore, planning for XG observatories should explicitly include support for long-term monitoring and mitigation of narrow noise features.

8.4 Discussion

In this work, we have outlined how the understanding of technical noise and data quality will be essential to realize the scientific objectives of XG gravitational-wave detectors. A summary of the key science themes that we identified as highly reliant on meeting our goals in these areas is shown in Table 12. As designs for XG detectors are still under development, effects from different types of technical noise in these new detectors have not been extensively studied. When possible, design work should weigh the contribution of technical noises from various design parameters, especially if those noises have the potential to significantly inhibit science goals.

It is important to remind the reader that the majority of the projections and discussions

Technical goal	Key science objectives realized
Low controls noise	<ul style="list-style-type: none"> • Significant increase in sensitivity to cosmological sources • Significant increase in signal-to-noise ratio of all detected signals • Significant increase in early warning time for CBC signals
Low glitch rate	<ul style="list-style-type: none"> • Increased chance of identification of remnants of Population III stars and PBH • Decreased risk of data quality issues preventing analysis of short, rare events such as SNe • Decreased risk of systematic biases for analyses of astrophysical signals
High uptime	<ul style="list-style-type: none"> • Reduced risk of missing rare events such as SNe, nearby multi-messenger astronomy events, or multi-band events • Decreased uncertainty in sky localization
Low line density	<ul style="list-style-type: none"> • Ability to search large frequency spaces for persistent gravitational waves

Table 12: Summary of the main science goals that are impacted by the different technical noise sources discussed in this work. A wide range of different science themes will be enabled by improvements in technical noise and data quality. Note that this table is not exhaustive and does not contain all of the science objectives that would be realized by addressing these goals.

in this work are based on the current state of technical noise, data quality, and calibration, and are not the full picture of the impact these challenges will have on XG science. Considering the improvements made in the last decade in gravitational-wave science, it is perhaps expected that significant further research will yield improvements in all of these areas by the XG detector era. For example, comparing the Enhanced LIGO noise budget [229] to the Advanced LIGO budget [44], there was over an order of magnitude improvement in the level of controls noise for both length and alignment controls. Similarly, we have seen orders of magnitude improvements in the glitch rate between observing runs, as was the case for Advanced Virgo between O2 and O3 [171]. Therefore, the conclusions in this work should be interpreted as a reminder that continued progress in technical noise, data quality, and calibration is integral to achieving XG science goals rather than suggesting that these challenges are insurmountable. Even while XG detectors remain in the design phase, current generation gravitational wave detectors continue to expand in their abilities; additional research and development on techniques to improve the sensitivity of detectors and their data quality will be essential to maximize the science potential of both the current and next generation of gravitational-wave astrophysics.

Chapter 9

Future Work

This dissertation has covered a wide range of topics related to improving interferometer controls and noise performance. Every part of the noise spectrum is addressed, from high power upgrades intended to target the high frequency sensitivity, studies in coating thermal noise targeting the mid-frequency band, to significant focus on the low frequency noises impact current gravitational-wave detectors. This chapter provides some outlooks on where future work can be performed to further improve gravitational-wave detector sensitivity and operation.

9.1 Coating Thermal Noise

The investigation into low-noise optical coatings is one that is being undertaken by many people across the field of gravitational-wave instrumentation. The coatings presented in this paper had some favorable properties, such as low sample loss and high index of refraction. While these preliminary results possibly indicate significant absorption in the samples, further iteration on mixtures and deposition techniques may allow further improvement to these parameters. Combined with the favorable loss witnessed so far, the high index of refraction means that a coating of this type would require a much thinner high reflective stack.

Even if some properties are unfavorable, such as absorption, there are methods of mitigation. A multi-material coating, making use of some lower absorption material on the top and the new material within the bulk of the stack, could combine the beneficial properties of several materials to meet current and future gravitational-wave detector requirements [230].

9.2 Alignment Sensing

This dissertation has already explored the benefits of including in-vacuum wavefront sensing to the POP port of the Advanced LIGO interferometers. This would be a valuable addition to the alignment sensing scheme that can mitigate problems experienced now, such as low SNR for the signal-recycling cavity alignment sensors, and better control of the power-recycling cavity. It may also prove a useful tool during future power-up commissioning, as it a sensor more robust against certain control evolutions, such as moving through critical coupling.

More investigation and attention is required to understand the noise limitations of the wavefront sensors, especially in the levels of dark noise and shot noise. Improvements that allow the WFS to become shot-noise limited could greatly improve the alignment controls. This will require efforts on multiple fronts: characterization of the WFS electronics is necessary, alongside further study of the optical gain effects in the interferometer.

The alignment sensors continue to be limited by excess noise from other sources, such as damping loops and optical table motion. The effect of these noises is significant, and alignment noise continues to be at risk of limiting gravitational-wave detectors. Lower noise local sensors and improved seismic isolation are two paths forward that can greatly improve the noises that impact the alignment sensing.

9.3 Alignment Control

While alignment modeling at high power proved a useful tool during high power commissioning, limitations in the SOFT loops still have the potential to impact the operation of the detector. A more in-depth study of the full control chain of the SOFT loops, including sensor cross-coupling at the transmission QPDs, and possible actuator cross coupling from thermal effects or other means could restabilize these fragile loops.

Furthermore, radiation pressure compensation should be reinvestigated in the context of SOFT loop stabilization. While RPC was not beneficial for HARD loop control, its effect on the SOFT loops could mitigate significant stability issues once Advanced LIGO reaches full design operating power. Studying this implementation now would confirm its efficacy for future use.

9.4 Future Detectors

This dissertation provides a first look at the impact of technical noise in future detectors, informed by current generation experiences. These projections are useful, but likely an incomplete look at the technical noise landscape in the detectors to come. However, the work highlights the necessity to consider auxiliary controls within the design work at this time. Performance requirements for the beamsplitter, for example, could greatly inform bigger design requirements such as the overall arm cavity finesse. Given the extensive experience with technical challenges in the current generation, every attention should be paid to how these challenges might manifest in future detectors.

Bibliography

- [1] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophys. J. Lett.*, 195:L51–L53, 1975.
- [2] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6):061102, February 2016.
- [3] Rainer Weiss. Republication of: Electromagnetically coupled broadband gravitational antenna. *Gen. Rel. Grav.*, 54(11):153, 2022.
- [4] Andreas Freise and Kenneth Strain. Interferometer Techniques for Gravitational-Wave Detection. *Living Rev. Rel.*, 13:1, 2010.
- [5] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, 2nd. ed. edition, 2017.
- [6] Matthew Evans et al. A Horizon Study for Cosmic Explorer: Science, Observatories, and Community. 9 2021.
- [7] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. 2 2017.
- [8] Bahaa E. A. Saleh and Malvin Carl Teich. *Fundamentals of photonics: Part 1: Optics; 3rd edition*. Wiley series in pure and applied optics. Wiley, Hoboken, 2019.
- [9] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.
- [10] J. Mizuno, K. A. Strain, P. G. Nelson, J. M. Chen, R. Schilling, A. Ruediger, W. Winkler, and K. Danzmann. Resonant sideband extraction: A New configuration for interferometric gravitational wave detectors. *Phys. Lett. A*, 175:273–276, 1993.

- [11] Chris L. Mueller et al. The Advanced LIGO Input Optics. *Rev. Sci. Instrum.*, 87(1):014502, 2016.
- [12] Tobin T. Fricke et al. DC readout experiment in Enhanced LIGO. *Class. Quant. Grav.*, 29:065005, 2012.
- [13] Rich Abbott, Rana Adhikari, Stefan Ballmer, Lisa Barsotti, Matt Evans, Peter Fritschel, Valera Frolov, Guido Mueller, Bram Slagmolen, and Sam Waldman. Advanced ligo length sensing and control final design. Technical Report T1000298, LIGO Laboratory, 2010.
- [14] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015.
- [15] T. Akutsu et al. KAGRA: 2.5 Generation Interferometric Gravitational Wave Detector. *Nature Astron.*, 3(1):35–40, 2019.
- [16] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31(2):97–105, 1983.
- [17] A. Staley et al. Achieving resonance in the Advanced LIGO gravitational-wave interferometer. *Class. Quant. Grav.*, 31(24):245010, 2014.
- [18] K. Arai, M. Ando, S. Moriwaki, K. Kawabe, and K. Tsubono. New signal extraction scheme with harmonic demodulation for power-recycled Fabry-Perot-Michelson interferometers. *Phys. Lett. A*, 273:15–24, 2000.
- [19] Jeffrey Kissel. Ligo seismic isolation and suspensions cartoon. Technical Report LIGO-G1200071, LIGO Laboratory, 2023.
- [20] L. Barsotti, P. Fritschel, M. Evans, and S. Gras. Updated Advanced LIGO sensitivity design curve. Technical Report DCC-T1800044, LIGO, 2 2018.
- [21] D. Ganapathy, W. Jia, M. Nakano, V. Xu, N. Aritomi, T. Cullen, N. Kijbunchoo, S. E. Dwyer, A. Mullavey, L. McCuller, R. Abbott, I. Abouelfettouh, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, S. M. Aston, M. Ball, S. W. Ballmer, D. Barker, L. Barsotti, B. K. Berger, J. Betzwieser, D. Bhattacharjee, G. Billingsley, S. Biscans,

N. Bode, E. Bonilla, V. Bossilkov, A. Branch, A. F. Brooks, D. D. Brown, J. Bryant, C. Cahillane, H. Cao, E. Capote, F. Clara, J. Collins, C. M. Compton, R. Cottingham, D. C. Coyne, R. Crouch, J. Csizmazia, L. P. Dartez, N. Demos, E. Dohmen, J. C. Driggers, A. Effler, A. Ejlli, T. Etzel, M. Evans, J. Feicht, R. Frey, W. Frischhertz, P. Fritschel, V. V. Frolov, P. Fulda, M. Fyffe, B. Gateley, J. A. Giaime, K. D. Giardina, J. Glanzer, E. Goetz, R. Goetz, A. W. Goodwin-Jones, S. Gras, C. Gray, D. Griffith, H. Grote, T. Guidry, E. D. Hall, J. Hanks, J. Hanson, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, D. Hoyland, H. Y. Huang, Y. Inoue, A. L. James, A. Jennings, S. Karat, S. Karki, M. Kasprzack, K. Kawabe, P. J. King, J. S. Kissel, K. Komori, A. Kontos, R. Kumar, K. Kuns, M. Landry, B. Lantz, M. Laxen, K. Lee, M. Lesovsky, F. Llamas, M. Lormand, H. A. Loughlin, R. Macas, M. MacInnis, C. N. Makarem, B. Mannix, G. L. Mansell, R. M. Martin, K. Mason, F. Matchard, N. Mavalvala, N. Maxwell, G. McCarrol, R. McCarthy, D. E. McClelland, S. McCormick, T. McRae, F. Mera, E. L. Merilh, F. Meylahn, R. Mittleman, D. Moraru, G. Moreno, T. J. N. Nelson, A. Neunzert, J. Notte, J. Oberling, T. O’Hanlon, C. Osthelder, D. J. Ottaway, H. Overmier, W. Parker, A. Pele, H. Pham, M. Pirello, V. Quetschke, K. E. Ramirez, J. Reyes, J. W. Richardson, M. Robinson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, A. Sanchez, E. J. Sanchez, L. E. Sanchez, R. L. Savage, D. Schaetzel, M. G. Schiworski, R. Schnabel, R. M. S. Schofield, E. Schwartz, D. Sellers, T. Shaffer, R. W. Short, D. Sigg, B. J. J. Slagmolen, C. Soike, S. Soni, V. Srivastava, L. Sun, D. B. Tanner, M. Thomas, P. Thomas, K. A. Thorne, C. I. Torrie, G. Traylor, A. S. Ubhi, G. Vajente, J. Vanosky, A. Vecchio, P. J. Veitch, A. M. Vibhute, E. R. G. von Reis, J. Warner, B. Weaver, R. Weiss, C. Whittle, B. Willke, C. C. Wipf, H. Yamamoto, L. Zhang, and M. E. Zucker. Broadband Quantum Enhancement of the LIGO Detectors with Frequency-Dependent Squeezing. *Physical Review X*, 13(4):041021, October 2023.

- [22] Albert Einstein. On the motion of small particles suspended in liquids at rest required by the molecular-kinetic theory of heat. *Annalen der physik*, 17(549-560):208, 1905.
- [23] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32:110–113, Jul 1928.
- [24] Richard F Greene and Herbert B Callen. On the formalism of thermodynamic fluctuation theory. *Physical Review*, 83(6):1231, 1951.

- [25] Herbert B. Callen and Richard F. Greene. On a theorem of irreversible thermodynamics. *Physical Review*, 86(5):702–710, June 1952.
- [26] Richard F Greene and Herbert B Callen. On a theorem of irreversible thermodynamics. ii. *Physical Review*, 88(6):1387, 1952.
- [27] Herbert B Callen and Theodore A Welton. Irreversibility and generalized noise. *Physical Review*, 83(1):34, 1951.
- [28] Peter R Saulson. Thermal noise in mechanical experiments. *Physical Review D*, 42(8):2437, 1990.
- [29] Yu. Levin. Internal thermal noise in the LIGO test masses: A direct approach. *Physical Review D: Particles and Fields*, 57(2):659–663, January 1998.
- [30] M. Punturo et al. The Einstein Telescope: A third-generation gravitational wave observatory. *Class. Quant. Grav.*, 27:194002, 2010.
- [31] Evan D. Hall et al. Gravitational-wave physics with Cosmic Explorer: Limits to low-frequency sensitivity. *Phys. Rev. D*, 103(12):122004, 2021.
- [32] Kentaro Somiya. Detector configuration of KAGRA: The Japanese cryogenic gravitational-wave detector. *Class. Quant. Grav.*, 29:124007, 2012.
- [33] R. X. Adhikari et al. A cryogenic silicon interferometer for gravitational-wave detection. *Class. Quant. Grav.*, 37(16):165003, 2020.
- [34] E J Daw, J A Giaime, D Lormand, M Lubinski, and J Zweizig. Long-term study of the seismic environment at LIGO. *Classical and Quantum Gravity*, 21(9):2255–2273, April 2004.
- [35] Varun Srivastava. *Detector Improvements and Optimization to Advance Gravitational-wave Astronomy*. PhD thesis, Syracuse University, 2022.
- [36] F Matichard, B Lantz, R Mittleman, K Mason, J Kissel, B Abbott, S Biscans, J McIver, R Abbott, S Abbott, et al. Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, aug 2015.

- [37] S M Aston, M A Barton, A S Bell, N Beveridge, B Bland, A J Brummitt, G Cagnoli, C A Cantley, L Carbone, A V Cumming, et al. Update on quadruple suspension design for Advanced LIGO. *Classical and Quantum Gravity*, 29(23):235004, oct 2012.
- [38] Michael Coughlin, Nikhil Mukund, Jan Harms, Jenne Driggers, Rana Adhikari, and Sanjit Mitra. Towards a first design of a Newtonian-noise cancellation system for Advanced LIGO. *Class. Quant. Grav.*, 33(24):244001, 2016.
- [39] M. W. Coughlin, J. Harms, J. Driggers, D. J. McManus, N. Mukund, M. P. Ross, B. J. J. Slagmolen, and K. Venkateswara. Implications of dedicated seismometer measurements on Newtonian-noise cancellation for Advanced LIGO. *Phys. Rev. Lett.*, 121(22):221104, 2018.
- [40] Jennifer C. Driggers, Matthew Evans, Keenan Pepper, and Rana Adhikari. Active noise cancellation in a suspended interferometer. *Rev. Sci. Instrum.*, 83:024501, 2012.
- [41] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Phys. Rev. D*, 78:102003, 2008.
- [42] Camilla Compton. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=74681>.
- [43] Naoki Aritomi. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=74788>.
- [44] A. Buikema, C. Cahillane, G. L. Mansell, C. D. Blair, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, J. S. Areeda, Y. Asali, S. M. Aston, C. Austin, A. M. Baer, M. Ball, S. W. Ballmer, S. Banagiri, D. Barker, L. Barsotti, J. Bartlett, B. K. Berger, J. Betzwieser, D. Bhattacharjee, G. Billingsley, S. Biscans, R. M. Blair, N. Bode, P. Booker, R. Bork, A. Bramley, A. F. Brooks, D. D. Brown, K. C. Cannon, X. Chen, A. A. Ciobanu, F. Clara, S. J. Cooper, K. R. Corley, S. T. Countryman, P. B. Covas, D. C. Coyne, L. E. H. Datrier, D. Davis, C. Di Fronzo, K. L. Dooley, J. C. Driggers, P. Dupej, S. E. Dwyer, A. Effler, T. Etzel, M. Evans, T. M. Evans, J. Feicht, A. Fernandez-Galiana, P. Fritschel, V. V. Frolov, P. Fulda, M. Fyffe, J. A. Giaime, K. D. Giardina, P. Godwin, E. Goetz, S. Gras, C. Gray, R. Gray, A. C. Green, E. K.

- Gustafson, R. Gustafson, J. Hanks, J. Hanson, T. Hardwick, R. K. Hasskew, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, J. D. Jones, S. Kandhasamy, S. Karki, M. Kasprzack, K. Kawabe, N. Kijbunchoo, P. J. King, J. S. Kissel, Rahul Kumar, M. Landry, B. B. Lane, B. Lantz, M. Laxen, Y. K. Lecoeuche, J. Leviton, J. Liu, M. Lormand, A. P. Lundgren, R. Macas, M. MacInnis, D. M. Macleod, S. Márka, Z. Márka, D. V. Martynov, K. Mason, T. J. Massinger, F. Matichard, N. Mavalvala, R. McCarthy, D. E. McClelland, S. McCormick, L. McCuller, J. McIver, T. McRae, G. Mendell, K. Merfeld, E. L. Merilh, F. Meylahn, T. Mistry, R. Mittleman, G. Moreno, C. M. Mow-Lowry, S. Mozzon, A. Mullavey, T. J. N. Nelson, P. Nguyen, L. K. Nuttall, J. Oberling, Richard J. Oram, B. O'Reilly, C. Osthelder, D. J. Ottaway, H. Overmier, J. R. Palamos, W. Parker, E. Payne, A. Pele, R. Penhorwood, C. J. Perez, M. Pirello, H. Radkins, K. E. Ramirez, J. W. Richardson, K. Riles, N. A. Robertson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, E. J. Sanchez, L. E. Sanchez, T. R. Saravanan, R. L. Savage, D. Schaetzl, R. Schnabel, R. M. S. Schofield, E. Schwartz, D. Sellers, T. Shaffer, D. Sigg, B. J. J. Slagmolen, J. R. Smith, S. Soni, B. Sorazu, A. P. Spencer, K. A. Strain, L. Sun, M. J. Szczepańczyk, M. Thomas, P. Thomas, K. A. Thorne, K. Toland, C. I. Torrie, G. Traylor, M. Tse, A. L. Urban, G. Vajente, G. Valdes, D. C. Vander-Hyde, P. J. Veitch, K. Venkateswara, G. Venugopalan, A. D. Viets, T. Vo, C. Vorvick, M. Wade, R. L. Ward, J. Warner, B. Weaver, R. Weiss, C. Whittle, B. Willke, C. C. Wipf, L. Xiao, H. Yamamoto, Hang Yu, Haocun Yu, L. Zhang, M. E. Zucker, and J. Zweizig. Sensitivity and performance of the Advanced LIGO detectors in the third observing run. *Physical Review D: Particles and Fields*, 102(6):062003, September 2020.
- [45] E. D. Hall. PhD thesis, California Institute of Technology, 2017.
- [46] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=76279>.
- [47] Camilla Compton. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=73773>.
- [48] L. Barsotti, M. Evans, and P. Fritschel. Alignment sensing and control in advanced LIGO. *Class. Quant. Grav.*, 27:084026, 2010.

- [49] L. Barsotti and M. Evans. of Alignment Sensing and Control for Advanced LIGO. Technical Report DCC-T0900511, LIGO, 10 2009.
- [50] Guido Mueller. Beam jitter coupling in advanced ligo. *Opt. Express*, 13(18):7118–7132, Sep 2005.
- [51] J. Abadie et al. Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO’s Sixth and Virgo’s Second and Third Science Runs. 3 2012.
- [52] Hsin-Yu Chen, Daniel E. Holz, John Miller, Matthew Evans, Salvatore Vitale, and Jolien Creighton. Distance measures in gravitational-wave astrophysics and cosmology. *Class. Quant. Grav.*, 38(5):055010, 2021.
- [53] Evan Hall. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=68424>.
- [54] L. Barsotti, L. McCuller, M. Evans, and P. Fritschel. The A+ design curve. Technical Report DCC-T1800042, LIGO, 2 2018.
- [55] James Alexander Clark, Andreas Bauswein, Nikolaos Stergioulas, and Deirdre Shoemaker. Observing Gravitational Waves From The Post-Merger Phase Of Binary Neutron Star Coalescence. *Class. Quant. Grav.*, 33(8):085003, 2016.
- [56] Mark G. Alford, Luke Bovard, Matthias Hanauske, Luciano Rezzolla, and Kai Schwenzer. Viscous Dissipation and Heat Conduction in Binary Neutron-Star Mergers. *Phys. Rev. Lett.*, 120(4):041101, 2018.
- [57] Carolyn A. Raithel and Elias R. Most. Characterizing the Breakdown of Quasi-universality in Postmerger Gravitational Waves from Binary Neutron Star Mergers. *Astrophys. J. Lett.*, 933(2):L39, 2022.
- [58] Nayyer Raza, Jess McIver, Gergely Dály, and Peter Raffai. Prospects for reconstructing the gravitational-wave signals from core-collapse supernovae with Advanced LIGO-Virgo and the BayesWave algorithm. *Phys. Rev. D*, 106(6):063014, 2022.
- [59] K. Ackley et al. Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global network. *Publ. Astron. Soc. Austral.*, 37:e047, 2020.

- [60] E. Cesarini, M. Lorenzini, E. Campagna, F. Martelli, F. Piergiovanni, F. Vetrano, G. Losurdo, and G. Cagnoli. A “gentle” nodal suspension for measurements of the acoustic attenuation in materials. *Rev. Sci. Instrum.*, 80(5):053904, 2009.
- [61] Gregory M. Harry et al. Titania-doped tantala/silica coatings for gravitational-wave detection. *Class. Quant. Grav.*, 24:405–416, 2007.
- [62] Massimo Granata et al. Mechanical loss in state-of-the-art amorphous optical coatings. *Phys. Rev. D*, 93(1):012007, 2016.
- [63] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, B. Allen, A. Allocca, P. A. Altin, S. B. Anderson, W. G. Anderson, K. Arai, M. C. Araya, C. C. Arceneaux, J. S. Areeda, N. Arnaud, K. G. Arun, S. Ascenzi, G. Ashton, M. Ast, S. M. Aston, P. Astone, P. Aufmuth, C. Aulbert, S. Babak, P. Bacon, M. K. M. Bader, P. T. Baker, F. Baldaccini, G. Ballardín, S. W. Ballmer, J. C. Barayoga, S. E. Barclay, B. C. Barish, D. Barker, F. Barone, B. Barr, L. Barsotti, M. Barsuglia, D. Barta, J. Bartlett, I. Bartos, R. Bassiri, A. Basti, J. C. Batch, C. Baune, V. Bavigadda, M. Bazzan, B. Behnke, M. Bejger, A. S. Bell, C. J. Bell, B. K. Berger, J. Bergman, G. Bergmann, C. P. L. Berry, D. Bersanetti, A. Bertolini, J. Betzwieser, S. Bhagwat, R. Bhandare, I. A. Bilenko, G. Billingsley, J. Birch, R. Birney, S. Biscans, A. Bisht, M. Bitossi, C. Biwer, M. A. Bizouard, J. K. Blackburn, C. D. Blair, D. G. Blair, R. M. Blair, S. Bloemen, O. Bock, T. P. Bodiya, M. Boer, G. Bogaert, C. Bogan, A. Bohe, P. Bojtos, C. Bond, F. Bondu, R. Bonnand, B. A. Boom, R. Bork, V. Boschi, S. Bose, Y. Bouffanais, A. Bozzi, C. Bradaschia, P. R. Brady, V. B. Braginsky, M. Branchesi, J. E. Brau, T. Briant, A. Brillet, M. Brinkmann, V. Brisson, P. Brockill, A. F. Brooks, D. A. Brown, D. D. Brown, N. M. Brown, C. C. Buchanan, A. Buikema, T. Bulik, H. J. Bulten, A. Buonanno, D. Buskulic, C. Buy, R. L. Byer, L. Cadonati, G. Cagnoli, C. Cahillane, J. Calderón Bustillo, T. Callister, E. Calloni, J. B. Camp, K. C. Cannon, J. Cao, C. D. Capano, E. Capocasa, F. Carbognani, S. Caride, J. Casanueva Diaz, C. Casentini, S. Caudill, M. Cavaglià, F. Cavalier, R. Cavalieri, G. Cella, C. B. Cepeda, L. Cerboni Baiardi, G. Cerretani, E. Cesarini, R. Chakraborty, T. Chalermongsak, S. J. Chamberlin, M. Chan, S. Chao, P. Charlton, E. Chassande-Mottin, H. Y. Chen, Y. Chen, C. Cheng, A. Chincarini, A. Chiummo, H. S. Cho,

M. Cho, J. H. Chow, N. Christensen, Q. Chu, S. Chua, S. Chung, G. Ciani, F. Clara, J. A. Clark, F. Cleva, E. Coccia, P.-F. Cohadon, A. Colla, C. G. Collette, L. Cominsky, M. Constancio, A. Conte, L. Conti, D. Cook, T. R. Corbitt, N. Cornish, A. Corsi, S. Cortese, C. A. Costa, M. W. Coughlin, S. B. Coughlin, J.-P. Coulon, S. T. Countryman, P. Couvares, E. E. Cowan, D. M. Coward, M. J. Cowart, D. C. Coyne, R. Coyne, K. Craig, J. D. E. Creighton, J. Cripe, S. G. Crowder, A. Cumming, L. Cunningham, E. Cuoco, T. Dal Canton, S. L. Danilishin, S. D'Antonio, K. Danzmann, N. S. Derman, V. Dattilo, I. Dave, H. P. Daveloza, M. Davier, G. S. Davies, E. J. Daw, R. Day, D. DeBra, G. Debreczeni, J. Degallaix, M. De Laurentis, S. Deléglise, W. Del Pozzo, T. Denker, T. Dent, H. Dereli, V. Dergachev, R. T. DeRosa, R. De Rosa, R. DeSalvo, S. Dhurandhar, M. C. Díaz, L. Di Fiore, M. Di Giovanni, A. Di Lieto, S. Di Pace, I. Di Palma, A. Di Virgilio, G. Dojcinoski, V. Dolique, F. Donovan, K. L. Dooley, S. Doravari, R. Douglas, T. P. Downes, M. Drago, R. W. P. Drever, J. C. Driggers, Z. Du, M. Ducrot, S. E. Dwyer, T. B. Edo, M. C. Edwards, A. Effler, H.-B. Eggenstein, P. Ehrens, J. Eichholz, S. S. Eikenberry, W. Engels, R. C. Essick, T. Etzel, M. Evans, T. M. Evans, R. Everett, M. Factourovich, V. Fafone, H. Fair, S. Fairhurst, X. Fan, Q. Fang, S. Farinon, B. Farr, W. M. Farr, M. Favata, M. Fays, H. Fehrmann, M. M. Fejer, I. Ferrante, E. C. Ferreira, F. Ferrini, F. Fidecaro, I. Fiori, D. Fiorucci, R. P. Fisher, R. Flaminio, M. Fletcher, J.-D. Fournier, S. Franco, S. Frasca, F. Frasconi, Z. Frei, A. Freise, R. Frey, V. Frey, T. T. Fricke, P. Fritschel, V. V. Frolov, P. Fulda, M. Fyffe, H. A. G. Gabbard, J. R. Gair, L. Gammaitoni, S. G. Gaonkar, F. Garufi, A. Gatto, G. Gaur, N. Gehrels, G. Gemme, B. Gendre, E. Genin, A. Genai, J. George, L. Gergely, V. Germain, Archisman Ghosh, S. Ghosh, J. A. Giaime, K. D. Giardino, A. Giazotto, K. Gill, A. Glaefke, E. Goetz, R. Goetz, L. Gondan, G. González, J. M. Gonzalez Castro, A. Gopakumar, N. A. Gordon, M. L. Gorodetsky, S. E. Gossan, M. Gosselin, R. Gouaty, C. Graef, P. B. Graff, M. Granata, A. Grant, S. Gras, C. Gray, G. Greco, A. C. Green, P. Groot, H. Grote, S. Grunewald, G. M. Guidi, X. Guo, A. Gupta, M. K. Gupta, K. E. Gushwa, E. K. Gustafson, R. Gustafson, J. J. Hacker, B. R. Hall, E. D. Hall, G. Hammond, M. Haney, M. M. Hanke, J. Hanks, C. Hanna, M. D. Hannam, J. Hanson, T. Hardwick, K. Haris, J. Harms, G. M. Harry, I. W. Harry, M. J. Hart, M. T. Hartman, C.-J. Haster, K. Haughian, A. Heidmann, M. C. Heintze, H. Heitmann, P. Hello, G. Hemming, M. Hendry, I. S. Heng, J. Hennig, A. W. Heptonstall, M. Heurs, S. Hild, D. Hoak, K. A. Hodge, D. Hofman, S. E.

Hollitt, K. Holt, D. E. Holz, P. Hopkins, D. J. Hosken, J. Hough, E. A. Houston, E. J. Howell, Y. M. Hu, S. Huang, E. A. Huerta, D. Huet, B. Hughey, S. Husa, S. H. Huttner, T. Huynh-Dinh, A. Idrisy, N. Indik, D. R. Ingram, R. Inta, H. N. Isa, J.-M. Isac, M. Isi, G. Islas, T. Isogai, B. R. Iyer, K. Izumi, T. Jacqmin, H. Jang, K. Jani, P. Jaranowski, S. Jawahar, F. Jiménez-Forteza, W. W. Johnson, D. I. Jones, R. Jones, R. J. G. Jonker, L. Ju, C. V. Kalaghatgi, V. Kalogera, S. Kandhasamy, G. Kang, J. B. Kanner, S. Karki, M. Kasprzack, E. Katsavounidis, W. Katzman, S. Kaufer, T. Kaur, K. Kawabe, F. Kawazoe, F. Kéfélian, M. S. Kehl, D. Keitel, D. B. Kelley, W. Kells, R. Kennedy, J. S. Key, A. Khalaidovski, F. Y. Khalili, I. Khan, S. Khan, Z. Khan, E. A. Khazanov, N. Kijbunchoo, C. Kim, J. Kim, K. Kim, Nam-Gyu Kim, Namjun Kim, Y.-M. Kim, E. J. King, P. J. King, D. L. Kinzel, J. S. Kissel, L. Kleybolte, S. Klimenko, S. M. Koehlenbeck, K. Kokeyama, S. Koley, V. Kondrashov, A. Kontos, M. Korobko, W. Z. Korth, I. Kowalska, D. B. Kozak, V. Kringel, A. Królak, C. Krueger, G. Kuehn, P. Kumar, L. Kuo, A. Kutynia, B. D. Lackey, M. Landry, J. Lange, B. Lantz, P. D. Lasky, A. Lazzarini, C. Lazzaro, P. Leaci, S. Leavey, E. O. Lebigot, C. H. Lee, H. K. Lee, H. M. Lee, K. Lee, A. Lenon, M. Leonardi, J. R. Leong, N. Leroy, N. Letendre, Y. Levin, B. M. Levine, T. G. F. Li, A. Libson, T. B. Littenberg, N. A. Lockerbie, J. Logue, A. L. Lombardi, J. E. Lord, M. Lorenzini, V. Lorette, M. Lormand, G. Losurdo, J. D. Lough, H. Lück, A. P. Lundgren, J. Luo, R. Lynch, Y. Ma, T. MacDonald, B. Machenschalk, M. MacInnis, D. M. Macleod, F. Magaña Sandoval, R. M. Magee, M. Mageswaran, E. Majorana, I. Maksimovic, V. Malvezzi, N. Man, I. Mandel, V. Mandic, V. Mangano, G. L. Mansell, M. Manske, M. Mantovani, F. Marchesoni, F. Marion, S. Márka, Z. Márka, A. S. Markosyan, E. Maros, F. Martelli, L. Martellini, I. W. Martin, R. M. Martin, D. V. Martynov, J. N. Marx, K. Mason, A. Masserot, T. J. Massinger, M. Masso-Reid, F. Matichard, L. Matone, N. Mavalvala, N. Mazumder, G. Mazzolo, R. McCarthy, D. E. McClelland, S. McCormick, S. C. McGuire, G. McIntyre, J. McIver, D. J. McManus, S. T. McWilliams, D. Meacher, G. D. Meadors, J. Meidam, A. Melatos, G. Mendell, D. Mendoza-Gandara, R. A. Mercer, E. Merilh, M. Merzougui, S. Meshkov, C. Messenger, C. Messick, P. M. Meyers, F. Mezzani, H. Miao, C. Michel, H. Middleton, E. E. Mikhailov, L. Milano, J. Miller, M. Millhouse, Y. Minenkov, J. Ming, S. Mirshekari, C. Mishra, S. Mitra,

V. P. Mitrofanov, G. Mitselmakher, R. Mittleman, A. Moggi, M. Mohan, S. R. P. Mohapatra, M. Montani, B. C. Moore, C. J. Moore, D. Moraru, G. Moreno, S. R. Morris, K. Mossavi, B. Mours, C. M. Mow-Lowry, C. L. Mueller, G. Mueller, A. W. Muir, Arunava Mukherjee, D. Mukherjee, S. Mukherjee, N. Mukund, A. Mullavey, J. Munch, D. J. Murphy, P. G. Murray, A. Mytidis, I. Nardecchia, L. Naticchioni, R. K. Nayak, V. Necula, K. Nedkova, G. Nelemans, M. Neri, A. Neunzert, G. Newton, T. T. Nguyen, A. B. Nielsen, S. Nissanke, A. Nitz, F. Nocera, D. Nolting, M. E. Normandin, L. K. Nuttall, J. Oberling, E. Ochsner, J. O'Dell, E. Oelker, G. H. Ogin, J. J. Oh, S. H. Oh, F. Ohme, M. Oliver, P. Oppermann, Richard J. Oram, B. O'Reilly, R. O'Shaughnessy, D. J. Ottaway, R. S. Ottens, H. Overmier, B. J. Owen, A. Pai, S. A. Pai, J. R. Palamos, O. Palashov, C. Palomba, A. Pal-Singh, H. Pan, C. Pankow, F. Pannarale, B. C. Pant, F. Paoletti, A. Paoli, M. A. Papa, H. R. Paris, W. Parker, D. Pascucci, A. Pasqualetti, R. Passaquieti, D. Passuello, B. Patricelli, Z. Patrick, B. L. Pearlstone, M. Pedraza, R. Pedurand, L. Pekowsky, A. Pele, S. Penn, A. Perreca, M. Phelps, O. Piccinni, M. Pichot, F. Piergiovanni, V. Pierro, G. Pillant, L. Pinard, I. M. Pinto, M. Pitkin, R. Poggiani, P. Popolizio, A. Post, J. Powell, J. Prasad, V. Predoi, S. S. Premachandra, T. Prestegard, L. R. Price, M. Prijatelj, M. Principe, S. Privitera, G. A. Prodi, L. Prokhorov, O. Puncken, M. Punturo, P. Puppo, M. Pürerer, H. Qi, J. Qin, V. Quetschke, E. A. Quintero, R. Quitzow-James, F. J. Raab, D. S. Rabeling, H. Radkins, P. Raffai, S. Raja, M. Rakhmanov, P. Rapagnani, V. Raymond, M. Razzano, V. Re, J. Read, C. M. Reed, T. Regimbau, L. Rei, S. Reid, D. H. Reitze, H. Rew, S. D. Reyes, F. Ricci, K. Riles, N. A. Robertson, R. Robie, F. Robinet, A. Rocchi, L. Rolland, J. G. Rollins, V. J. Roma, R. Romano, G. Romanov, J. H. Romie, D. Rosińska, S. Rowan, A. Rüdiger, P. Ruggi, K. Ryan, S. Sachdev, T. Sadecki, L. Sadeghian, L. Salconi, M. Saleem, F. Salemi, A. Samajdar, L. Sammut, E. J. Sanchez, V. Sandberg, B. Sandeen, J. R. Sanders, B. Sassolas, B. S. Sathyaprakash, P. R. Saulson, O. Sauter, R. L. Savage, A. Sawadsky, P. Schale, R. Schilling, J. Schmidt, P. Schmidt, R. Schnabel, R. M. S. Schofield, A. Schönbeck, E. Schreiber, D. Schuette, B. F. Schutz, J. Scott, S. M. Scott, D. Sellers, A. S. Sengupta, D. Sentenac, V. Sequino, A. Sergeev, G. Serna, Y. Setyawati, A. Sevigny, D. A. Shaddock, S. Shah, M. S. Shahriar, M. Shaltev, Z. Shao, B. Shapiro, P. Shawhan, A. Sheperd, D. H. Shoemaker, D. M. Shoemaker, K. Siellez, X. Siemens, D. Sigg, A. D. Silva, D. Simakov, A. Singer, L. P. Singer, A. Singh, R. Singh, A. Singhal, A. M. Sintes, B. J. J. Slagmolen, J. R. Smith,

N. D. Smith, R. J. E. Smith, E. J. Son, B. Sorazu, F. Sorrentino, T. Souradeep, A. K. Srivastava, A. Staley, M. Steinke, J. Steinlechner, S. Steinlechner, D. Steinmeyer, B. C. Stephens, R. Stone, K. A. Strain, N. Straniero, G. Stratta, N. A. Strauss, S. Strigin, R. Sturani, A. L. Stuver, T. Z. Summerscales, L. Sun, P. J. Sutton, B. L. Swinkels, M. J. Szczepańczyk, M. Tacca, D. Talukder, D. B. Tanner, M. Tápai, S. P. Tarabrin, A. Taracchini, R. Taylor, T. Theeg, M. P. Thirugnanasambandam, E. G. Thomas, M. Thomas, P. Thomas, K. A. Thorne, K. S. Thorne, E. Thrane, S. Tiwari, V. Tiwari, K. V. Tokmakov, C. Tomlinson, M. Tonelli, C. V. Torres, C. I. Torrie, D. Töyrä, F. Travasso, G. Traylor, D. Trifirò, M. C. Tringali, L. Trozzo, M. Tse, M. Turconi, D. Tuyenbayev, D. Ugolini, C. S. Unnikrishnan, A. L. Urban, S. A. Usman, H. Vahlbruch, G. Vajente, G. Valdes, N. van Bakel, M. van Beuzekom, J. F. J. van den Brand, C. Van Den Broeck, D. C. Vander-Hyde, L. van der Schaaf, J. V. van Heijningen, A. A. van Veggel, M. Vardaro, S. Vass, M. Vasúth, R. Vaulin, A. Vecchio, G. Vedovato, J. Veitch, P. J. Veitch, K. Venkateswara, D. Verkindt, F. Vetrano, A. Viceré, S. Vinciguerra, D. J. Vine, J.-Y. Vinet, S. Vitale, T. Vo, H. Vocca, C. Vorvick, D. Voss, W. D. Voudsen, S. P. Vyatchanin, A. R. Wade, L. E. Wade, M. Wade, M. Walker, L. Wallace, S. Walsh, G. Wang, H. Wang, M. Wang, X. Wang, Y. Wang, R. L. Ward, J. Warner, M. Was, B. Weaver, L.-W. Wei, M. Weinert, A. J. Weinstein, R. Weiss, T. Welborn, L. Wen, P. Weßels, T. Westphal, K. Wette, J. T. Whelan, S. E. Whitcomb, D. J. White, B. F. Whiting, R. D. Williams, A. R. Williamson, J. L. Willis, B. Willke, M. H. Wimmer, W. Winkler, C. C. Wipf, H. Wittel, G. Woan, J. Worden, J. L. Wright, G. Wu, J. Yablon, W. Yam, H. Yamamoto, C. C. Yancey, M. J. Yap, H. Yu, M. Yvert, A. Zadrożny, L. Zangrando, M. Zanolin, J.-P. Zendri, M. Zevin, F. Zhang, L. Zhang, M. Zhang, Y. Zhang, C. Zhao, M. Zhou, Z. Zhou, X. J. Zhu, M. E. Zucker, S. E. Zuraw, and J. Zweigig. Gw150914: The advanced ligo detectors in the era of first discoveries. *Phys. Rev. Lett.*, 116:131103, Mar 2016.

- [64] Gregg Harry. Advanced LIGO Input Test Mass Coating Specification. Technical Report LIGO-E0900041, LIGO Laboratory, 2009.
- [65] Rand Dannenberg. Advanced LIGO End Test Mass (ETM) Coating Specification. Technical Report LIGO-E0900068, LIGO Laboratory, 2009.
- [66] A. Amato, G. Cagnoli, M. Granata, B. Sassolas, J. Degallaix, D. Forest, C. Michel, L. Pinard, N. Demos, S. Gras, M. Evans, A. Di Michele, and M. Canepa. Optical

- and mechanical properties of ion-beam-sputtered Nb₂O₅ and TiO₂-Nb₂O₅ thin films for gravitational-wave interferometers and an improved measurement of coating thermal noise in Advanced LIGO. *Physical Review D: Particles and Fields*, 103(7):072001, April 2021.
- [67] P. Fritschel, K. Kuns, J. Driggers, A. Effler, B. Lantz, D. Ottoway, S. Ballmer, K. Dooley, R. Adhikari, M. Evans, B. Farr, G. Gonzalez, P. Schmidt, and S. Raja. Report of the LSC Post-O5 Study Group. Technical Report DCC-T2200287, LIGO, 8 2022.
- [68] Gabriele Vajente. Notes on measuring coatings with a nodal suspension. Technical Report LIGO-T1900276, LIGO, 2019.
- [69] Gregory M. Harry et al. Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings. *Class. Quant. Grav.*, 19:897–918, 2002.
- [70] Nicholas Didio. *AlGaAs Coating Studies for Present and Future Gravitational Wave Detectors*. PhD thesis, Syracuse University, 2023.
- [71] National Instruments. <https://www.ni.com/en-us/shop/labview.html>.
- [72] Tianjun Li, Felipe A Aguilar Sandoval, Mickael Geitner, Ludovic Bellon, Gianpietro Cagnoli, Jérôme Degallaix, Vincent Dolique, Raffaele Flaminio, Danièle Forest, Massimo Granata, et al. Measurements of mechanical thermal noise and energy dissipation in optical dielectric coatings. *Physical Review D*, 89(9):092004, 2014.
- [73] COMSOL. <http://www.comsol.com/>.
- [74] Massimo Granata et al. Amorphous optical coatings of present gravitational-wave interferometers. *Class. Quant. Grav.*, 37(9):095004, 2020.
- [75] J. Steinlechner. Development of mirror coatings for gravitational-wave detectors. *Phil. Trans. Roy. Soc. Lond. A*, 376(2120):20170282, 2018.
- [76] Mariana A. Fazio, Gabriele Vajente, Alena Ananyeva, Ashot Markosyan, Riccardo Bassiri, Martin M. Fejer, and Carmen S. Menoni. Structure and morphology of low mechanical loss tio₂-doped ta₂o₅. *Opt. Mater. Express*, 10(7):1687–1703, Jul 2020.
- [77] I. W. Martin et al. Effect of heat treatment on mechanical dissipation in Ta₂O₅ coatings. *Class. Quant. Grav.*, 27:225020, 2010.

- [78] Roger P. Netterfield, Mark Gross, Fred N. Baynes, Katie L. Green, Gregory M. Harry, Helena Armandula, Sheila Rowan, Jim Hough, David R. M. Crooks, Martin M. Fejer, Roger Route, and Steven D. Penn. Low mechanical loss coatings for LIGO optics: progress report. In Michael L. Fulton and Jennifer D. T. Kruschwitz, editors, *Advances in Thin-Film Coatings for Optical Applications II*, volume 5870, page 58700H. International Society for Optics and Photonics, SPIE, 2005.
- [79] Elenna M. Capote, Amy Gleckl, Jazlyn Guerrero, Michael Rezac, Robert Wright, and Joshua R. Smith. In-vacuum measurements of optical scatter versus annealing temperature for amorphous ta₂o₅ and tio₂:ta₂o₅ thin films. *J. Opt. Soc. Am. A*, 38(4):534–541, Apr 2021.
- [80] Émile Lalande, Alexandre W. Lussier, Carl Lévesque, Marianne Ward, Bill Baloukas, Ludvik Martinu, Gabriele Vajente, Garilynn Billingsley, Alena Ananyeva, Riccardo Bassiri, Martin M. Fejer, and François Schiettekatte. Zirconia-titania-doped tantala optical coatings for low mechanical loss Bragg mirrors. *Journal of Vacuum Science and Technology A*, 39(4):043416, 06 2021.
- [81] K. Prasai, J. Jiang, A. Mishkin, B. Shyam, S. Angelova, R. Birney, D. A. Drabold, M. Fazio, E. K. Gustafson, G. Harry, S. Hoback, J. Hough, C. Lévesque, I. MacLaren, A. Markosyan, I. W. Martin, C. S. Menoni, P. G. Murray, S. Penn, S. Reid, R. Robie, S. Rowan, F. Schiettekatte, R. Shink, A. Turner, G. Vajente, H-P. Cheng, M. M. Fejer, A. Mehta, and R. Bassiri. High precision detection of change in intermediate range order of amorphous zirconia-doped tantala thin films due to annealing. *Phys. Rev. Lett.*, 123:045501, Jul 2019.
- [82] F. Samiee, K. Raeissi, and M. A. Golozar. Nanoindentation testing of pulse electrodeposited thin zirconia coatings. *Surface Engineering*, 29(10):726–730, 2013.
- [83] O. Anderson, C. R. Ottermann, R. Kuschnerreit, P. Hess, and K. Bange. Density and Young’s modulus of thin TiO₂ films. *Fresenius’ Journal of Analytical Chemistry*, 358(1):315–318, 1997.
- [84] E. Hecht. *Optics*. Pearson, 2012.

- [85] Gabriele Vajente et al. Low Mechanical Loss TiO₂:GeO₂ Coatings for Reduced Thermal Noise in Gravitational Wave Interferometers. *Phys. Rev. Lett.*, 127(7):071101, 2021.
- [86] Fabian Magaña Sandoval, Thomas Vo, Daniel Vander-Hyde, J. R. Sanders, and Stefan W. Ballmer. Sensing Optical Cavity Mismatch with a Mode-Converter and Quadrant Photodiode. *Phys. Rev. D*, 100(10):102001, 2019.
- [87] H. Kogelnik and T. Li. Laser beams and resonators. *Appl. Opt.*, 5(10):1550–1567, Oct 1966.
- [88] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=62531>.
- [89] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=62828>.
- [90] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=62883>.
- [91] Nergis Mavalvala, Daniel Sigg, and David Shoemaker. Experimental test of an alignment-sensing scheme for a gravitational-wave interferometer. *Appl. Opt.*, 37(33):7743–7746, Nov 1998.
- [92] "Daniel Brown". Arm asc yaw and pitch degrees of freedom, 2022. "LIGO-D2200425".
- [93] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=64545>.
- [94] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=67334>.
- [95] Wenxuan Jia and Matthew Evans. Non-invasive beam spot position monitoring using cameras - LHO F2F Commissioning Meeting. Technical Report LIGO-G2202043, LIGO Laboratory, 2022.
- [96] Daniel Brown. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=66197>.

- [97] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=66462>.
- [98] Elenna Capote. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=67099>.
- [99] Terra Hardwick and Craig Ingram. aligo optical sensors graphic (psl 1064 nm beam). Technical Report LIGO-G1601619, LIGO Laboratory, 2016.
- [100] Julius S. Bendat and Allan G. Piersol. *Random Data: Analysis and Measurement Procedures*. Wiley, 2010.
- [101] Elenna Capote. Alignment sensing and control using in-vacuum pop wavefront sensors. Technical Report LIGO-T2400054, LIGO Hanford Observatory, 2024.
- [102] Daniel Sigg Jenne Driggers. The need for an in-vacuum pop wfs. Technical report, LIGO, 2016. <https://dcc.ligo.org/LIGO-T1600580>.
- [103] <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=67831>.
- [104] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=67706>.
- [105] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=74106>.
- [106] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=43259>.
- [107] <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=63158>.
- [108] <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=67831>.
- [109] Finesse3. <https://finesse.ifosim.org>.
- [110] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=72137>.
- [111] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=72199>.
- [112] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=68385>.
- [113] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=68456>.
- [114] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=63624>.

- [115] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=66734>.
- [116] Tobin Fricke. *Homodyne detection for laser-interferometric gravitational wave detectors*. PhD thesis, Louisiana State University, 2011.
- [117] Craig Russell Cahillane. *Controlling and Calibrating Interferometric Gravitational Wave Detectors*. PhD thesis, California Institute of Technology, 2021.
- [118] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=44947>.
- [119] Rich Abbott and Peter Fritschel. aligo i&q rf demodulator design. Technical Report LIGO-T1000044, LIGO Laboratory, 2010.
- [120] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=46509>.
- [121] Evan Hall. Comparison of refl wfs error signals between hanford and livingston in july 2022. Technical Report LIGO-T2400107, LIGO Laboratory, 2022.
- [122] L. Carbone et al. Sensors and Actuators for the Advanced LIGO Mirror Suspensions. *Class. Quant. Grav.*, 29:115005, 2012.
- [123] G. Ciani et al. Small optic suspensions for Advanced LIGO input optics and other precision optical experiments. *Rev. Sci. Instrum.*, 87(11):114504, 2016.
- [124] Norna A Robertson, Derek Bridges, Mark Barton, Jay Heefner, Janeen Romie, Calum Torrie, and Jeff Kissel. HAM Large Triple Suspension (HLTS) Final Design Document. Technical Report DCC-T1000012, LIGO, 10 2013.
- [125] Hang Yu. *Astrophysical signatures of neutron stars in compact binaries and experimental improvements on gravitational-wave detectors*. PhD thesis, Massachusetts Institute of Technology, 2019.
- [126] Elenna Capote and Craig Cahillane. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=67069>.
- [127] Gabriele Vajente, Yiwen Huang, Maximiliano Isi, Jenne C. Driggers, Jeffrey S. Kissel, Marek J. Szczepanczyk, and Salvatore Vitale. Machine-learning nonstationary noise out of gravitational-wave detectors. *Phys. Rev. D*, 101(4):042003, 2020.

- [128] Gabriele Vajente. Subtraction of non-stationary noise couplings. Technical Report LIGO-T1800525, LIGO, 1 2019.
- [129] Gabriele Vajente. <https://git.ligo.org/gabriele-vajente/nonsens>.
- [130] Gabriele Vajente. Training NonSENS models with offline data for online feed-forward. Technical Report LIGO-T2300064, LIGO, 2 2023.
- [131] Aidan F. Brooks, Gabriele Vajente, Hiro Yamamoto, Rich Abbott, Carl Adams, Rana X. Adhikari, Alena Ananyeva, Stephen Appert, Koji Arai, Joseph S. Areeda, Yasmeen Asali, Stuart M. Aston, Corey Austin, Anne M. Baer, Matthew Ball, Stefan W. Ballmer, Sharan Banagiri, David Barker, Lisa Barsotti, Jeffrey Bartlett, Beverly K. Berger, Joseph Betzwieser, Dripta Bhattacharjee, Garilynn Billingsley, Sebastien Biscans, Carl D. Blair, Ryan M. Blair, Nina Bode, Phillip Booker, Rolf Bork, Alyssa Bramley, Daniel D. Brown, Aaron Buikema, Craig Cahillane, Kipp C. Cannon, Huy Tuong Cao, Xu Chen, Alexei A. Ciobanu, Filiberto Clara, Camilla Compton, Sam J. Cooper, Kenneth R. Corley, Stefan T. Countryman, Pep B. Covas, Dennis C. Coyne, Laurence E. Datrier, Derek Davis, Chiara D. Difronzo, Katherine L. Dooley, Jenne C. Driggers, Peter Dupej, Sheila E. Dwyer, Anamaria Effler, Todd Etzel, Matthew Evans, Tom M. Evans, Jon Feicht, Alvaro Fernandez-Galiana, Peter Fritschel, Valery V. Frolov, Paul Fulda, Michael Fyffe, Joe A. Giaime, Dwayne D. Giardina, Patrick Godwin, Evan Goetz, Slawomir Gras, Corey Gray, Rachel Gray, Anna C. Green, Anchal Gupta, Eric K. Gustafson, Dick Gustafson, Evan Hall, Jonathan Hanks, Joe Hanson, Terra Hardwick, Raine K. Hasskew, Matthew C. Heintze, Adrian F. Helmling-Cornell, Nathan A. Holland, Kiamu Izmui, Wenxuan Jia, Jeff D. Jones, Shivaraj Kandhasamy, Sudarshan Karki, Marie Kasprzack, Keita Kawabe, Nutsinee Kijbunchoo, Peter J. King, Jeffrey S. Kissel, Rahul Kumar, Michael Landry, Benjamin B. Lane, Brian Lantz, Michael Laxen, Yannick K. Lecoeuuche, Jessica Leviton, Liu Jian, Marc Lormand, Andrew P. Lundgren, Ronaldas Macas, Myron Macinnis, Duncan M. Macleod, Georgia L. Mansell, Szabolcs Marka, Zsuzsanna Marka, Denis V. Martynov, Ken Mason, Thomas J. Massinger, Fabrice Matchard, Nergis Mavalvala, Richard McCarthy, David E. McClelland, Scott McCormick, Lee McCuller, Jessica McIver, Terry McRae, Gregory Mendell, Kara Merfeld, Edmond L. Merilh, Fabian Meylahn, Timesh Mistry, Richard Mittleman, Gerardo Moreno, Conor M. Mow-Lowry, Simone Mozzon, Adam Mullavey, Timothy J. Nelson, Philippe Nguyen, Laura K. Nuttall, Jason Oberling,

Richard J. Oram, Charles Osthelder, David J. Ottaway, Harry Overmier, Jordan R. Palamos, William Parker, Ethan Payne, Arnaud Pele, Reilly Penhorwood, Carlos J. Perez, Marc Pirello, Hugh Radkins, Karla E. Ramirez, Jonathan W. Richardson, Keith Riles, Norna A. Robertson, Jameson G. Rollins, Chandra L. Romel, Janeen H. Romie, Michael P. Ross, Kyle Ryan, Travis Sadecki, Eduardo J. Sanchez, Luis E. Sanchez, Saravanan R. Tiruppatturajamanikkam, Richard L. Savage, Dean Schaetzel, Roman Schnabel, Robert M. Schofield, Eyal Schwartz, Danny Sellers, Thomas Shaffer, Daniel Sigg, Bram J. Slagmolen, Joshua R. Smith, Siddharth Soni, Borja Sorazu, Andrew P. Spencer, Ken A. Strain, Ling Sun, Marek J. Szczepanczyk, Michael Thomas, Patrick Thomas, Keith A. Thorne, Karl Toland, Calum I. Torrie, Gary Traylor, Maggie Tse, Alexander L. Urban, Guillermo Valdes, Daniel C. Vander-Hyde, Peter J. Veitch, Krishna Venkateswara, Gautam Venugopalan, Aaron D. Viets, Thomas Vo, Cheryl Vorvick, Madeline Wade, Robert L. Ward, Jim Warner, Betsy Weaver, Rainer Weiss, Chris Whittle, Benno Willke, Christopher C. Wipf, Liting Xiao, Hang Yu, Haocun Yu, Liyuan Zhang, Michael E. Zucker, and John Zweizig. Point absorbers in Advanced LIGO. *Applied Optics*, 60(13):4047–4063, May 2021.

- [132] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=69082>.
- [133] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=76312>.
- [134] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=76363>.
- [135] John A. Sidles and Daniel Sigg. Optical torques in suspended Fabry–Perot interferometers. *Phys. Lett. A*, 354:167–172, 2006.
- [136] Edgard Bonilla. Technical note on bhqs sigg-sidles modes. Technical Report LIGO-T2300150, LIGO, 2023.
- [137] Craig Cahillane. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=58327>.
- [138] Jeffrey Kissel. Suspensions controls design summary table. Technical Report LIGO-G1100968, LIGO Laboratory, 2011.
- [139] LIGO-Virgo-KAGRA Collaboration. Gravitational-wave candidate event database. <https://gracedb.ligo.org/superevents/public/>, 2024.

- [140] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, (3):031040, 2019.
- [141] Henning Vahlbruch, Simon Chelkowski, Boris Hage, Alexander Franzen, Karsten Danzmann, and Roman Schnabel. Coherent control of vacuum squeezing in the gravitational-wave detection band. *Phys. Rev. Lett.*, 97:011101, Jul 2006.
- [142] Chris Whittle, Kentaro Komori, Dhruva Ganapathy, Lee McCuller, Lisa Barsotti, Nergis Mavalvala, and Matthew Evans. Optimal detuning for quantum filter cavities. *Phys. Rev. D*, 102(10):102002, November 2020.
- [143] Jameson Graef Rollins. Distributed state machine supervision for long-baseline gravitational-wave detectors. *Review of Scientific Instruments*, 87(9):094502, 09 2016.
- [144] Carlton Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, 23(8):1693–1708, April 1981.
- [145] M. Tse, Haocun Yu, N. Kijbunchoo, A. Fernandez-Galiana, P. Dupej, L. Barsotti, C. D. Blair, D. D. Brown, S. E. Dwyer, A. Effler, M. Evans, P. Fritschel, V. V. Frolov, A. C. Green, G. L. Mansell, F. Matichard, N. Mavalvala, D. E. McClelland, L. McCuller, T. McRae, J. Miller, A. Mullavey, E. Oelker, I. Y. Phinney, D. Sigg, B. J. J. Slagmolen, T. Vo, R. L. Ward, C. Whittle, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, J. S. Areeda, Y. Asali, S. M. Aston, C. Austin, A. M. Baer, M. Ball, S. W. Ballmer, S. Banagiri, D. Barker, J. Bartlett, B. K. Berger, J. Betzwieser, D. Bhattacharjee, G. Billingsley, S. Biscans, R. M. Blair, N. Bode, P. Booker, R. Bork, A. Bramley, A. F. Brooks, A. Buikema, C. Cahillane, K. C. Cannon, X. Chen, A. A. Ciobanu, F. Clara, S. J. Cooper, K. R. Corley, S. T. Countryman, P. B. Covas, D. C. Coyne, L. E. H. Datrier, D. Davis, C. Di Fronzo, J. C. Driggers, T. Etzel, T. M. Evans, J. Feicht, P. Fulda, M. Fyffe, J. A. Giaime, K. D. Giardina, P. Godwin, E. Goetz, S. Gras, C. Gray, R. Gray, Anchal Gupta, E. K. Gustafson, R. Gustafson, J. Hanks, J. Hanson, T. Hardwick, R. K. Hasskew, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, J. D. Jones, S. Kandhasamy, S. Karki, M. Kasprzack, K. Kawabe, P. J. King, J. S. Kissel, Rahul Kumar, M. Landry, B. B. Lane, B. Lantz, M. Laxen, Y. K. Lecoeuche, J. Leviton, J. Liu, M. Lormand, A. P. Lundgren, R. Macas,

- M. MacInnis, D. M. Macleod, S. Márka, Z. Márka, D. V. Martynov, K. Mason, T. J. Massinger, R. McCarthy, S. McCormick, J. McIver, G. Mendell, K. Merfeld, E. L. Merilth, F. Meylahn, T. Mistry, R. Mittleman, G. Moreno, C. M. Mow-Lowry, S. Mozzon, T. J. N. Nelson, P. Nguyen, L. K. Nuttall, J. Oberling, R. J. Oram, B. O'Reilly, C. Osthelder, D. J. Ottaway, H. Overmier, J. R. Palamos, W. Parker, E. Payne, A. Pele, C. J. Perez, M. Pirello, H. Radkins, K. E. Ramirez, J. W. Richardson, K. Riles, N. A. Robertson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, E. J. Sanchez, L. E. Sanchez, T. R. Saravanan, R. L. Savage, D. Schaetzl, R. Schnabel, R. M. S. Schofield, E. Schwartz, D. Sellers, T. J. Shaffer, J. R. Smith, S. Soni, B. Sorazu, A. P. Spencer, K. A. Strain, L. Sun, M. J. Szczepańczyk, M. Thomas, P. Thomas, K. A. Thorne, K. Toland, C. I. Torrie, G. Traylor, A. L. Urban, G. Vajente, G. Valdes, D. C. Vander-Hyde, P. J. Veitch, K. Venkateswara, G. Venugopalan, A. D. Viets, C. Vorvick, M. Wade, J. Warner, B. Weaver, R. Weiss, B. Willke, C. C. Wipf, L. Xiao, H. Yamamoto, M. J. Yap, Hang Yu, L. Zhang, M. E. Zucker, and J. Zweizig. Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy. *Physical Review Letters*, 123(23):231107, December 2019.
- [146] S. Gras and M. Evans. Direct measurement of coating thermal noise in optical resonators. *Physical Review D: Particles and Fields*, 98(12):122001, December 2018.
- [147] Wenxuan Jia, Hiroaki Yamamoto, Kevin Kuns, Anamaria Effler, Matthew Evans, Peter Fritschel, et al. Point absorber limits to future gravitational-wave detectors. *Phys. Rev. Lett.*, 127:241102, Dec 2021.
- [148] S. Biscans, S. Gras, C. D. Blair, J. Driggers, M. Evans, P. Fritschel, T. Hardwick, and G. Mansell. Suppressing parametric instabilities in ligo using low-noise acoustic mode dampers. *Phys. Rev. D*, 100:122003, Dec 2019.
- [149] B. Gustavsen and A. Semlyen. Rational approximation of frequency domain responses by vector fitting. *IEEE Transactions on Power Delivery*, 14(3):1052–1061, 1999.
- [150] <https://git.ligo.org/lee-mcculler/iirrational>.
- [151] B. Swinkels, E. Campagna, G. Vajente, L. Barsotti, and M. Evans. Longitudinal noise subtraction: the alpha-, beta- and gamma-technique. Technical report, Virgo Collaboration, 2008.

- [152] L. McCuller, S. E. Dwyer, A. C. Green, Haocun Yu, K. Kuns, L. Barsotti, C. D. Blair, D. D. Brown, A. Effler, M. Evans, A. Fernandez-Galiana, P. Fritschel, V. V. Frolov, N. Kijbunchoo, G. L. Mansell, F. Matichard, N. Mavalvala, D. E. McClelland, T. McRae, A. Mullavey, D. Sigg, B. J. J. Slagmolen, M. Tse, T. Vo, R. L. Ward, C. Whittle, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, J. S. Areeda, Y. Asali, S. M. Aston, C. Austin, A. M. Baer, M. Ball, S. W. Ballmer, S. Banagiri, D. Barker, J. Bartlett, B. K. Berger, J. Betzwieser, D. Bhattacharjee, G. Billingsley, S. Biscans, R. M. Blair, N. Bode, P. Booker, R. Bork, A. Bramley, A. F. Brooks, A. Buikema, C. Cahillane, K. C. Cannon, X. Chen, A. A. Ciobanu, F. Clara, C. M. Compton, S. J. Cooper, K. R. Corley, S. T. Countryman, P. B. Covas, D. C. Coyne, L. E. H. Datrier, D. Davis, C. Di Fronzo, K. L. Dooley, J. C. Driggers, T. Etzel, T. M. Evans, J. Feicht, P. Fulda, M. Fyffe, J. A. Giaime, K. D. Giardina, P. Godwin, E. Goetz, S. Gras, C. Gray, R. Gray, E. K. Gustafson, R. Gustafson, J. Hanks, J. Hanson, T. Hardwick, R. K. Hasskew, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, J. D. Jones, S. Kandhasamy, S. Karki, M. Kasprzack, K. Kawabe, P. J. King, J. S. Kissel, Rahul Kumar, M. Landry, B. B. Lane, B. Lantz, M. Laxen, Y. K. Lecoecuche, J. Leviton, J. Liu, M. Lormand, A. P. Lundgren, R. Macas, M. MacInnis, D. M. Macleod, S. Márka, Z. Márka, D. V. Martynov, K. Mason, T. J. Massinger, R. McCarthy, S. McCormick, J. McIver, G. Mendell, K. Merfeld, E. L. Merilh, F. Meylahn, T. Mistry, R. Mittleman, G. Moreno, C. M. Mow-Lowry, S. Mozzon, T. J. N. Nelson, P. Nguyen, L. K. Nuttall, J. Oberling, Richard J. Oram, C. Osthelder, D. J. Ottaway, H. Overmier, J. R. Palamos, W. Parker, E. Payne, A. Pele, R. Penhorwood, C. J. Perez, M. Pirello, H. Radkins, K. E. Ramirez, J. W. Richardson, K. Riles, N. A. Robertson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, E. J. Sanchez, L. E. Sanchez, T. R. Saravanan, R. L. Savage, D. Schaetzl, R. Schnabel, R. M. S. Schofield, E. Schwartz, D. Sellers, T. Shaffer, J. R. Smith, S. Soni, B. Sorazu, A. P. Spencer, K. A. Strain, L. Sun, M. J. Szczepańczyk, M. Thomas, P. Thomas, K. A. Thorne, K. Toland, C. I. Torrie, G. Traylor, A. L. Urban, G. Vajente, G. Valdes, D. C. Vander-Hyde, P. J. Veitch, K. Venkateswara, G. Venugopalan, A. D. Viets, C. Vorvick, M. Wade, J. Warner, B. Weaver, R. Weiss, B. Willke, C. C. Wipf, L. Xiao, H. Yamamoto, Hang Yu, L. Zhang, M. E. Zucker, and J. Zweizig. LIGO's quantum response to squeezed states. *Physical Review D: Particles and Fields*, 104(6):062006, September 2021.

- [153] Haocun Yu, L. McCuller, M. Tse, N. Kijbunchoo, L. Barsotti, N. Mavalvala, J. Betzwieser, C. D. Blair, S. E. Dwyer, A. Effler, M. Evans, A. Fernandez-Galiana, P. Fritschel, V. V. Frolov, F. Matichard, D. E. McClelland, T. McRae, A. Mullavey, D. Sigg, B. J. J. Slagmolen, C. Whittle, A. Buikema, Y. Chen, T. R. Corbitt, R. Schnabel, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, J. S. Areeda, Y. Asali, S. M. Aston, C. Austin, A. M. Baer, M. Ball, S. W. Ballmer, S. Banagiri, D. Barker, J. Bartlett, B. K. Berger, D. Bhattacharjee, G. Billingsley, S. Biscans, R. M. Blair, N. Bode, P. Booker, R. Bork, A. Bramley, A. F. Brooks, D. D. Brown, C. Cahillane, K. C. Cannon, X. Chen, A. A. Ciobanu, F. Clara, S. J. Cooper, K. R. Corley, S. T. Countryman, P. B. Covas, D. C. Coyne, L. E. H. Datrier, D. Davis, C. Di Fronzo, K. L. Dooley, J. C. Driggers, P. Dupej, T. Etzel, T. M. Evans, J. Feicht, P. Fulda, M. Fyffe, J. A. Giaime, K. D. Giardina, P. Godwin, E. Goetz, S. Gras, C. Gray, R. Gray, A. C. Green, Anchal Gupta, E. K. Gustafson, R. Gustafson, J. Hanks, J. Hanson, T. Hardwick, R. K. Hasskew, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, J. D. Jones, S. Kandhasamy, S. Karki, M. Kasprzack, K. Kawabe, P. J. King, J. S. Kissel, Rahul Kumar, M. Landry, B. B. Lane, B. Lantz, M. Laxen, Y. K. Lecoecuche, J. Leviton, J. Liu, M. Lormand, A. P. Lundgren, R. Macas, M. MacInnis, D. M. Macleod, G. L. Mansell, S. Márka, Z. Márka, D. V. Martynov, K. Mason, T. J. Massinger, R. McCarthy, S. McCormick, J. McIver, G. Mendell, K. Merfeld, E. L. Merilh, F. Meylahn, T. Mistry, R. Mittleman, G. Moreno, C. M. Mow-Lowry, S. Mozzon, T. J. N. Nelson, P. Nguyen, L. K. Nuttall, J. Oberling, Richard J. Oram, C. Osthelder, D. J. Ottaway, H. Overmier, J. R. Palamos, W. Parker, E. Payne, A. Pele, C. J. Perez, M. Pirello, H. Radkins, K. E. Ramirez, J. W. Richardson, K. Riles, N. A. Robertson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, E. J. Sanchez, L. E. Sanchez, T. R. Saravanan, R. L. Savage, D. Schaetzl, R. M. S. Schofield, E. Schwartz, D. Sellers, T. Shaffer, J. R. Smith, S. Soni, B. Sorazu, A. P. Spencer, K. A. Strain, L. Sun, M. J. Szczepańczyk, M. Thomas, P. Thomas, K. A. Thorne, K. Toland, C. I. Torrie, G. Traylor, A. L. Urban, G. Vajente, G. Valdes, D. C. Vander-Hyde, P. J. Veitch, K. Venkateswara, G. Venugopalan, A. D. Viets, T. Vo, C. Vorvick, M. Wade, R. L. Ward, J. Warner, B. Weaver, R. Weiss, B. Willke, C. C. Wipf, L. Xiao, H. Yamamoto, Hang Yu, L. Zhang, M. E. Zucker, J. Zweizig, and members of the LIGO Scientific Collaboration. Quantum correlations between light and the kilogram-mass mirrors of LIGO. *Nature*, 583(7814):43–47, July 2020.

- [154] W. Jia et al. Ligo operates with quantum noise below the standard quantum limit. *in preparation*.
- [155] Varun Srivastava, Georgia Mansell, Camille Makarem, Minkyun Noh, Richard Abbott, Stefan Ballmer, GariLynn Billingsley, Aidan Brooks, Huy Tuong Cao, Peter Fritschel, Don Griffith, Wenxuan Jia, Marie Kasprzack, Myron MacInnis, Sebastian Ng, Luis Sanchez, Calum Torrie, Peter Veitch, and Fabrice Matichard. Piezo-deformable mirrors for active mode matching in advanced LIGO. *Opt. Express, OE*, 30(7):10491–10501, March 2022.
- [156] Huy Tuong Cao, Sebastian W. S. Ng, Minkyun Noh, Aidan Brooks, Fabrice Matichard, and Peter J. Veitch. Enhancing the dynamic range of deformable mirrors with compression bias. *Opt. Express, OE*, 28(26):38480–38490, December 2020.
- [157] Huy Tuong Cao, Aidan Brooks, Sebastian W. S. Ng, David Ottaway, Antonio Perreca, Jonathan W. Richardson, Aria Chaderjian, and Peter J. Veitch. High dynamic range thermally actuated bimorph mirror for gravitational wave detectors. *Appl. Opt., AO*, 59(9):2784–2790, March 2020.
- [158] L McCuller and L Barsotti. Design Requirement Document of the A+ filter cavity and relay optics for frequency dependent squeezing. *Tech. Report No. LIGO-T1800447*.
- [159] A. Viets and M. Wade. Subtracting Narrow-band Noise from LIGO Strain Data in the Third Observing Run. Technical Report DCC-T2100058, LIGO, 8 2021.
- [160] Gabriele Vajente. Data mining and machine learning improve gravitational-wave detector sensitivity. *Phys. Rev. D*, 105:102005, May 2022.
- [161] M. Evans, L. Barsotti, and P. Fritschel. A general approach to optomechanical parametric instabilities. *Phys. Lett. A*, 374(4):665–671, 2010.
- [162] Matthew Evans, Slawek Gras, Peter Fritschel, John Miller, Lisa Barsotti, Denis Martynov, Aidan Brooks, Dennis Coyne, Rich Abbott, Rana X. Adhikari, Koji Arai, Rolf Bork, Bill Kells, Jameson Rollins, Nicolas Smith-Lefebvre, Gabriele Vajente, Hiroaki Yamamoto, Carl Adams, Stuart Aston, Joseph Betzweiser, Valera Frolov, Adam Mullavey, Arnaud Pele, Janeen Romie, Michael Thomas, Keith Thorne, Sheila Dwyer, Kiwamu Izumi, Keita Kawabe, Daniel Sigg, Ryan Derosa, Anamaria Effler, Keiko

- Kokeyama, Stefan Ballmer, Thomas J. Massinger, Alexa Staley, Matthew Heinze, Chris Mueller, Hartmut Grote, Robert Ward, Eleanor King, David Blair, Li Ju, and Chunnong Zhao. Observation of parametric instability in advanced ligo. *Phys. Rev. Lett.*, 114:161102, Apr 2015.
- [163] Carl Blair, Slawek Gras, et al. First demonstration of electrostatic damping of parametric instability at advanced ligo. *Phys. Rev. Lett.*, 118:151102, Apr 2017.
- [164] Hartmut Grote, Michael Weinert, Rana X. Adhikari, Christoph Affeldt, Volker Kringel, Jonathan Leong, James Lough, Harald Lück, Emil Schreiber, Kenneth A. Strain, Henning Vahlbruch, and Holger Wittel. High power and ultra-low-noise photodetector for squeezed-light enhanced gravitational wave detectors. *Opt. Express*, 24(18):20107–20118, Sep 2016.
- [165] Aidan F. Brooks, Benjamin Abbott, Muzammil A. Arain, Giacomo Ciani, Ayodele Cole, Greg Grabeel, Eric Gustafson, Chris Guido, Matthew Heintze, Alastair Heptonstall, Mindy Jacobson, Won Kim, Eleanor King, Alexander Lynch, Stephen O’Connor, David Ottaway, Ken Mailand, Guido Mueller, Jesper Munch, Virginio Sannibale, Zhenhua Shao, Michael Smith, Peter Veitch, Thomas Vo, Cheryl Vorvick, and Phil Willems. Overview of advanced ligo adaptive optics. *Appl. Opt.*, 55(29):8256–8265, Oct 2016.
- [166] T. Hardwick, V. J. Hamedan, C. Blair, A. C. Green, and D. Vander-Hyde. Demonstration of dynamic thermal compensation for parametric instability suppression in Advanced LIGO. *Class. Quant. Grav.*, 37(20):205021, 2020.
- [167] Elenna Capote, Louis Dartez, and Derek Davis. Technical Noise, Data Quality, and Calibration Requirements for Next-Generation Gravitational-Wave Science. 4 2024.
- [168] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.
- [169] R. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *Phys. Rev. X*, 11:021053, 2021.

- [170] R. Abbott et al. GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. 8 2021. arXiv:2108.01045.
- [171] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run. *Phys. Rev. X*, 13(4):041039, 2023.
- [172] F. Acernese et al. Virgo detector characterization and data quality: results from the O3 run. *Class. Quant. Grav.*, 40(18):185006, 2023.
- [173] H. Abe et al. Performance of the KAGRA detector during the first joint observation with GEO 600 (O3GK). *PTEP*, 2023(10):10A101, 2023.
- [174] T. Akutsu et al. Overview of KAGRA: Calibration, detector characterization, physical environmental monitors, and the geophysics interferometer. *PTEP*, 2021(5):05A102, 2021.
- [175] B. P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quant. Grav.*, 33(13):134001, 2016.
- [176] L. K. Nuttall. Characterising Transient Noise in the LIGO Detectors. *Phil. Trans. Roy. Soc. Lond.*, A376:20170286, 2018.
- [177] D. Davis et al. LIGO detector characterization in the second and third observing runs. *Class. Quant. Grav.*, 38(13):135014, 2021.
- [178] B. P. Abbott et al. A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals. *Class. Quant. Grav.*, 37(5):055002, 2020.
- [179] Derek Davis and Marissa Walker. Detector Characterization and Mitigation of Noise in Ground-Based Gravitational-Wave Interferometers. *Galaxies*, 10(1):12, 2022.
- [180] B P Abbott et al. Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO’s first observing run. *Class. Quant. Grav.*, 35(6):065010, 2018.

- [181] J. Powell. Parameter Estimation and Model Selection of Gravitational Wave Signals Contaminated by Transient Detector Noise Glitches. *Class. Quant. Grav.*, 35(15):155017, 2018.
- [182] J. Y. L. Kwok, K. L. Lo, A. J. Weinstein, and T. G. F. Li. Investigation of the effects of non-Gaussian noise transients and their mitigation in parameterized gravitational-wave tests of general relativity. *Phys. Rev. D*, 105(2):024066, 2022.
- [183] Simone Mozzon, Gregory Ashton, Laura K. Nuttall, and Andrew R. Williamson. Does nonstationary noise in LIGO and Virgo affect the estimation of H_0 ? *Phys. Rev. D*, 106(4):043504, 2022.
- [184] R. Macas, J. Pooley, L. K. Nuttall, D. Davis, M. J. Dyer, Y. Leconte, J. D. Lyman, J. McIver, and K. Rink. Impact of noise transients on low latency gravitational-wave event localization. *Phys. Rev. D*, 105(10):103021, 2022.
- [185] Sophie Hourihane, Katerina Chatziioannou, Marcella Wijngaarden, Derek Davis, Tyson Littenberg, and Neil Cornish. Accurate modeling and mitigation of overlapping signals and glitches in gravitational-wave data. *Phys. Rev. D*, 106(4):042006, 2022.
- [186] Sudarshan Ghonge, Joshua Brandt, J. M. Sullivan, Margaret Millhouse, Katerina Chatziioannou, James A. Clark, Tyson Littenberg, Neil Cornish, Sophie Hourihane, and Laura Cadonati. Characterizing the efficacy of methods to subtract terrestrial transient noise near gravitational wave events and the effects on parameter estimation. 11 2023.
- [187] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton. Subtracting glitches from gravitational-wave detector data during the third LIGO-Virgo observing run. *Class. Quant. Grav.*, 39(24):245013, 2022.
- [188] B. K. Berger et al. Searching for the causes of anomalous Advanced LIGO noise. *Applied Physics Letters*, 122(18):184101, May 2023.
- [189] Benjamin P. Abbott et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93(11):112004, 2016. [Addendum: *Phys.Rev.D* 97, 059901 (2018)].

- [190] H. Abe et al. Performance of the KAGRA detector during the first joint observation with GEO 600 (O3GK). *PTEP*, 2023(10):10A101, 2023.
- [191] Varun Srivastava, Derek Davis, Kevin Kuns, Philippe Landry, Stefan Ballmer, Matthew Evans, Evan D. Hall, Jocelyn Read, and B. S. Sathyaprakash. Science-driven Tunable Design of Cosmic Explorer Detectors. *Astrophys. J.*, 931(1):22, 2022.
- [192] Grant David Meadors, Keita Kawabe, and Keith Riles. Increasing LIGO sensitivity by feedforward subtraction of auxiliary length control noise. *Class. Quant. Grav.*, 31:105014, 2014.
- [193] S. J. Cooper, A. C. Green, C. J. Collins, D. Hoyland, C. C. Speake, A. Freise, and C. M. Mow-Lowry. A compact, large-range interferometer for precision measurement and inertial sensing. *Class. Quant. Grav.*, 35(9):095007, 2018.
- [194] Marie Kasprzack. ASC noise budgeting for LLO. Technical Report DCC-G2100751, LIGO, 2021.
- [195] Riccardo Maggiore et al. Angular control noise in Advanced Virgo and implications for the Einstein Telescope. 1 2024.
- [196] K Kuns et al. Cosmic Explorer strain sensitivity. Technical Report CE-T2000017, Cosmic Explorer, 2020.
- [197] J. Rollins and J. Creighton. inspiral-range: version 0.9.1. <https://git.ligo.org/gwinc/inspiral-range/tree/0.9.1>.
- [198] Geraint Pratten et al. Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes. *Phys. Rev. D*, 103(10):104056, 2021.
- [199] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
- [200] Thibault Damour, Alessandro Nagar, and Loic Villain. Measurability of the tidal polarizability of neutron stars in late-inspiral gravitational-wave signals. *Phys. Rev. D*, 85:123007, 2012.

- [201] Man Leong Chan, Chris Messenger, Ik Siong Heng, and Martin Hendry. Binary Neutron Star Mergers and Third Generation Detectors: Localization and Early Warning. *Phys. Rev. D*, 97(12):123014, 2018.
- [202] Yufeng Li, Ik Siong Heng, Man Leong Chan, Chris Messenger, and Xilong Fan. Exploring the sky localization and early warning capabilities of third generation gravitational wave detectors in three-detector network configurations. *Phys. Rev. D*, 105(4):043010, 2022.
- [203] Ssohrab Borhanian and B. S. Sathyaprakash. Listening to the Universe with Next Generation Ground-Based Gravitational-Wave Detectors. 2 2022.
- [204] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, 1991.
- [205] Alexander M. Grant and David A. Nichols. Outlook for detecting the gravitational-wave displacement and spin memory effects with current and future gravitational-wave detectors. *Phys. Rev. D*, 107(6):064056, 2023. [Erratum: *Phys.Rev.D* 108, 029901 (2023)].
- [206] Amber K. Lenon, Duncan A. Brown, and Alexander H. Nitz. Eccentric binary neutron star search prospects for Cosmic Explorer. *Phys. Rev. D*, 104(6):063011, 2021.
- [207] Pankaj Saini. Resolving the eccentricity of stellar mass binary black holes with next generation gravitational wave detectors. 8 2023.
- [208] P. B. Covas et al. Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97(8):082002, 2018.
- [209] R. Abbott et al. All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems. *Phys. Rev. D*, 103(6):064017, 2021. [Erratum: *Phys.Rev.D* 108, 069901 (2023)].
- [210] Benjamin P. Abbott et al. All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data. *Phys. Rev. D*, 96(6):062002, 2017.

- [211] Rodrigo Tenorio, David Keitel, and Alicia M. Sintes. Search Methods for Continuous Gravitational-Wave Signals from Unknown Sources in the Advanced-Detector Era. *Universe*, 7(12):474, 2021.
- [212] Keith Riles. Searches for continuous-wave gravitational radiation. *Living Rev. Rel.*, 26(1):3, 2023.
- [213] Karl Wette. Searches for continuous gravitational waves from neutron stars: A twenty-year retrospective. *Astropart. Phys.*, 153:102880, 2023.
- [214] R N Manchester, G B Hobbs, A Teoh, and M Hobbs. The Australia Telescope National Facility pulsar catalogue. *Astron. J.*, 129:1993, 2005.
- [215] ATNF pulsar catalogue, v1.71. <https://www.atnf.csiro.au/research/pulsar/psrcat/index.html?version=1.71>. Accessed: 2024-01-08.
- [216] R. Abbott et al. Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO. *Astrophys. J. Suppl.*, 267(2):29, 2023.
- [217] O3 instrumental lines. <https://gwosc.org/O3/o3speclines/>. Accessed: 2024-01-08.
- [218] D. Tuyenbayev et al. Improving LIGO calibration accuracy by tracking and compensating for slow temporal variations. *Class. Quant. Grav.*, 34(1):015002, 2017.
- [219] Aaron Viets et al. Reconstructing the calibrated strain signal in the Advanced LIGO detectors. *Class. Quant. Grav.*, 35(9):095015, 2018.
- [220] Reed Essick. Calibration uncertainty’s impact on gravitational-wave observations. *Phys. Rev. D*, 105(8):082002, 2022.
- [221] Craig Cahillane et al. Calibration uncertainty for Advanced LIGO’s first and second observing runs. *Phys. Rev.*, D96(10):102001, 2017.
- [222] Ling Sun et al. Characterization of systematic error in Advanced LIGO calibration. *Class. Quant. Grav.*, 37(22):225008, 2020.
- [223] T Starkman. aLIGO LHO Logbook. [71501](#), 2023.

- [224] N. A. Robertson, P. Fritschel, B. Shapiro, C. I. Torrie, and M. Evans. Design of a tuned mass damper for high quality factor suspension modes in Advanced LIGO. *Rev. Sci. Instrum.*, 88(3):035117, 2017.
- [225] J.C. Driggers et al. Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO. *Phys. Rev. D*, 99(4):042001, 2019.
- [226] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall. Improving the sensitivity of advanced ligo using noise subtraction. *Class. Quant. Grav.*, 36(5):055011, 2019.
- [227] Gabriele Vajente, Yiwen Huang, Maximiliano Isi, Jenne C. Driggers, Jeffrey S. Kissel, Marek J. Szczepanczyk, and Salvatore Vitale. Machine-learning nonstationary noise out of gravitational-wave detectors. *Phys. Rev. D*, 101(4):042003, 2020.
- [228] H. Yu and R. X. Adhikari. Nonlinear Noise Cleaning in Gravitational-Wave Detectors With Convolutional Neural Networks. *Front. Artif. Intell.*, 5:811563, 2022.
- [229] B. P. Abbott et al. LIGO: The Laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.*, 72:076901, 2009.
- [230] Simon Tait. *Studies of mechanical and optical properties of thin film coatings for future gravitational wave detectors*. PhD thesis, University of Glasgow, 2021.

Elenna Capote

Doctoral Candidate in Physics at Syracuse University

✉ ecapote@syr.edu ✉ elenna.capote@ligo.org

EDUCATION

Syracuse University <i>Ph.D. in Physics</i> <i>Research Advisor: Stefan Ballmer</i>	Current <i>June 2024</i>
California State University Fullerton <i>M.S. Physics</i> <i>Research Advisor: Joshua Smith</i>	May 2020
Massachusetts Institute of Technology <i>S.B. Physics</i> <i>Concentration in Particle-Astrophysics</i>	Feb 2019

AWARDS, FELLOWSHIPS, AND RESEARCH SUPPORT

NSF-KAGRA Exchange Program <i>Institute for Cosmic Ray Research, University of Tokyo</i>	Sept.-Nov. 2023
OzGrav Visitor Program <i>University of Adelaide</i>	July-Aug. 2023
LIGO Scientific Collaboration Fellows Program <i>LIGO Hanford Observatory</i>	Feb. 2022-Aug. 2023
STEM Graduate Fellowship <i>Syracuse University Physics Department</i> \$56,000	2020, 2022
PAARE Graduate Fellowship <i>Syracuse University Physics Department</i> \$34,000	2021
Social Justice Award <i>Syracuse University Physics Department</i> \$200	2021
Outstanding Master's Student Award <i>CSUF Physics Department</i> \$300	2020
Nancy Goodhue-McWilliams Fellowship <i>Gravitational Wave Physics and Astronomy Center, CSUF</i> \$10,000	2019,2020

PUBLICATIONS

Short Author Publications

- In-Vacuum Measurements of Optical Scatter Versus Annealing Temperature for Amorphous Ta₂O₅ and TiO₂:Ta₂O₅ Thin Films. Capote, et al. JOSA A 38 pp. 534-541 (2021).
<https://arxiv.org/abs/2011.14013> <https://opg.optica.org/josaa/abstract.cfm?uri=josaa-38-4-534>
(refereed)

Capote, CV

- Technical Noise, Data Quality, and Calibration Requirements for Next-Generation Gravitational-Wave Science. Capote, Darte, and Davis. (2024) *Submitted to Class. Quantum Grav.* <https://arxiv.org/abs/2404.04761>
- Sensitivity and performance of the Advanced LIGO detectors in the Fourth Observing Run. Capote, Jia, Aritomi, Nakano, Xu, et al. *In Preparation*
- Cryogenic Mechanical Loss Studies of Large-Area Substrate-Transferred GaAs/AlGaAs Crystalline Coatings. Didio, Tanioka, Capote, et al. *In Preparation*

Collaboration Publications

- GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. The LIGO Scientific Collaboration, the Virgo Collaboration. (2021) Submitted to Phys. Rev. D. - <https://arxiv.org/abs/2108.01045>
- GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration. (2021) Submitted to Phys. Rev. X. - <http://arxiv.org/abs/2111.03606>
- I am also an author on 29 additional papers as a member of the LIGO Scientific Collaboration.

RESEARCH EXPERIENCE

Fellow, LIGO Hanford Observatory | Feb. 2022 – Aug. 2023
LIGO Scientific Collaboration Fellows Program

- Key member of the advanced gravitational wave interferometer commissioning efforts to improve sensitivity and duty cycle
- Commissioned high power operation, achieving highest operating power in a gravitational wave interferometer
- Modeled and developed improved controls to increase low-frequency detector sensitivity
- Prototyped new control methods for future observing runs

Research Assistant, Syracuse University | Aug. 2020 – Feb 2022
Department of Physics

- Measured mechanical loss of single-layer amorphous and crystalline optical coatings using a gentle nodal suspension experiment
- Designed analysis methods for multi-modal ringdown measurements implementing digital filtering and a heterodyne method
- Quantified coating thermal noise by fitting substrate and coating loss measurements

Research Assistant, CSUF | June 2019 – Aug. 2020
Gravitational Wave Physics and Astronomy Center

- Measured scattered light from sample optics to test coating properties of current and future LIGO mirror coatings
- Collaborated with physics and engineering students to develop in-vacuum annealing experiment
- Performed data acquisition and analysis via LabVIEW programs and MATLAB scripts

MIT International Science and Technology Fellow | June – Aug. 2013
Department of Astrophysics, Universidad Catolica de Chile

- Searched for Classical Cepheids in the VISTA Variables survey of the Milky Way
- Classified Ks-band light curves in the variable star database
- Calculated distance of Cepheids to better map 3-D structure of the Milky Way

TEACHING EXPERIENCE AND OUTREACH

Teaching Associate | *Department of Physics, CSUF* 2019-2020

- Guided students through astronomy lab projects related to coordinate systems, seasons, moon phases, solar system models, light spectra, telescopes, HR diagrams and Hubble's law
- Prepared and led mini-lectures to prepare students to complete lab work in class, wrote quizzes
- Led outdoor observation sessions
- Graded lab reports and in-class quizzes

High School and College Tutor | *North County San Diego, CA* 2016-2019

- Met with high school and college students one-on-one for subject and test preparation tutoring
- Taught students in all levels of mathematics, physics, chemistry, biology, and Spanish
- Customized personal sessions with students to best suit their learning style and needs (homework help, study habits, test preparation)
- Prepared students for standardized tests such as the ACT or SAT

IYPT-Chile Physics Teacher and Program Coordinator | *Instituto Nacional, Santiago, Chile* June-Aug. 2015

- Planned and implemented an experimental physics and presentation skills curriculum for a 6-week program
- Advised a team of five students in three different physics experiments

Outreach Tour Guide | *LIGO Hanford Observatory* 2022-2023

- Led tours of gravitational wave observatory for general public

SELECTED CONFERENCE TALKS AND SEMINARS

- The Sensitivity and Performance of Current and Future Ground-Based Gravitational-Wave Detectors - Gravitational Wave Physics and Astronomy Workshop, University of Birmingham, UK, May 2024 - **Invited**
- The Performance of the Advanced LIGO Detectors during the Fourth Observing Run - American Physical Society Conference, Sacramento, CA, April 2024 - **Invited**
- Modeling Alignment Sensing and Control for Advanced LIGO Observing Run 4 - Caltech LIGO Seminar, Pasadena, CA, Aug. 2023 - **Invited**
- Advanced LIGO Detector Commissioning for Observing Run 4 - Gravitational Wave Advanced Detector Workshop, Elba, Italy May 2023 - **Invited**
- Advanced LIGO Detector Commissioning for Observing Run 4 - Gravitational Wave Research Exchange Meeting, University of Tokyo (remote), July 2023 - **Invited**
- Advanced LIGO Detector Commissioning for Observing Run 4 - Caltech LIGO Seminar, Pasadena, CA, June 2023 - **Invited**
- Advanced LIGO Detector Commissioning for Observing Run 4 - LIGO-Virgo-KAGRA Collaboration Conference Instrument Science Plenary, Evanston, IL, March 2023 - **Invited**
- Elastic Loss in Amorphous Zirconia-Titania Coatings - LIGO-Virgo-KAGRA Collaboration Conference Coatings Workshop, Evanston, IL, March 2023
- Elastic Loss in Amorphous Zirconia-Titania Coatings - Gravitational Wave Science Symposium, Padova, Italy, June 2022
- In-Vacuum Measurements of Optical Scatter Versus Annealing Temperature for Amorphous Ta₂O₅ and TiO₂:Ta₂O₅ Thin Films - American Physical Society Conference (remote), April 2020
- In-Vacuum Measurements of Optical Scatter Versus Annealing Temperature for Amorphous Ta₂O₅ and TiO₂:Ta₂O₅ Thin Films - LIGO-Virgo-KAGRA Collaboration Conference Optics Working Group (remote), March 2020

Capote, CV