

Mitigating Noise in Interferometric Gravitational Wave Detectors

Thesis by
William Zachary Korth

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

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

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2019
Defended February 12, 2019

© 2019

William Zachary Korth
ORCID: 0000-0002-4422-1070

All rights reserved

ACKNOWLEDGEMENTS

I have spent well over a third of my life in LIGO. It has been a long road, and one that's been as challenging as it has been rewarding. Over my years in graduate school, I can barely count the number of choices I made that prolonged the journey significantly. I regret none of them.

To Guido, Dave R. and Dave T., thank you for bringing me into this remarkable family, and for giving me the opportunity to be a part of something truly great.

To Eric B. and Ken, thanks for initiating me into the LIGO Lab and for an unforgettable summer at Caltech.

To Rana, thank you for believing that I had what it takes and for helping me remember when I forgot. You have been an inspiration to me from the beginning, as you are to so many others. It has been a privilege to work under you.

To Alan, Yanbei and Keith, thank you for your unwavering support through thick and thin. I could not have asked for a better committee. Also to Alan and Yanbei, thank you for all you have taught me, whether in a classroom or over a beer or three.

To Julie, Tess and Sofie, thank you for all your help with every form, deadline and submission over the years.

To Eric G., thank you for always having my back and for letting some of those regular reports slide. I will miss stopping by your office to talk about whatever strikes our fancy.

To Aidan, Alastair, Dmass, Jenne, Frank, Jan, Gabriele, Jamie, Nic, Chris W., Larry, Rory, Tara, Matt, Alberto, Andrew, Kiwamu, Jeff K., and everyone who came before me, thank you for all you've taught me; I will always cherish the time we spent together. To Eric, Evan, Den and those who came *after* me but still somehow got out before me, I forgive you!

To Aaron, Brittany and Johannes, thanks for sharing the Cryo lab with me for these past couple years, and for all your help in finally getting me out of it! To Gautam and Craig, it was a pleasure sharing a softball championship with you.

To Valera, Ryan, Chris M., Anamaria, Jan, Keiko, Adam, Den, Suresh, and the rest of the Swamp Donkeys, I think you already know how I felt about leaving Louisiana; I feel the same way every day. Thank you for being a part of making that year and

a half or so among the best of my life. To Thomas, thanks for letting me crash your lab for a while and for teaching me a great deal.

To Haixing, thank you for the countless hours we spent talking about quantum noise and about life. I will always be moved by your intellect and your philosophy.

To Rich, thanks for teaching me damn near all I know about electronics. I am forever grateful for the number of afternoons you spent with me at your whiteboard after I showed up unannounced. Thanks, too, for all the parts I pilfered from the Downs lab over the years.

To Koji, thank you for the sheer amount of knowledge, intuition, and general experimental wisdom you have bestowed upon me over the years. I would not be the person I am today without the constant support and guidance you have given me from my first day to my last.

To Ernie, thank you for feeding me! By my count, I probably visited your magical truck of deliciousness over 2000 times. Don't think for a second I won't be making occasional Thursday trips back to Pasadena just for a spicy beef fix!

To my beautiful wife, Gretchen, thank you for dragging me into an unheated pool in the dead of winter, then down a mountain, then out a perfectly good plane and into your life forever. Thanks for hopping in a car a few days after your defense and driving across the country to a new home, sight-unseen. Thanks for always believing in me and for everything you've done to get us where we are today. I can't wait to see where our next adventure takes us. I love you.

To my little Max, thanks for playing your own tiny part to make sure this thesis was finished in under a decade. We've only just met you, but you've already brought us so much joy. I love you, too.

To Tiny, Mufasa, Scotty, Peeps, Dory and Frankie, thanks for occasionally taking time out of your day—and mine—to remind me that sometimes snuggles are more important than really getting much of anything done.

Last but not least, to my incredibly loving family, thank you for making me the man I am. Thank you for giving me the tools I needed to succeed in this world, and the will to do it. I love you all so very much.

To the many people I'm sure to have forgotten here, thank you for being a part of my life and for letting the little things slide.

ABSTRACT

Gravitational waves, first predicted by Einstein in 1916, eluded detection for nearly a century. These faint ripples in the fabric of spacetime, with typical strain amplitudes at the Earth on the order of $|h| \sim 10^{-22}$, carry secrets of the universe untold by electromagnetic radiation. Following decades of research and development, a network of terrestrial interferometric detectors succeeded in measuring the passing of a gravitational wave ([GW150914](#)) for the first time in 2015. Individual detectors within this network are currently said to be operating in a “second-generation” configuration; over the next decade, planned upgrades will take these detectors beyond this into a new generation. This thesis concerns the characterization and reduction of noise in one of these second-generation detectors, Advanced LIGO, as well as efforts underway to improve its sensitivity in the coming years.

The first part of this thesis is a detailed overview of gravitational waves, the history of gravitational wave detection, and a reasonably thorough description of the Advanced LIGO detector. Particular attention is paid to a pedagogical motivation of the optical configuration of Advanced LIGO with reference to its forebears. This part ends with an overview of the sources of noise limiting the sensitivity of Advanced LIGO, and an exposition of plans to reduce their influence in the future.

The second part describes the development of a laser gyroscope for use in tilt sensing in Advanced LIGO, starting with a motivation of the work based on limitations in the area of seismic noise sensing and cancellation.

The third part recounts the design, fabrication, testing, installation and commissioning of an important component of the Advanced LIGO detector: the output mode cleaner (OMC).

The fourth part outlines a proposed scheme for reduction of quantum noise in gravitational wave detectors and other experiments. In particular, this scheme allows for the operation of a so-called “optical spring” cavity in such a way as to be largely immune from the deleterious effects of quantum radiation pressure noise.

The fifth and final part describes progress towards a direct measurement of thermal noise in thin silicon ribbons, which is pertinent to the design of suspensions in future cryogenic gravitational wave detectors.

This thesis has the internal LIGO document number [P1900035](#).

PUBLISHED CONTENT AND CONTRIBUTIONS

Korth, W. Z., A. Heptonstall, E. D. Hall, K. Arai, E. K. Gustafson, and R. X. Adhikari (2016). “Passive, free-space heterodyne laser gyroscope”. In: *Classical and Quantum Gravity* **33** 035004. doi: <https://doi.org/10.1088/0264-9381/33/3/035004>.

W.Z.K. participated in the conception of the experiment, conducted it, and wrote the paper.

Korth, W. Z., H. Miao, T. Corbitt, G. D. Cole, Y. Chen, and R. X. Adhikari (2013). “Suppression of quantum-radiation-pressure noise in an optical spring”. In: *Physical Review A* **88** 033805. doi: <https://doi.org/10.1103/PhysRevA.88.033805>.

W.Z.K. devised the concept with H. Miao and wrote the paper.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vi
Table of Contents	vii
List of Illustrations	xi
List of Tables	xv
I Background and Theory	1
Chapter I: Gravitational waves	2
1.1 Basic gravitational wave theory	2
1.1.1 General relativity	2
1.1.2 Gravitational waves	3
1.2 History of detection science	6
1.2.1 First estimates: Hertz-type experiment	6
1.2.2 Early detection attempts: Weber’s resonant bars	7
1.2.3 Astronomical observations	8
1.2.4 Laser interferometers	9
Chapter II: State of the art: Advanced LIGO	11
2.1 Description of the instrument	11
2.1.1 Optical configuration	11
2.1.2 Systems and control	29
2.2 Noise	38
2.2.1 Seismic and gravity-gradient (Newtonian) noise	39
2.2.2 Laser noise	40
2.2.3 Scattered light noise	41
2.2.4 Residual gas noise	41
2.2.5 Electronic noise	42
2.2.6 Quantum noise	42
2.2.7 Thermal noise	45
2.3 The path forward	46
2.3.1 Reducing noise in LIGO upgrades	46
2.3.2 Putting it all together: the case for silicon in LIGO Voyager	49
2.3.3 Outstanding questions about silicon	50
II Laser Gyroscope	52
Chapter III: Tilt sensing in Advanced LIGO	54
Chapter IV: Passive, free-space heterodyne laser gyroscope	57

4.1	Introduction	57
4.1.1	Laser gyroscopes	57
4.1.2	Current laser gyroscope sensitivities	58
4.1.3	Rotation sensing	59
4.2	The passive, free-space laser gyroscope	60
4.2.1	Overview	60
4.2.2	Detailed experimental design	62
4.2.3	Realized performance and noise analysis	65
4.3	Possible future designs	72
4.4	Conclusion	73
III Advanced LIGO Output Mode Cleaner		75
Chapter V: Background		76
5.1	DC Readout	76
5.2	Mode cleaners in detail	77
5.2.1	Transverse modes and Gaussian beams	77
5.2.2	Cavity transverse mode spacing and higher-order mode transmission spectrum	80
Chapter VI: The Advanced LIGO OMC		84
6.1	Optical design	84
6.2	Mechanical design, fabrication and benchtop testing	86
6.2.1	Overall construction and components	86
6.2.2	Mirror characterization and bench assembly	91
6.3	Installation	93
6.3.1	OMC suspension assembly	93
6.3.2	Detector electronics	94
6.4	Integration	98
6.4.1	Bring-up	98
6.4.2	Control system	99
6.4.3	Advanced LIGO commissioning with the OMC	103
IV Suppression of Quantum Radiation Pressure Noise in an Optical Spring		107
Chapter VII: Suppression of Quantum Radiation Pressure Noise in an Optical Spring		109
7.1	Introduction	109
7.2	Optical spring	111
7.3	Evading quantum radiation pressure noise	113
7.4	Residual radiation pressure noise	114
7.5	Experimental realization with double optical spring	116
7.6	Conclusion	117

V Thermal Noise in Silicon Ribbons	119
Chapter VIII: Thermal noise theory, resonator design, and numerical modeling	121
8.1 Thermal noise in detail	121
8.1.1 The loss angle	121
8.1.2 The fluctuation-dissipation theorem (again)	123
8.1.3 Sources of loss	123
8.2 Suspension thermal noise in LIGO Voyager	126
8.2.1 Background	127
8.2.2 The ribbon cross section	127
8.2.3 Suspension thermal noise	127
8.2.4 Conclusion	131
8.3 Ribbon cantilever design and numerical modeling	132
8.3.1 Design	132
8.3.2 Finite-element analysis (FEA) numerical model	134
Chapter IX: Cantilever fabrication	138
9.1 Fabrication process	138
9.1.1 Overview	138
9.1.2 Detailed recipe	138
Chapter X: Characterizing loss: Ringdown measurements	148
10.1 The apparatus	148
10.1.1 The clamp	148
10.1.2 The cryostat	149
10.2 Operation	150
10.2.1 Thermal control	150
10.2.2 Resonator readout and control	152
10.3 Measurements	155
10.3.1 Expectations	155
10.3.2 Results	156
10.3.3 Analysis	157
Chapter XI: Interferometric interrogation of cantilever displacement noise	159
11.1 Overview	159
11.1.1 Design	159
11.1.2 Target sensitivity	161
11.2 Detailed description	162
11.2.1 Optical system	162
11.2.2 Physical system	164
11.2.3 Electrical system	167
11.3 Experimental progress	169
11.3.1 Phase I: Test cavities	169
11.3.2 Phase II: Dummy cantilever cavities	173
11.3.3 Phase III: Real cantilever cavities	181
11.4 Conclusion	184
Bibliography	187
Appendix A: Feedback control	203
A.1 Overview	203

A.2 Stability	206
Appendix B: Pound-Drever-Hall locking	207
Appendix C: Supplemental optical spring radiation pressure noise calculations	210
C.0.1 Ideal situation—no optical loss and leading-order terms . . .	210
C.1 Realistic situation—optical loss and next-order corrections	214
C.2 Feedback and closed-loop response	216
C.3 Proposed experimental setup	219

LIST OF ILLUSTRATIONS

<i>Number</i>		<i>Page</i>
1.1	The two orthogonal gravitational wave polarization modes	5
2.1	Illustration of a gravitational wave's effect on the orthogonal arms of a Michelson interferometer.	13
2.2	A Michelson interferometer with Fabry-Pérot arm cavities.	15
2.3	Diagram of a Fabry-Pérot cavity.	16
2.4	A Michelson interferometer with Fabry-Pérot arm cavities and power recycling.	19
2.5	Transfer function from input laser power noise to apparent GW signal in a PRFPMI	25
2.6	Effective model for the common-mode coupled cavity.	25
2.7	The full core aLIGO optical configuration	27
2.8	More accurate diagram of the aLIGO optical layout.	30
2.9	Simple block diagram of a typical digital control loop in aLIGO. . . .	31
2.10	Advanced LIGO PSL.	33
2.11	Advanced LIGO seismic isolation system.	34
2.12	A simple, single-stage pendulum.	35
2.13	Detail of the Advanced LIGO quadruple suspension.	37
2.14	Advanced LIGO noise budget.	38
2.15	Quantum noise spectral density.	44
2.16	Squeezed states.	48
2.17	Coefficient of thermal expansion (CTE) of monocrystalline silicon. . .	50
3.1	Simplified demonstration of the seismometer tilt-coupling problem. . .	55
3.2	The aLIGO tilt-sensing requirement definition.	56
4.1	Performance for a number of different rotation sensors.	61
4.2	Simplified diagram of the laser gyroscope.	62
4.3	Photo of the laser gyroscope experiment.	63
4.4	Noise budget of the laser gyroscope.	65
4.5	Diagram of the out-of-loop RAM monitor setup.	72
4.6	Proposed fiber-distributed gyroscope array scheme.	72
5.1	Diagram of a Gaussian beam.	78
5.2	Higher-order mode spectrum for a cavity with $g_1, g_2 = 0.98$	82

6.1	HOM spectrum for a cavity with the coarse parameters chosen for the OMC.	85
6.2	Power-law models of the eLIGO dark port field spatial structure. . . .	87
6.3	Simulated length scan of the OMC.	88
6.4	Simulated radius of curvature scan of the OMC.	88
6.5	As-built diagram of the L1 OMC.	89
6.6	Diagram of the OMC mirror curvature measurement setup.	91
6.7	Diagram of the OMC test setup.	93
6.8	Photo of the interior of HAM6.	95
6.9	Bird's-eye diagram of the HAM6 layout.	96
6.10	Evaluation of the input-referred noise of the in-vacuum DCPD preamps.	97
6.11	Evaluation of the voltage transfer function of the in-vacuum DCPD preamps.	98
6.12	Simple diagram of the digital model for the OMC length and angular control.	100
6.13	OMC control screen.	101
6.14	Dither locking diagram.	102
6.15	Length scan of the OMC while closely aligned to the interferometer output beam.	104
6.16	Noise in the low-noise power-recycled Michelson interferometer (PRMI) during aLIGO commissioning.	106
7.1	Simplified experimental layout.	112
8.1	A simple mass-spring system with damping.	121
8.2	Illustration of the thermoelastic loss mechanism.	124
8.3	Possible LIGO Voyager suspension ribbon cross sections.	127
8.4	Projection of LIGO Voyager suspension thermal noise as a function of ribbon aspect ratio.	129
8.5	Normalized strain energy density along the bending direction of a ribbon.	131
8.6	Total loss factor as a function of surface layer loss and thickness, for the two extremal aspect ratios. In each case, the acceptable total loss contour is highlighted in green.	132
8.7	Design of the silicon cantilever.	133
8.8	The first few cantilever mode shapes.	134
8.9	Cantilever mechanical transfer function.	136
8.10	Cantilever thermal displacement noise.	137

9.1	Cantilever fabrication process.	139
9.2	Wafer before and after PECVD nitride deposition.	141
9.3	Demonstration of the operation of the contact aligner for photolithography.	142
9.4	Mask for the photolithography process.	143
9.5	Demonstration of the appearance and transfer of the photolithography pattern during development and dry etching.	143
9.6	The KOH through-etch.	145
9.7	The second round of etching.	146
10.1	Stainless steel clamp used to hold the cantilever.	149
10.2	Clamp alignment jig.	150
10.3	Photo of the ringdown measurement experimental volume.	151
10.4	Control screen of the mode ringer module.	153
10.5	Calculated thermoelastic loss vs. temperature for the cantilevers.	156
10.6	Typical ringdown measurement results.	157
11.1	Optical layout for the interferometric experiment.	160
11.2	Initial noise budget estimate for the interferometric experiment.	161
11.3	Drawing of the cryostat used to hold the ribbon cantilever cavities.	165
11.4	Diagram of the suspension from which the optical breadboard is hung within the cryostat.	166
11.5	Modal analysis of the suspension.	167
11.6	Two electronic chains used in the differential cavity beat note readout.	168
11.7	Test cavities.	170
11.8	Evolution of the beat spectrum over part of the test cavity phase.	171
11.9	Detailed FEA model of the test cavity payload.	172
11.10	Test cavity mode identification.	173
11.11	Inter-experiment beat measurement.	174
11.12	Optical contacting of the mirrors to the cantilevers.	175
11.13	Demonstration of the adapter piece between the cantilever clamps and the suspended breadboard.	176
11.14	Demonstration of the radiation pressure cantilever actuation.	178
11.15	In-situ cantilever ringdown using radiation pressure excitation.	179
11.16	Noise in the dummy cantilever cavities compared with expected thermal noise, along with a projection of thermal noise in the real cantilever cavities.	180

11.17	Comparison of the mechanical responses of the dummy and real cantilevers.	181
11.18	Monolithic suspended payload design.	182
11.19	Comparison of the measured noise in the real and dummy cantilever cases, with thermal noise estimates.	183
11.20	Attempt to explain the elevated noise level using anomalously low- Q higher-order bending modes.	185
A.1	Diagram of a generic feedback loop.	204
B.1	A Pound-Drever-Hall locking setup.	207
C.1	A graphical representation of the feedback model described in the text.	217
C.2	The proposed experimental setup.	220

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1	37
7.1	117
8.1	129
8.2	130

Part I

Background and Theory

Chapter 1

GRAVITATIONAL WAVES

1.1 Basic gravitational wave theory

Here, I will give a *very* brief summary of general relativity and gravitational waves.

1.1.1 General relativity

Gravity, a fundamental force of nature, is currently best understood using Einstein’s theory of general relativity (GR) [1]. In contrast to the earlier Newtonian theory, which treated the force of gravity as an instantaneous attractive interaction between any two massive objects, general relativity describes the phenomenon as arising from the interaction of matter and energy with the continuum of spacetime. Specifically, the presence of matter or energy induces a local curvature in the spacetime manifold, which in turn affects the trajectories of objects through it.¹

In this way, the effect of gravity can be understood as a purely geometrical one: objects under the sole influence of gravity must travel along an extremal 4-dimensional (4-D) path—known as a **geodesic**—that is determined everywhere by the local spacetime curvature.

General relativity is therefore a field theory of gravity, governed by the **Einstein field equation**²:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (1.1)$$

On the right-hand side of this equation, G is Newton’s gravitational constant, c is the speed of light in vacuum, and $T_{\mu\nu}$, known as the **stress-energy tensor**, represents matter and energy present in a region of spacetime. The quantity on the left-hand side, $G_{\mu\nu}$, is called the **Einstein tensor**, and it encapsulates the curvature of spacetime. More concretely, the Einstein tensor is constructed as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R. \quad (1.2)$$

Here, the quantities $R_{\mu\nu}$ and R , known as the **Ricci tensor** and the **Ricci scalar**, respectively, are both derived by contraction of a higher-rank tensor, $R^\rho{}_{\sigma\mu\nu}$, termed

¹In the words of Misner, et al. [2], “spacetime tells matter how to move; matter tells spacetime how to curve.”

²This equation and ones that follow in this chapter contain tensors. The reader should be familiar with tensor notation and the Einstein summation convention.

the **Riemann curvature tensor**:

$$R_{\mu\nu} = R^{\lambda}{}_{\mu\lambda\nu} \quad (1.3)$$

$$R = R^{\lambda}{}_{\lambda} . \quad (1.4)$$

The quantity $g_{\mu\nu}$, called the **metric tensor** (or simply “the metric”), directly encodes the geometry of spacetime and appears in GR whenever this geometry is invoked. For example, the metric is used when calculating tensor contractions (e.g., $R = R^{\lambda}{}_{\lambda} = g^{\mu\nu} R_{\mu\nu}$)³. It is also used in the definition of the Riemann curvature tensor,

$$R^{\rho}{}_{\sigma\mu\nu} = \frac{\partial}{\partial x^{\mu}} \Gamma^{\rho}{}_{\nu\sigma} - \frac{\partial}{\partial x^{\nu}} \Gamma^{\rho}{}_{\mu\sigma} + \Gamma^{\rho}{}_{\mu\lambda} \Gamma^{\lambda}{}_{\nu\sigma} - \Gamma^{\rho}{}_{\nu\lambda} \Gamma^{\lambda}{}_{\mu\sigma} , \quad (1.5)$$

via the **Christoffel symbols**:

$$\Gamma^{\lambda}{}_{\mu\nu} = \frac{1}{2} g^{\lambda\rho} \left(\frac{\partial g_{\rho\mu}}{\partial x^{\nu}} + \frac{\partial g_{\rho\nu}}{\partial x^{\mu}} - \frac{\partial g_{\mu\nu}}{\partial x^{\rho}} \right) . \quad (1.6)$$

Most pointedly, the metric determines the **spacetime interval** (or **line element**),

$$ds^2 = g_{\mu\nu} x^{\mu} x^{\nu} , \quad (1.7)$$

which generalizes the concept of distance in three dimensions to 4-D. By extremizing an action defined as the length of a curve in spacetime (i.e., $S = \int ds$), one arrives at the **geodesic equation**:

$$\frac{d^2 x^{\mu}}{d\tau^2} + \Gamma^{\mu}{}_{\rho\sigma} \frac{dx^{\rho}}{d\tau} \frac{dx^{\sigma}}{d\tau} = 0 , \quad (1.8)$$

where the proper time τ is defined with respect to the spacetime interval as

$$d\tau^2 \equiv - ds^2 . \quad (1.9)$$

1.1.2 Gravitational waves

In regions far from any matter or energy, we have $T_{\mu\nu} \approx 0$. Here, defining the **Minkowski metric** that describes flat spacetime,

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} , \quad (1.10)$$

³The quantity $g^{\mu\nu}$ found here is known as the **inverse metric tensor** and is defined through the equation $g^{\mu\lambda} g_{\lambda\nu} = \delta^{\mu}_{\nu}$, where δ^{μ}_{ν} is the Kronecker delta.

we can write

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.11)$$

and treat spacetime as flat, save for a small perturbation $h_{\mu\nu}$, called the spacetime **strain**. In this regime, known as the **weak-field limit**, one can search for solutions to a wave equation in $h_{\mu\nu}$ of the form [2]

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0. \quad (1.12)$$

After appropriate gauge fixing⁴, a solution to this equation can be found of the form

$$h_{\mu\nu}(t, z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos \left[\omega_{\text{GW}} \left(t - \frac{z}{c} \right) \right], \quad (1.13)$$

where the orthogonal polarization amplitudes h_+ and h_\times are known as “ h -plus” and “ h -cross”, respectively. This equation describes a gravitational wave of angular frequency ω_{GW} traveling in the $+\hat{z}$ direction.

Using (1.7), (1.11) and the form of $h_{\mu\nu}$ we can see that, at $t, z = 0$ in the presence of this perturbation, the line element becomes

$$ds^2 = -c^2 dt^2 + (1 + h_+) dx^2 + (1 - h_+) dy^2 + 2h_\times dx dy + dz^2. \quad (1.14)$$

Apparently, spacetime has been *stretched* in the \hat{x} direction and *compressed* in the \hat{y} direction by the factor $\sqrt{1 + h_+} \approx (1 + \frac{h_+}{2})$. Due to the oscillatory nature of (1.13), it is clear that the opposite will be true at a GW phase shift of π from this point. The orthogonal polarization amplitude h_\times has the identical effect, only along the lines $y = x$ and $y = -x$, as can be confirmed by a coordinate rotation of 45° about the z axis. These polarizations can be combined with an arbitrary phase shift to produce more complicated waveforms (e.g., “circularly polarized” GWs, etc.). A diagram of these modes can be found in Fig. 1.1.

1.1.2.1 GW sources and amplitudes

Gravitational waves are generated by any object with a fluctuating **quadrupole moment**. Given a source object A with quadrupole moment $I_{\mu\nu}^{(A)}$, the spacetime

⁴The choice of gauge used here, called the **transverse traceless (TT) gauge**, assigns fixed coordinates to objects traveling along spacetime geodesics (i.e., objects that are “freely falling”).




Figure 1.1: The two orthogonal gravitational wave polarization modes, h_+ and h_\times , as a function of phase. It should be easy visualize the “circular polarization” that results from a combination of these two with a phase shift of $\pi/2$.

strain caused by A evaluated a distance r away is given by [3]

$$h_{\mu\nu}^{(A)}(r, t) = \frac{2G}{rc^4} \ddot{I}_{\mu\nu}^{(A)}(t - \frac{r}{c}), \quad (1.15)$$

where the dots denote time derivatives.

How large are gravitational waves? By definition, the strain amplitudes are far smaller than unity ($h \equiv |h_{\mu\nu}| \ll 1$), but by how much? As I’ll show in more detail in Sec. 1.2.1, we can use (1.15) to estimate the amplitude of a GW reaching the earth from some hypothetical source. For instance, a binary star system located 30 Mpc away, whose constituents are two solar-mass neutron stars orbiting one another nearly close enough to touch (on the order of 10 km), would produce a peak strain on earth of $h \approx 10^{-21}$! Due to their extreme faintness, then, GWs from all but the most violent events in the universe bear no hope of being detected here on earth.

Gravitational wave sources can be roughly divided into four categories:

- **Continuous:** Steady sinusoidal signals produced by individual rotating massive objects (e.g., a neutron star with a mountain on its surface).
- **Compact Binary Inspiral:** Quasi-sinusoidal signals produced by binary systems of dense objects like neutron stars and black holes.

- **Burst:** Brief, wideband signals theoretically emitted by some energetic objects like supernovae.
- **Stochastic:** Cosmic background produced by phase changes in the early universe.

To date, all GW detections have been in the compact binary inspiral category.

1.2 History of detection science

Even as he first predicted them, Einstein was already thinking about the possibility of detecting gravitational waves [4]. Given how hard they would be to measure, however, he himself had little confidence we would ever do so. It took a number of years for the scientific community to reach a consensus that the GWs were real and produced an observable effect [5].

1.2.1 First estimates: Hertz-type experiment

In 1879, the Berlin Academy of Sciences offered a prize to anyone who could prove experimentally the relationship between electromagnetic forces and the dielectric polarization of insulators. Heinrich Hertz pursued this prize at the insistence of his mentor, Hermann von Helmholtz [6]. In a series of experiments throughout the 1880s, Hertz built scientific apparatuses to generate electromagnetic waves and then detect them, all within the same laboratory. In doing so, he experimentally confirmed many of James Clerk Maxwell's theoretical predictions on electromagnetism.

Given the clear parallels, one might ask if it is possible to perform the analogous experiment with GWs. For example, if we took a barbell-like object with heavy masses on either end and then rotated it rapidly, what would be the magnitude of the GWs generated? Following the analysis in [3], we can use (1.15) to answer that question.

The quadrupole moment is defined as

$$I_{\mu\nu} \equiv \int (x_\mu x_\nu - \frac{1}{3} \delta_{\mu\nu} r^2) \rho(\mathbf{r}) dV . \quad (1.16)$$

If we assume our radiator is a barbell of point mass M on each side, separated by a distance d and rotating about its center of mass at a frequency f_{rot} in the $x - y$

plane, this gives⁵:

$$I_{xx} = \frac{Md^2}{2} \left[\cos^2(2\pi f_{\text{rot}}t) - \frac{1}{3} \right] \quad (1.17)$$

$$I_{yy} = \frac{Md^2}{2} \left[\sin^2(2\pi f_{\text{rot}}t) - \frac{1}{3} \right] \quad (1.18)$$

$$I_{xy} = I_{yx} = \frac{Md^2}{2} \cos(2\pi f_{\text{rot}}t) \sin(2\pi f_{\text{rot}}t). \quad (1.19)$$

Therefore, using (1.15):

$$h = |h_{\mu\nu}| \sim \frac{8\pi^2 G}{c^4} \frac{Md^2 f_{\text{rot}}^2}{z}, \quad (1.20)$$

where z is the distance from the $x - y$ plane at which we make our detection. Let's say we could manage to make a device with $M = 1000$ kg, $d = 1$ m, and $f_{\text{rot}} = 1$ kHz. That would give

$$h \sim 6.5 \times 10^{-34} \times \left(\frac{1 \text{ m}}{z} \right). \quad (1.21)$$

As if that weren't bad enough, in order to distinguish a GW from near-field effects, we would need to be far from the source in the **wave zone** (i.e., $z \gtrsim \lambda_{\text{GW}} = \frac{2\pi c}{\omega_{\text{GW}}}$). Even at the nearer edge with $z = \lambda_{\text{GW}} \approx 300$ km, this gives

$$h_{\text{best}} \sim 2 \times 10^{-39}, \quad (1.22)$$

which is very, very small.

1.2.2 Early detection attempts: Weber's resonant bars

As it was clear that the only gravitational waves we would stand a chance of detecting would be astrophysical in origin, experimenters started devising apparatuses to make such a measurement. In the 1960s, Joseph Weber began searching for GWs with large aluminum cylinders [7], which he realized could be used as mechanical resonant amplifiers.

To understand how bar detectors work, it helps to introduce the concept of treating a GW as a tidal force acting on the detector system. Imagine two free point masses of mass M separated in space by a distance L . If a GW were to act on these masses with strain amplitude $h(t)$ along the line between them, then—relative to the midpoint between them—an experimenter would see that each object is displaced by a distance

$$x_{\text{GW}}(t) = \frac{L}{2} \frac{h(t)}{2}. \quad (1.23)$$

⁵The z components of $I_{\mu\nu}$ either vanish or are constant, and therefore do not contribute to GW generation.

Given this displacement response, one can write down an apparent gravitational wave force of magnitude

$$|F_{\text{GW}}| \equiv M\ddot{x}_{\text{GW}} = \frac{ML}{4}\ddot{h} \quad (1.24)$$

acting on each object.

Even though the calculation above is done for two free masses, it would still be valid to use F_{GW} to describe the passing GW even if we were to connect the two masses using a spring. In that case, we would expect the system's response to GWs, x/h , to be augmented by the resonant mechanical response of the system to external forces. In fact, this force approach can be used to calculate the effect of a passing GW on an arbitrarily constructed distribution of mass⁶. Taking advantage of this effect, one can use high-quality mechanical resonators as bells, ringing at their resonant frequencies when acted upon even by the paltry force of a gravitational wave.

Weber experimented with bar detectors for many years, and though he claimed to have made GW observations, they were never corroborated by other researchers [8]. Nevertheless, he is credited with sparking interest in the field of gravitational wave detection, and bar detectors were still in use many decades later [9].

1.2.3 Astronomical observations

Several attempts have been made to observe the effects of gravitational waves by studying the motion of objects and other phenomena in the sky.

1.2.3.1 The Hulse-Taylor binary

As a form of radiation, gravitational waves carry energy. Therefore, as an astrophysical system emits GWs, it must lose the energy carried away. In the case of a binary star system, this loss of energy translates into an **orbital decay**, or the gradual inspiral of each constituent towards the system's barycenter.

In the 1979, Hulse and Taylor described [10] a binary star system (PSR B1913+16) with a pulsar as one constituent star. By measuring the doppler shift of the pulsar emissions due to the binary orbit, they were able to accurately track the orbital period of the system. In doing so, they observed that it was spinning down at a rate consistent with that predicted by GR via gravitational radiation.

This discovery earned Hulse and Taylor the 1993 Nobel Prize in Physics and catalyzed a strong scientific pursuit of a direct detection.

⁶This tidal force formalism is valid in the long-wave limit (i.e., $L \ll \lambda_{\text{GW}}$).

1.2.3.2 CMB polarization

According to the theory of inflation, the universe experienced a period of rapid expansion shortly after its inception at the big bang. As it expanded, it carried with it quantum fluctuations that were present at its previous tiny scales. Once the universe had expanded and cooled enough to transmit light, photons encoded with these fluctuations were emitted in all directions. These photons, constituting the **cosmic microwave background (CMB)**, are perceptible from all angles on the sky, and provide deep insight into the history of the universe.

In studying the anisotropy of the CMB, experts make a distinction between curl-free “E-modes” and divergence-free “B-modes.” While many fields can be responsible for the former, the latter can only be accounted for by interaction with primordial gravitational waves. Therefore, some researchers endeavor to detect the effect of gravitational waves by searching for B-mode anisotropy in the CMB [11].

1.2.3.3 Timing arrays

In the 1970s, some realized [12, 13] that low-frequency gravitational waves could be detected by precise timing of known pulsars in the sky. By observing the minute variations in the apparent arrival of pulsar emissions from multiple stars, their relative velocities with respect to earth could be tracked. In doing so, one could observe the effect of a gravitational wave passing through the intervening space.

Around the same time, it was suggested that this method could be applied to man-made spacecraft, as well [14].

1.2.4 Laser interferometers

The following is a ridiculously brief history of interferometric gravitational wave detectors. For an excellent review of the field, see [15].

1.2.4.1 First generation

In 1972, Rainer Weiss proposed using a Michelson interferometer to search for gravitational waves [16]. Roughly 20 years later, following a major multi-national research, development and prototyping effort, the Laser Interferometer Gravitational-wave Observatory (LIGO) was born [17].

By the mid-2000s, the twin LIGO observatories in Hanford, WA and Livingston, LA were operating at the design sensitivity of their first incarnation, termed **Initial**

LIGO (iLIGO) [18]. These first-generation detectors reached a peak strain sensitivity of about $2 \times 10^{-23} \text{ 1}/\sqrt{\text{Hz}}$ near 100 Hz, and—along with a network of partnering detectors worldwide [19–21]—proved the principle of interferometric gravitational wave detection.

While a second-generation detector had been planned for LIGO from the start, it was decided during iLIGO to install an incremental upgrade known as **Enhanced LIGO (eLIGO)** [22]. This upgrade, which included a more powerful laser and a novel GW signal readout scheme, resulted in a modest strain sensitivity improvement of roughly a factor of two over most of the operational frequency band.

1.2.4.2 Second generation

In late 2010, the LIGO detectors went down and began installation of their second incarnation, **Advanced LIGO (aLIGO)** [23]. This *major* upgrade included a completely new, much-higher-power laser source, a new multi-stage seismic isolation system, an exquisitely more complicated multi-stage test mass suspension including much larger mirrors, a modified interferometer topology, new sensors and actuators, and a host of additional auxiliary control systems. As a result of these improvements, aLIGO will eventually have a strain sensitivity nearly ten times better than that of iLIGO.

On September 14, 2015, while operating in an “engineering run” prior to its first true observational run, aLIGO made the world’s first direct detection of a gravitational wave [24]. This detection, christened **GW150914**, was a loud ($\text{SNR} \approx 24$) event caused by the merger of two relatively massive black holes. For this achievement, the 2017 Nobel Prize in Physics was awarded to LIGO pioneers Kip Thorne, Rainer Weiss, and Barry Barish.

Since that initial detection and up to the date of this publication, LIGO and its partners have detected ten subsequent GW events [25], including one from a binary neutron star coalescence [26] that was paired with electromagnetic observations in a major breakthrough for **multi-messenger astronomy**.

Advanced LIGO continues to improve its sensitivity and is scheduled to meet design specifications within the next few years [27].

Chapter 2

STATE OF THE ART: ADVANCED LIGO

This chapter contains a detailed look at the Advanced LIGO detector, divided into three parts. In Sec. 2.1, I give a thorough description of the interferometer, including the optical configuration and the control strategy. Sec. 2.2 contains an overview of the various noise sources that conspire to limit the sensitivity of aLIGO. Finally, in Sec. 2.3, I give a brief synopsis of the path forward for improving aLIGO and terrestrial GW detectors in general, which will motivate the main work of this thesis in Part V.

2.1 Description of the instrument

Advanced LIGO is a complicated machine. This section gives a detailed description of the aLIGO optical configuration, as well as the control strategy needed to make it work.

2.1.1 Optical configuration

2.1.1.1 The Michelson interferometer

At the core of Advanced LIGO is a Michelson interferometer. To understand why, consider the response of such an apparatus to a gravitational wave (illustrated in Fig. 2.1). Defining the input field

$$E_{\text{in}} = E_0 e^{i\omega_0 t}, \quad (2.1)$$

where ω_0 is the laser angular frequency, we can calculate the **input-output relation** of this interferometer. First, the fields leaving the beamsplitter and entering the arms are

$$\begin{aligned} E_{\text{xi}} &= \frac{i}{\sqrt{2}} E_{\text{in}} \\ E_{\text{yi}} &= \frac{1}{\sqrt{2}} E_{\text{in}}, \end{aligned} \quad (2.2)$$

where we will use the so-called ‘‘Siegman convention’’ [28], wherein amplitude transmission coefficients carry the complex factor i , while reflection coefficients do

not. Propagating these fields down the arms, reflecting off the end mirrors, and then propagating them back to the beamsplitter gives

$$\begin{aligned} E_{x0} &= e^{-ikL_x} \cdot (1) \cdot e^{-ikL_x} E_{xi} = \frac{ie^{-2ikL_x}}{\sqrt{2}} E_{in} \\ E_{y0} &= e^{-ikL_y} \cdot (1) \cdot e^{-ikL_y} E_{yi} = \frac{e^{-2ikL_y}}{\sqrt{2}} E_{in}, \end{aligned} \quad (2.3)$$

where $k = \frac{2\pi}{\lambda} = \frac{\omega_0}{c}$ is the laser wave number and L_x and L_y are the lengths of the x and y arms, respectively. Finally, the two output fields are:

$$\begin{aligned} E_{out} &= \frac{E_{x0}}{\sqrt{2}} + \frac{iE_{y0}}{\sqrt{2}} = \frac{iE_{in}}{2} (e^{-2ikL_x} + e^{-2ikL_y}) \\ &= iE_0 e^{i[\omega_0 t - k(L_x + L_y)]} \cos[k(L_x - L_y)] \end{aligned} \quad (2.4)$$

$$\begin{aligned} E_{refl} &= \frac{iE_{x0}}{\sqrt{2}} + \frac{E_{y0}}{\sqrt{2}} \\ &= -iE_0 e^{i[\omega_0 t - k(L_x + L_y)]} \sin[k(L_x - L_y)]. \end{aligned} \quad (2.5)$$

Now, a photodetector doesn't directly measure the *amplitude* of the electromagnetic field, but rather its *power*, i.e.:

$$\begin{aligned} P_{out} &= E_{out}^* E_{out} = P_0 \cos^2[k(L_x - L_y)] \\ P_{refl} &= E_{refl}^* E_{refl} = P_0 \sin^2[k(L_x - L_y)], \end{aligned} \quad (2.6)$$

where $P_0 \equiv E_0^2$. Thus, the transmission or reflection of input light by the interferometer is dependent on the differential length of its arms.

Note that the quantity kL has units of radians, and it quantifies a phase traversed by the optical field over a length L . Examining the equations above, we see that what really matters for the interference condition is the phase. Changing the physical length of the arms is one way to affect this phase, but it is not the only way; for instance, one could change the optical path length by adding a refractive medium into the arm, or a gravitational wave could cause the space itself to stretch or shrink. In general terms of this phase, $\Delta\phi$, we can rewrite (2.6):

$$\begin{aligned} P_{out} &= P_0 \cos^2 \frac{\Delta\phi}{2} = \frac{P_0}{2} (1 + \cos \Delta\phi) \\ P_{refl} &= P_0 \cos^2 \frac{\Delta\phi}{2} = \frac{P_0}{2} (1 - \cos \Delta\phi). \end{aligned} \quad (2.7)$$




Figure 2.1: Illustration of a gravitational wave's effect on the orthogonal arms of a Michelson interferometer.

2.1.1.2 Michelson response to a gravitational wave

As discussed briefly in Sec. 1.2.2, in the limit of long, low-frequency GWs, it is valid to treat the wave as a force F_{GW} acting on the free end masses, resulting in displacements

$$L_x(t) = L_{0x} + \Delta L_{\text{GW}}(t) \approx L_{0x} \left(1 + \frac{h(t)}{2} \right) \quad (2.8)$$

$$L_y(t) = L_{0y} - \Delta L_{\text{GW}}(t) \approx L_{0y} \left(1 - \frac{h(t)}{2} \right).$$

In the case where the macroscopic lengths of the arms are equal (i.e., $L_{0x} = L_{0y} \equiv L$), the differential signal produced by the GW is

$$L_x(t) - L_y(t) = Lh(t). \quad (2.9)$$

As the GW's wavelength shortens and becomes comparable to the length of the detector, one must use a more accurate approach. Recall that in the TT gauge,

freely falling objects (like our mirrors, at least in the $x - y$ plane) are by definition at fixed spatial coordinates. However, as the spacetime line element is modulated by a passing GW, a light beam traveling between two such objects must traverse a varying path length as the metric perturbation passes.

Following [3], we can compute the accurate response of a +-polarized GW by integrating the square root of (1.14), noting that the spacetime interval between any two events linked by a beam of light is zero. For the outbound leg down the X arm:

$$\begin{aligned} \int_0^{\tau_{x,\text{out}}} dt &= \frac{1}{c} \int_0^L \sqrt{1 + h(t)} dx \\ &= \frac{1}{c} \int_0^L \sqrt{1 + h_+ \cos(\omega_{\text{GW}} t)} dx \\ &\approx \frac{1}{c} \int_0^L \left(1 + \frac{h_+ \cos(\omega_{\text{GW}} t)}{2} \right) dx . \end{aligned} \quad (2.10)$$

A similar integral can be constructed for the return leg to determine $\tau_{x,\text{ret}}$. The time dependence of the integrand on the dx side should be rewritten in terms of the position of a single phase front (i.e., $t = \frac{x}{c}$ for the outbound leg and $t = \frac{2L-x}{c}$ for the return leg).

After finding the total, roundtrip travel time $\tau_{x,\text{RT}} = \tau_{x,\text{out}} + \tau_{x,\text{ret}}$ and the corresponding time for the Y arm, one arrives at the differential phase

$$\begin{aligned} \Delta\phi(t) &= \frac{2\pi c}{\lambda} (\tau_{x,\text{RT}} - \tau_{y,\text{RT}}) \\ &= \frac{2\pi L}{\lambda} h_+ \text{sinc}\left(\frac{\omega_{\text{GW}} L}{c}\right) \cos(\omega_{\text{GW}} t) . \end{aligned} \quad (2.11)$$

Plugging this into (2.7) gives the accurate power response of the interferometer to a sinusoidal GW $h(t)$. This treatment is important when considering higher-frequency gravitational waves; for the remainder of this thesis, we will use the approximation in (2.9), and we will regularly consider the displacement response and noise of the instrument rather than directly referring to the strain.

2.1.1.3 Arm cavities

As we have seen, the core function of the Michelson interferometer is to convert phase differences in the light returning from its orthogonal arms into power modulation at the reflection and output ports, which will hereafter sometimes be referred to as the **symmetric** and **anti-symmetric (AS)** ports. Given this fact, we can seek out

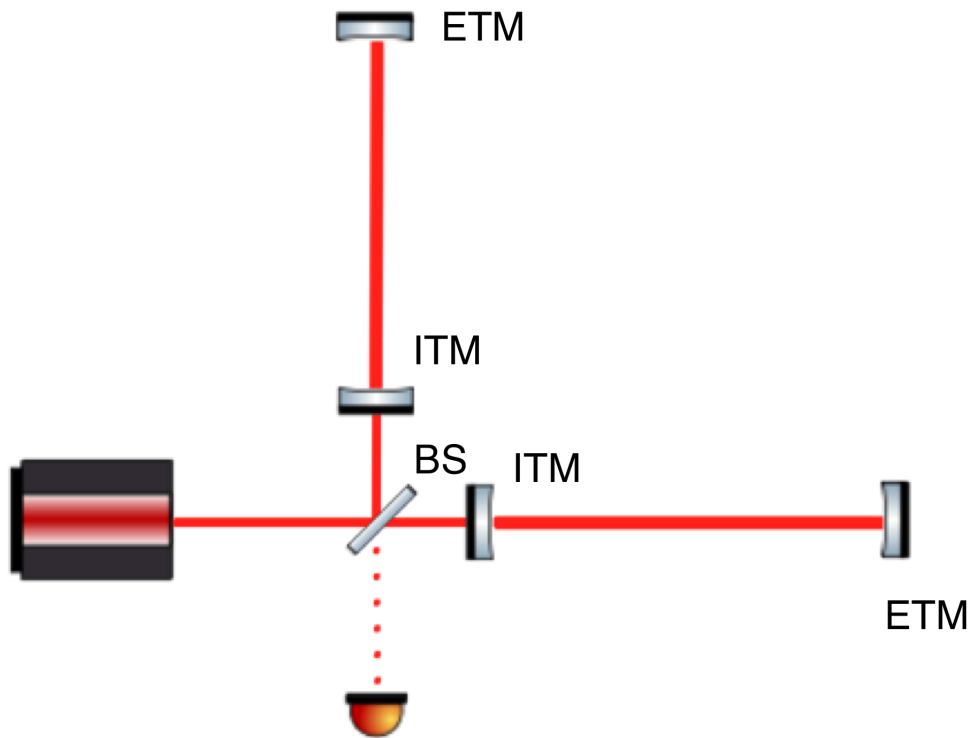


Figure 2.2: A Michelson interferometer with Fabry-Pérot arm cavities. Labeled are the **input test masses (ITMs)**, **end test masses (ETMs)**, and the **beamsplitter (BS)**.

methods of increasing the displacement-to-phase response of the arms. One way to do this is to install Fabry-Pérot cavities into the arms, as shown in Fig. 2.2.

To understand how this works, consider the cavity drawn schematically in Fig. 2.3, and let us compute the fields E_{refl} and E_{trans} as a function of the input field $E_{\text{in}} = E_0 e^{i\omega_0 t}$ (i.e., the input-output relation). The field just inside the input mirror is

$$E_1 = it_i E_{\text{in}} + r_i E_3. \quad (2.12)$$

But we also know that

$$E_3 = r_e e^{2ikL} E_1, \quad (2.13)$$

so

$$E_1 = it_i E_{\text{in}} + r_i r_e e^{2ikL} E_1. \quad (2.14)$$

Therefore, recursively,

$$E_1 = \left(\frac{it_i}{1 - r_i r_e e^{2ikL}} \right) E_{\text{in}}. \quad (2.15)$$

Using this, we can find the intracavity **circulating power**:

$$P_{\text{circ}} = |E_1|^2 = \left(\frac{t_i^2}{1 + r_i^2 r_e^2 - 2r_i r_e \cos(2kL)} \right) |E_{\text{in}}|^2. \quad (2.16)$$

In particular, in the **resonant** condition where $2kL = n \times 2\pi$,

$$G \equiv \frac{P_{\text{circ}}}{P_{\text{in}}}\bigg|_{\text{res}} = \left(\frac{t_i}{1 - r_i r_e} \right)^2, \quad (2.17)$$

where $P_{\text{in}} = |E_{\text{in}}|^2$ and G is known as the cavity **power gain**.




Figure 2.3: Diagram of a Fabry-Pérot cavity.

We can also now find the transmitted and reflected fields:

$$\begin{aligned} E_{\text{trans}} &= it_e E_2 = it_e e^{ikL} E_1 \\ &= \left(\frac{-t_i t_e e^{ikL}}{1 - r_i r_e e^{2ikL}} \right) E_{\text{in}} \end{aligned} \quad (2.18)$$

$$\begin{aligned} E_{\text{refl}} &= r_i E_{\text{in}} + it_i E_3 = r_i E_{\text{in}} + it_i r_e e^{2ikL} E_1 \\ &= \left[r_i - \left(\frac{t_i^2 r_e e^{2ikL}}{1 - r_i r_e e^{2ikL}} \right) \right] E_{\text{in}} \\ &= \left(\frac{r_i - r_e e^{2ikL}}{1 - r_i r_e e^{2ikL}} \right) E_{\text{in}}. \end{aligned} \quad (2.19)$$

Consider the case $r_i = r_e \equiv \sqrt{R}$, where R is the power reflection coefficient, so that $t_i = t_r \equiv \sqrt{T} = \sqrt{1 - R}$. This case, where the input and end mirrors are identical, is called **critical coupling**. Here, we have

$$\frac{P_{\text{trans}}}{P_{\text{in}}} = \frac{|E_{\text{trans}}|^2}{|E_{\text{in}}|^2} = \frac{T^2}{1 + R^2 - 2R \cos(2kL)} \quad (2.20)$$

$$\frac{P_{\text{refl}}}{P_{\text{in}}} = \frac{|E_{\text{refl}}|^2}{|E_{\text{in}}|^2} = \frac{2R [1 - \cos(2kL)]}{1 + R^2 - 2R \cos(2kL)}. \quad (2.21)$$

Therefore, on resonance,

$$\left. \frac{P_{\text{trans}}}{P_{\text{in}}} \right|_{\text{res}} = \frac{T^2}{(1-R)^2} = 1 \quad (2.22)$$

$$\left. \frac{P_{\text{refl}}}{P_{\text{in}}} \right|_{\text{res}} = 0, \quad (2.23)$$

and the cavity is perfectly transmitting.

Now consider the case when the end mirror is perfectly reflecting (i.e., $r_e = 1$)¹. Assuming there are no optical losses in the system, all the power must be reflected regardless of the resonance condition. Using (2.19), however, we can see that the phase changes as we sweep the cavity from anti-resonant to resonant:

$$E_{\text{refl}} = \begin{cases} E_{\text{in}} & \text{anti-resonant} \\ -E_{\text{in}} & \text{resonant} \end{cases}. \quad (2.24)$$

Differentiating that equation with respect to L and evaluating on resonance, we find:

$$\left. \frac{dE_{\text{refl}}}{dL} \right|_{\text{res}} = 2ik \left(\frac{1+r_i}{1-r_i} \right) E_{\text{in}}. \quad (2.25)$$

As expected, the derivative is imaginary, meaning the response is in phase. Comparing this result with what we would get for a simple reflecting mirror with no input mirror (i.e., $r_i = 0$), we see that the phase response is increased by the factor in the parentheses². That is,

$$\left(\frac{d\phi_{\text{refl}}}{dL} \right)_{\text{cavity}} = \left(\frac{1+r_i}{1-r_i} \right) \left(\frac{d\phi_{\text{refl}}}{dL} \right)_{\text{mirror}}. \quad (2.26)$$

Thus, via (2.7), the response of the Michelson to gravitational waves is increased by this factor when Fabry-Pérot cavities are introduced into the arms.

2.1.1.4 Power recycling

As we'll see in more detail in below, the SNR of the interferometer with respect to certain sources of noise originating from the laser is maximized very close to the **dark fringe**, where the differential arm phase is $\Delta\phi \approx \pi$. Here, according to (2.7), $P_{\text{out}} \approx 0$ and nearly all the light entering the interferometer is eventually reflected

¹A cavity for which $r_e > r_i$ is called **overcoupled**; if $r_e < r_i$, it is called **undercoupled**, though this case is not often used.

²Note that this factor is equal to the power gain G .

back towards the laser³. Given this obvious waste, one can ask how we might reuse this power leaving the interferometer.

In reality, not all the light entering the interferometer is sent back to the laser, even on the dark fringe; since all optical systems exhibit some loss (e.g., from scattering and absorption, etc.), some light is lost in the arms. This loss is enhanced by the arm cavities: if the power loss over a single roundtrip of the arm cavity is ϵ_{arm} , the light power being constantly lost from the overall system is

$$P_{\text{lost}} = 2\epsilon_{\text{arm}}P_{\text{arm}} = \epsilon_{\text{arm}}G_{\text{arm}}P_{\text{in}}. \quad (2.27)$$

With this in mind, we can imagine that the entire Fabry-Pérot Michelson interferometer (FPMI) is just a mirror with amplitude transmission coefficient

$$t_{\text{FPMI}} = \sqrt{\epsilon_{\text{arm}}G_{\text{arm}}}. \quad (2.28)$$

Therefore—by analogy with (2.22) and (2.23)—if we place a new mirror between the laser and the BS (as shown in Fig. 2.4) and set its transmission coefficient to $t_{\text{PRM}} = t_{\text{FPMI}}$, we critically couple the light to these integrated losses, maximizing both the storage time and circulating power within the interferometer. In this configuration, no wasted light emerges back at the bright port.

To calculate the power recycling enhancement, we need the arm cavity power gain:

$$G_{\text{arm}} \approx \left(\frac{t_{\text{ITM}}}{1 - r_{\text{ITM}}} \right)^2 \approx \frac{4}{t_{\text{ITM}}^2}, \quad (2.29)$$

where we have used the approximation $\epsilon_{\text{arm}} \ll t_{\text{ITM}}^2 \ll 1$. Therefore, we choose $t_{\text{PRM}} = \frac{2\sqrt{\epsilon_{\text{arm}}}}{t_{\text{ITM}}}$, giving⁴

$$G_{\text{PRC}} = \frac{1}{T_{\text{PRM}}} = \frac{T_{\text{ITM}}}{4\epsilon_{\text{arm}}}, \quad (2.30)$$

where “PRC” refers to the **power recycling cavity** created between the PRM and the average of the two ITMs.

Since the optical power incident on the BS is now $G_{\text{PRC}}P_{\text{in}}$, we have increased the sensitivity of the instrument to GWs by this PRC gain via (2.7). Note also that the circulating power in the arms is now

$$P_{\text{arm}}^{\text{PRFPMI}} = \frac{G_{\text{PRC}}G_{\text{arm}}P_{\text{in}}}{2}. \quad (2.31)$$

³Due to this choice of operating point, the symmetric and anti-symmetric ports are also known as the **bright port** and **dark port**, respectively.

⁴Here and elsewhere we will use the notation $T_X \equiv t_X^2$ and $R_X \equiv r_X^2$, with upper-case letters referring to power transmission/reflection coefficients and lower-case letters referring to their amplitude counterparts.

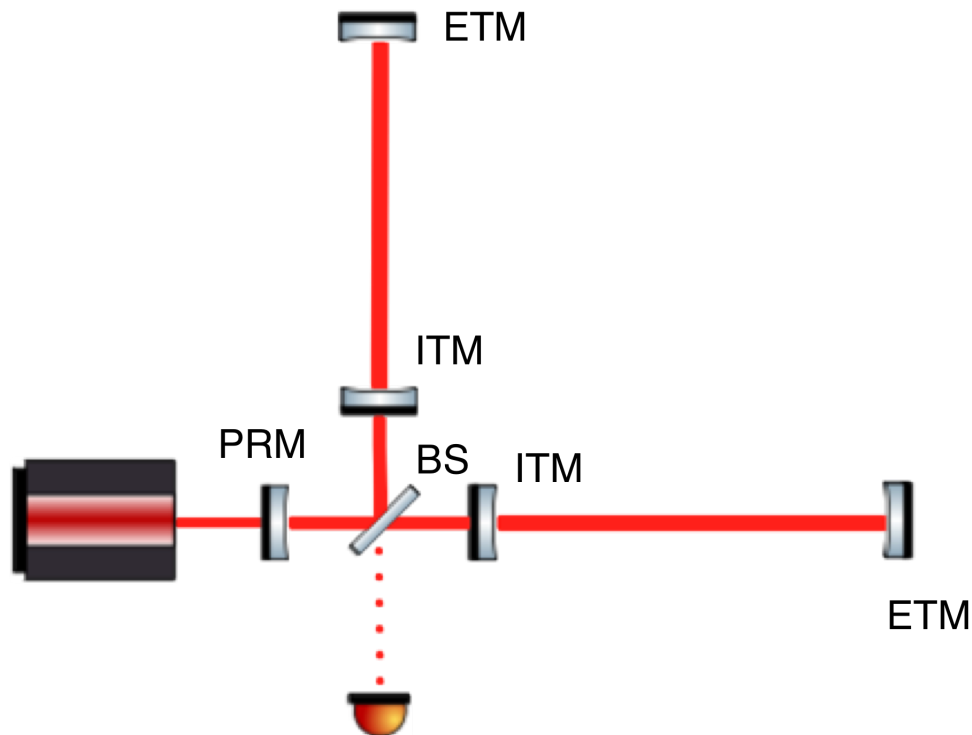


Figure 2.4: A Michelson interferometer with Fabry-Pérot arm cavities and power recycling. The added mirror between the laser and the BS is called the **power recycling mirror (PRM)**.

2.1.1.5 The gain-bandwidth tradeoff

The improvements that come from resonant enhancement of the detector do not come without a cost. Simply put, the increase in gain is the result of photons circulating for a longer period of time within the optical system; this increase in **storage time** results in a corresponding decrease in the detection bandwidth.

To understand this effect, it is useful to introduce the concept of modulation. Consider a laser with input field amplitude E_0 and angular frequency ω_0 passing through an element that had the ability to apply a phase shift $\Delta\phi_m(t) = \Gamma \cos(\omega_m t)$. This results in an output field

$$E_{\text{out}} = E_0 e^{i\omega_0 t + \Gamma \cos(\omega_m t)}. \quad (2.32)$$

Using the Jacobi-Anger expansion,

$$e^{iz \cos \theta} = \sum_{n=-\infty}^{\infty} i^n J_n(z) e^{in\theta}, \quad (2.33)$$

where $J_n(z)$ is the n^{th} Bessel function of the first kind, this gives the first-order expansion:

$$\begin{aligned} E_{\text{out}} &\approx E_0 e^{i\omega_0 t} \left(J_0(\Gamma) + iJ_1(\Gamma) e^{i\omega_m t} + iJ_1(\Gamma) e^{-i\omega_m t} \right) \\ &\approx E_0 e^{i\omega_0 t} \left(1 + \frac{i\Gamma}{2} e^{i\omega_m t} + \frac{i\Gamma}{2} e^{-i\omega_m t} \right), \end{aligned} \quad (2.34)$$

where we have used the approximations $J_0(z) \approx 1$ and $J_1(z) \approx \frac{z}{2}$ for $z \ll 1$. Examining this expression and boosting into the frame of the carrier phasor rotating in the complex plane at angular frequency ω_0 , we see that there are three components: 1) a DC field with amplitude E_0 ; 2) a field at relative angular frequency $+\omega_m$ with amplitude $\frac{E_0\Gamma}{2}$; 3) a field at relative angular frequency $-\omega_m$, also with an amplitude $\frac{E_0\Gamma}{2}$.

The intuitive time-domain picture of this modulated field is a phasor whose amplitude is not changing, but whose phase is fluctuating—i.e., deviating from a linear accumulation of $\omega_0 t$ in time—by the sinusoidal factor $\Delta\phi_m(t)$. However, this expansion allows us to treat the phase modulation **sidebands (SBs)** as independent frequency components of the laser field, and these sidebands can be used to compute the dynamic response of a system.

For example, one way of modulating the phase of a laser is to reflect it off of a moving mirror. Indeed, we have already been considering how a length-change-induced phase shift between two orthogonal arms in a Michelson interferometer can lead to power modulations (i.e., *signal*) at its output port, and how installing resonant cavities in those arms can amplify this displacement-to-phase response. In doing so, we found in (2.26) that this resonant enhancement was proportional to the arm cavity power gain G_{arm} . However, this was the gain evaluated *on resonance*. Since the carrier is resonant, the sidebands are necessarily *not* on resonance, and this will affect the AC response of the cavity to mirror displacement.

Using (2.16), (2.17) and the relation $k = \frac{\omega}{c}$ we can define the frequency-dependent cavity power gain

$$G(\omega) = \left(\frac{t_i^2}{1 + r_i^2 r_e^2 - 2r_i r_e \cos\left(\frac{2\omega L}{c}\right)} \right) = \left(\frac{\left(\frac{t_i}{1-r_i r_e}\right)^2}{1 + \frac{4r_i r_e}{(1-r_i r_e)^2} \sin^2\left(\frac{\omega L}{c}\right)} \right) \equiv \frac{G(0)}{1 + F(\omega)}. \quad (2.35)$$

It is clear from this expression that the low-frequency gain $G(0)$ is linked to the rapidness of the high-frequency rolloff, embodied by $F(\omega)$. We can use $F(\omega)$ to find the so-called **half-width half-maximum (HWHM)** frequency at which $G(\omega)$ is reduced by a factor of two. To do this, we set $F(2\pi f_{\text{HWHM}}) = 1$, and therefore

$$f_{\text{HWHM}} = \frac{c}{2\pi L} \arcsin\left(\frac{1 - r_i r_e}{2\sqrt{r_i r_e}}\right). \quad (2.36)$$

In general, cavities can be built with an arbitrary macroscopic length L . Therefore, we would like to have a length-invariant measure of the sharpness of a cavity's resonance. To get this, we can compare the HWHM frequency to another frequency scale: the **free spectral range (FSR)**. This quantity, equal to

$$\nu_{\text{FSR}} = \frac{c}{2L}, \quad (2.37)$$

is the spacing in optical frequency between adjacent cavity resonances. Using these two frequency scales, we can define the **finesse**:

$$\mathcal{F} \equiv \frac{\nu_{\text{FSR}}}{2f_{\text{HWHM}}} = \frac{\pi}{2 \arcsin\left(\frac{1 - r_i r_e}{2\sqrt{r_i r_e}}\right)} \approx \frac{\pi\sqrt{r_i r_e}}{1 - r_i r_e} \approx \frac{2\pi}{\epsilon_{\text{tot}}}, \quad (2.38)$$

where $\epsilon_{\text{tot}} = T_i + T_e = (1 - r_i^2) + (1 - r_e^2) \ll 1$ is the total roundtrip power loss of the cavity⁵. Finally, in the high-finesse limit, the resonant power gain is related to the finesse by

$$G(0) = \begin{cases} \frac{\mathcal{F}}{\pi} & \text{critically coupled } (r_e = r_i) \\ \frac{2\mathcal{F}}{\pi} & \text{fully overcoupled } (r_e = 1). \end{cases} \quad (2.39)$$

Now, in (2.26), we saw that the DC length-to-phase response of the arm cavity is proportional to $G(0)$. Naively then, we might expect the AC response to go as $G(\omega)$, but we would be wrong. To see why, imagine that the modulated laser from (2.32) was injected into the arm cavity before the modulation was applied. The resonant carrier field amplitude becomes

$$|E_{0,\text{cav}}| = \sqrt{G(0)}E_0 \quad (2.40)$$

so the sidebands are generated with seed amplitude

$$|E_{\pm}^{(0)}| = \frac{\sqrt{G(0)}E_0\Gamma}{2}. \quad (2.41)$$

⁵At this point, we are only considering the ‘‘losses’’ we intentionally introduce in the form of the transmissions of the cavity mirrors; in practice, there are always true losses (e.g., from scatter and absorption) that must be included in ϵ_{tot} .

Since they are born within the cavity, they are resonantly enhanced by the factor $\sqrt{G(\omega_m)}$, giving the output signal SB amplitude⁶

$$|E_{\pm}(\omega_m)| = \frac{\sqrt{G(0)G(\omega_m)}E_0\Gamma}{2} \propto \frac{G(0)}{1 + \frac{\omega_m}{2\pi f_{\text{HWHM}}}} \quad (\text{for } \omega_m \ll 2\pi\nu_{\text{FSR}}) \quad . \quad (2.42)$$

Therefore, the DC response of the system to phase or length fluctuations is proportional to the cavity power gain, however, the AC frequency dependence is that of the amplitude gain. This high-frequency rolloff approximates a single-pole low-pass filter at frequencies well below ν_{FSR} . For this reason, the HWHM frequency is often referred to as the **cavity pole** frequency, f_p .

Most saliently, we can see from (2.38) and (2.39) that the product $G(0)f_p$ is a constant for a given fixed cavity length—this is the essence of the gain-bandwidth tradeoff.

2.1.1.6 Coupled cavities

Treating the PRC as a cavity with the PRM as an input mirror and the FPPI as an effective end mirror is not the only way to view the system. From the perspective of a photon circulating in the arm cavity, for example, things look wholly different. To such a photon, the effect of the PRM depends very much on whether there is a complementary photon in the other arm cavity: since the Michelson is being held on the dark fringe, *common-mode* signals generated in the two arms will interfere constructively when they reach the BS and be sent towards the symmetric port, while *differential-mode* signals will interfere destructively and be sent towards the anti-symmetric port.

In other words, photons corresponding to common signals between the two arms (e.g., signals originating from the laser in the first place, or signals caused by common-mode displacements of the arm cavity lengths) will see the influence of the PRM, while differential-mode signals (e.g., changes in the length of only one of the arm cavities, or—as we'll see later—optical field fluctuations entering the interferometer from the anti-symmetric port), have no way of knowing it's there; on the other side of the ITM, a differential-mode signal simply sees the dark port. Therefore, to understand this **coupled cavity** system, we must consider these two types of signals separately.

⁶By analogy with (2.15), the resonant sideband field enhancement is actually $\frac{\sqrt{G(\omega_m)}}{it_i}$, but this intracavity field must be multiplied by the input mirror field transmission it_i to get the signal sideband field (2.42) leaving the cavity and returning to the BS.

In the case of the common-mode signals, due to the gain-bandwidth relationship discussed in the previous section (i.e., since $f_p \propto \frac{1}{G}$), we can calculate the **coupled cavity pole**:

$$f_{p,cc} = \frac{f_p}{2G_{\text{PRC}}}. \quad (2.43)$$

That is, the increase in DC gain by the factor G_{PRC} comes with a reduction in bandwidth by the same factor⁷. In this case, there isn't much of a tradeoff: since the signal enhancements we have been calculating come from the amplification of the DC power, P_0 , we don't really care that the bandwidth of this coupled cavity has been reduced; as we'll see in a moment, the differential-mode signal SBs are unchanged by this effect.

In fact, the narrowing of this cavity is actually beneficial in many ways. For example, consider the effect of input laser power fluctuations. Rewriting the input power as

$$P_0 \rightarrow P(t) = P_0 + \delta P(t), \quad (2.44)$$

we can examine the effect of $\delta P(t)$ on the measured interferometer signal at the AS port. First, consider the simple Michelson case. Using (2.7) and recalling that we are operating at a small offset $\delta\phi$ from the dark fringe (i.e., $\Delta\phi = \pi + \delta\phi$, with $\delta\phi \ll \pi$), we have

$$P_{\text{out}}^{\text{MI}}(t) \approx (P_0 + \delta P(t)) \frac{\delta\phi^2}{4} \quad (2.45)$$

Therefore, the power fluctuation $\delta P(t)$ is passed directly to the output port by the factor $\delta\phi^2/4$. Comparing this to the effect of a true signal (e.g., by writing $\Delta\phi(t) = \pi + \delta\phi + 2kLh(t)$), we measure

$$P_{\text{out}}^{\text{MI}}(t) \approx P_0 \frac{\delta\phi^2}{4} + \frac{\delta\phi^2}{4} \delta P(t) + 2kLP_0\delta\phi h(t). \quad (2.46)$$

Analyzing the fluctuating terms, we see that the power noise couples to the apparent signal as⁸

$$h_{\text{app.}}^{\text{MI}}(t) = \frac{\delta\phi}{8kL} \frac{\delta P(t)}{P_0}. \quad (2.47)$$

⁷The factor of two comes from the fact that the ratio of G to \mathcal{F} changes by a factor of two when going from a fully overcoupled cavity to a critically coupled cavity.

⁸This result shows two important features of laser intensity noise: 1) the coupling is smaller as we approach the dark fringe, and 2) the important quantity is not the absolute noise $\delta P(t)$, but rather the *relative* quantity $\frac{\delta P(t)}{P_0}$.

Now, let's consider what happens when we introduce power recycling and arm cavities. In this case, switching to frequency space, (2.46) becomes^{9,10}

$$P_{\text{out}}^{\text{PRFPMI}}(f) \approx G_{\text{PRC}} P_0 \frac{\delta\phi^2}{4} + \left(\frac{G_{\text{PRC}}}{1 + \frac{if}{f_{p,cc}}} \right) \frac{\delta\phi^2}{4} \delta P(f) + \left(\frac{G_{\text{PRC}} G_{\text{arm}}}{1 + \frac{if}{f_p}} \right) 2kLP_0 \delta\phi h(f), \quad (2.48)$$

where here and for most of the rest of this thesis we will use the natural Fourier frequency, $f = \frac{\omega}{2\pi}$. Comparing (2.46), (2.48) and (2.47), we discover:

$$\frac{h_{\text{app.}}^{\text{PRFPMI}}(f)}{h_{\text{app.}}^{\text{MI}}(f)} = \frac{1}{G_{\text{arm}}} \left(\frac{1 + \frac{if}{f_p}}{1 + \frac{if}{f_{p,cc}}} \right). \quad (2.49)$$

Apart from the expected factor of G_{arm} due to the arm cavities, we also see an improvement in the region $f_{p,cc} < f < f_p$, where the coupling of laser power fluctuations to the output has started falling off due to the coupled cavity pole, but the interferometer's response to GWs has not started falling off due to the arm cavity pole. In this case, then, the reduction of the coupled cavity bandwidth has helped us. This improvement in laser power noise coupling is plotted in Fig. 2.5.

2.1.1.7 Signal recycling/extraction

Using the common/differential-mode selectivity of the Michelson on its dark fringe to create two decoupled optical systems begs the question: by extension of the power recycling concept, can we apply the same technique to the differential mode by adding a new mirror at the dark port, as shown in Fig. 2.7? The answer is “yes.”

To understand what is possible with this extra **signal recycling mirror (SRM)**, it is helpful to consider the power recycling case from a different angle. In the previous section, we determined $f_{p,cc}$ heuristically by appealing to the gain-bandwidth relation. However, we can arrive at the coupled cavity dynamics more directly by looking at it as a cavity made of the ETM on one side, and a compound mirror formed by the ITM and the PRM on the other side, as shown in Fig. 2.6.¹¹

⁹While it is a power quantity, the middle term here proportional to $\delta P(f)$ is still modified by the *amplitude* factor in the parentheses. This is because $\delta P(t)$ comes from a cross term between the DC carrier field and a fluctuating amplitude modulation SB. Only the latter part of the product carries frequency dependence, while both carry the DC amplitude gain of $\sqrt{G_{\text{PRC}}}$.

¹⁰The phase information here is inserted by analogy with that required by the Kramers-Kronig relations for the minimum-phase single-pole low-pass filter that the cavities approximate in the low-frequency limit.

¹¹Here, we are treating the common mode of the arms as a single effective cavity, and therefore the BS is essentially nonexistent.




Figure 2.5: Transfer function from input laser power noise to apparent GW signal in a PRFPMI, normalized by that for a simple Michelson.




Figure 2.6: Effective model for the common-mode coupled cavity.

Using (2.18), we can calculate the power transmission of the PRC as

$$T_{\text{PRC}}(\omega) = t_{\text{PRC}}^2(\omega) = \frac{T_{\text{PRM}}T_{\text{ITM}}}{1 + R_{\text{PRM}}R_{\text{ITM}} - 2\sqrt{R_{\text{PRM}}R_{\text{ITM}}}\cos\left(\frac{2\omega L_{\text{PRC}}}{c}\right)}. \quad (2.50)$$

When the PRC is anti-resonant, this becomes

$$T_{\text{PRC}}\Big|_{\text{anti-res}} \approx \frac{T_{\text{PRM}}T_{\text{ITM}}}{4}. \quad (2.51)$$

Now, in this picture, the coupled cavity uses the anti-resonant PRC as compound mirror designed to critically couple the input light to the arm losses ϵ_{arm} . Therefore, we expect the finesse to be

$$\mathcal{F}_{cc} = \frac{2\pi}{T_{\text{PRC}} + \epsilon_{\text{arm}}} = \frac{2\pi}{2T_{\text{PRC}}} \approx \frac{4\pi}{T_{\text{PRM}}T_{\text{ITM}}}, \quad (2.52)$$

and, since this is a critically coupled cavity, the power gain is

$$G_{cc} = \frac{\mathcal{F}_{cc}}{\pi} \approx \frac{4}{T_{\text{PRM}}T_{\text{ITM}}}. \quad (2.53)$$

If the two pictures are consistent, this should be equal to $G_{\text{PRC}}G_{\text{arm}}$. Using (2.29) and (2.30), we indeed find:

$$G_{\text{PRC}}G_{\text{arm}} = \frac{1}{T_{\text{PRM}}} \frac{4}{T_{\text{ITM}}} = G_{cc}. \quad (2.54)$$

As for the coupled cavity pole, we can compute the expected value as

$$f_{p,cc} = \frac{\mathcal{F}_{\text{arm}}}{\mathcal{F}_{cc}} f_p = \frac{2\pi}{T_{\text{ITM}}} \frac{T_{\text{PRM}}T_{\text{ITM}}}{4\pi} = \frac{T_{\text{PRM}}}{2} f_p = \frac{f_p}{2G_{\text{PRC}}}, \quad (2.55)$$

which is what we found in (2.43).

From this new perspective, we can see how adding an SRM will affect the differential mode. For example, if we choose $T_{\text{SRM}} = T_{\text{PRM}}$ and the same anti-resonant condition for the SRC, then the differential-mode signal sidebands will experience the same augmentation as the carrier field and the common-mode fluctuation SBs. Namely, they will be enhanced at low frequencies by the extra factor $\sqrt{G_{\text{PRC}}}$, but they will also be band-limited by the same coupled cavity pole $f_{p,cc} = \frac{f_p}{2G_{\text{PRC}}}$. This application is the eponymous **signal recycling**, and it has been used in the non-Fabry-Pérot-enhanced GW interferometer GEO600 [21]¹².

On the other hand, if we make the SRC *resonant*, this makes it highly transmissive from the dark port into the arms. In this case, the differential coupled cavity is

¹²Due to GEO600's relatively short arms and lack of arm cavities, true signal recycling results in manageable recycling cavity pole frequencies.

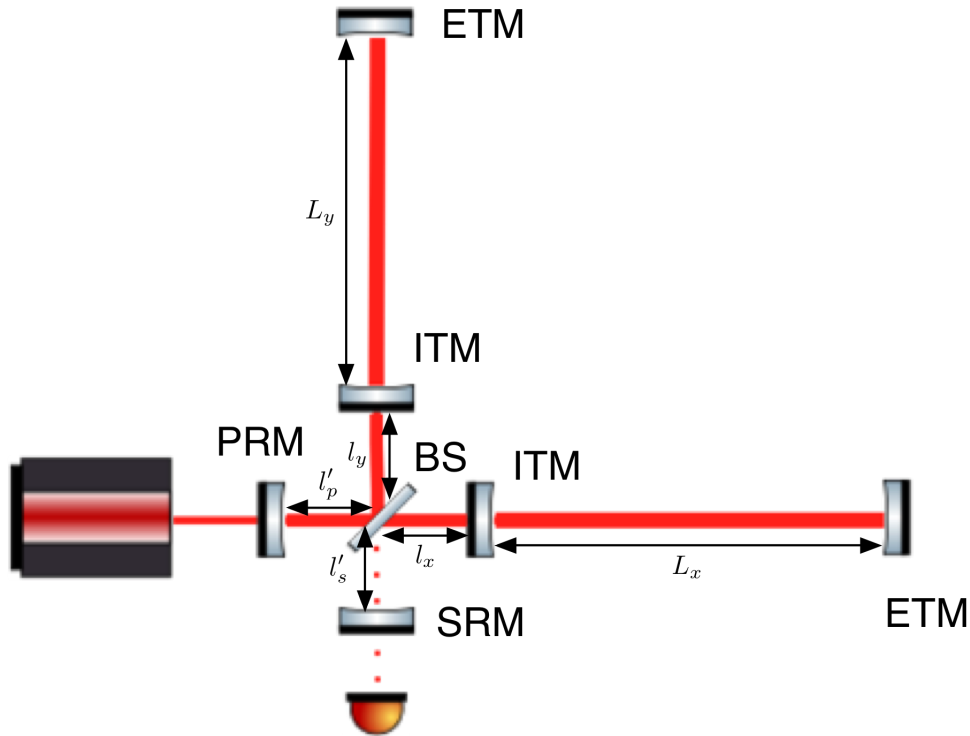


Figure 2.7: The full core aLIGO optical configuration: a dual-recycled Fabry-Pérot Michelson interferometer (DRFPMI). The lengths L_x , L_y , l_x , l_y , l'_p and l'_s will be used to define the interferometer's length DoFs in Sec. 2.1.2.4.

strongly overcoupled with very low finesse. Therefore, the low-frequency gain will be reduced by the factor $\sqrt{G_{\text{SRC}}}$ —less than one due to the anti-resonant state—and the bandwidth will be increased in proportion to $\frac{1}{G_{\text{SRC}}}$. Since, in this case, the SRC would be a relatively high-finesse cavity, the reduction in gain would be too great, even in light of the increased signal bandwidth. However, by choosing $T_{\text{SRM}} \gg T_{\text{PRM}}$, we can use this effect to a reasonable degree, increasing the differential coupled cavity pole frequency (with respect to the naked arm) by

$$\frac{f_{p,cd}}{f_p} = \frac{1}{G_{\text{SRC}}} \quad (2.56)$$

while only reducing the signal gain by the factor $\sqrt{G_{\text{SRC}}}$. This application is called **resonant sideband extraction (RSE)** [29], and it is how the SRM is primarily used in aLIGO. A significant advantage to this scheme is that it allows one to increase the finesse of the arm cavities—decreasing the power required on the BS for the same

sensitivity—without sacrificing bandwidth. This alleviates the thermal heating of the BS and ITM substrates, which can lead to optical aberrations.

Finally, the SRC can also be operated in a phase condition which is neither resonant nor anti-resonant. This configuration, called **detuned signal recycling**, has two main consequences. First, due to the displacement of the cavity resonance from the carrier frequency, the interferometer will have a heightened response to differential-mode signals of a frequency equal to this detuning (since either the upper or lower sideband will coincide with the coupled cavity resonance). This results in an interferometer response that is peaked at the detuning frequency, but diminished everywhere else with respect to the tuned RSE case. This allows for tuning the interferometer to sources of a specific, narrow spectral range.

Additionally, detuning causes a modification of the optomechanical dynamics of the system. [30, 31]. It is well known that an optical beam of power P exerts a force

$$F_{\text{RP}} = \frac{2P}{c} \quad (2.57)$$

on a mirror when reflected off of it at normal incidence. This **radiation pressure** effect exerts a static force

$$F_{\text{RP},0} = \frac{2P_{\text{circ}}}{c} = \frac{G_{\text{PRC}}G_{\text{arm}}P_0}{c} \quad (2.58)$$

on each of the test masses. Since P_{circ} is a function of the resonance condition, however, detuning of the arm cavities can modulate the radiation pressure force. This force modulation can in turn modulate the detuning, and so on and so forth. This position-dependent restoring force mimics a spring (or anti-spring), and is therefore known as an **optical spring**.

While the optical spring effect is more readily understood in the case of a single cavity, or with the common mode due to the carrier power dependence, the same principle applies to the differential signal sidebands, which beat with the strong carrier field and influence the interferometer's response to differential mode signals. Radiation pressure will be discussed more in Secs. 2.2.2 and 2.2.6.

2.1.1.8 Mode cleaning cavities

Finally, there are two independent cavities that serve as **mode cleaners**. As we have seen, Fabry-Pérot cavities can be used as optical filters, allowing only resonant fields to be transmitted, while spectral components outside the cavity bandwidth are

rejected. As we will discuss more later, cavities also act as spatial filters, as an input beam will only be transmitted if it matches the spatial mode of the cavity.

The first of these cavities, called the **input mode cleaner (IMC)**, is a longer, triangular, suspended cavity that sits between the laser and the PRM. Its purpose is to filter out noise on the input light before it enters the core interferometer (including amplitude noise, frequency/phase noise, pointing and polarization). The IMC will not be discussed much further in this thesis, but more details and an overview of the entire aLIGO input optics system can be found in [32].

The second, called the **output mode cleaner (OMC)**, is a shorter, monolithic cavity that sits between the SRM and the main readout photodetectors. Its purpose is to remove RF control sidebands that are intentionally placed on the laser, as well as so-called “junk light” created by aberrations in the core interferometer, before the final detection process. This thesis comprises—in part—the design, construction and testing of the aLIGO OMC, as well as its installation and integration at the LIGO Livingston Observatory (LLO). A description of this work, as well as a more detailed analysis of mode cleaners in general, is given in Part III.

A more accurate depiction of the core aLIGO optical layout, including the mode cleaners, folding mirrors inside of the PRC and SRC, and various other optical elements¹³, is shown in Fig 2.8. Each mirror (and nearly every other element) is held by a multiple-stage pendulum suspension, itself attached to a seismic platform equipped with passive and active isolation capabilities. These systems are discussed further in Sec. 2.1.2.3.

2.1.2 Systems and control

In this section, I will give a relatively brief overview of the general aLIGO control strategy, as well as of some of the major subsystems that conspire to make the overall machine work. This is *not* an exhaustive list; for example, no mention is made of the Angular Sensing and Control (ASC) subsystem that holds all the optical elements aligned with one another, nor of the Thermal Compensation System (TCS) [33] that measures and mitigates the effects of thermal loading on the interferometer optics, among many others. Excellent and more thorough treatments can be found in [15, 23, 34–36]. For those unfamiliar, an introduction to feedback control is given in Appendix A.

¹³See [23] for more details on these other elements.

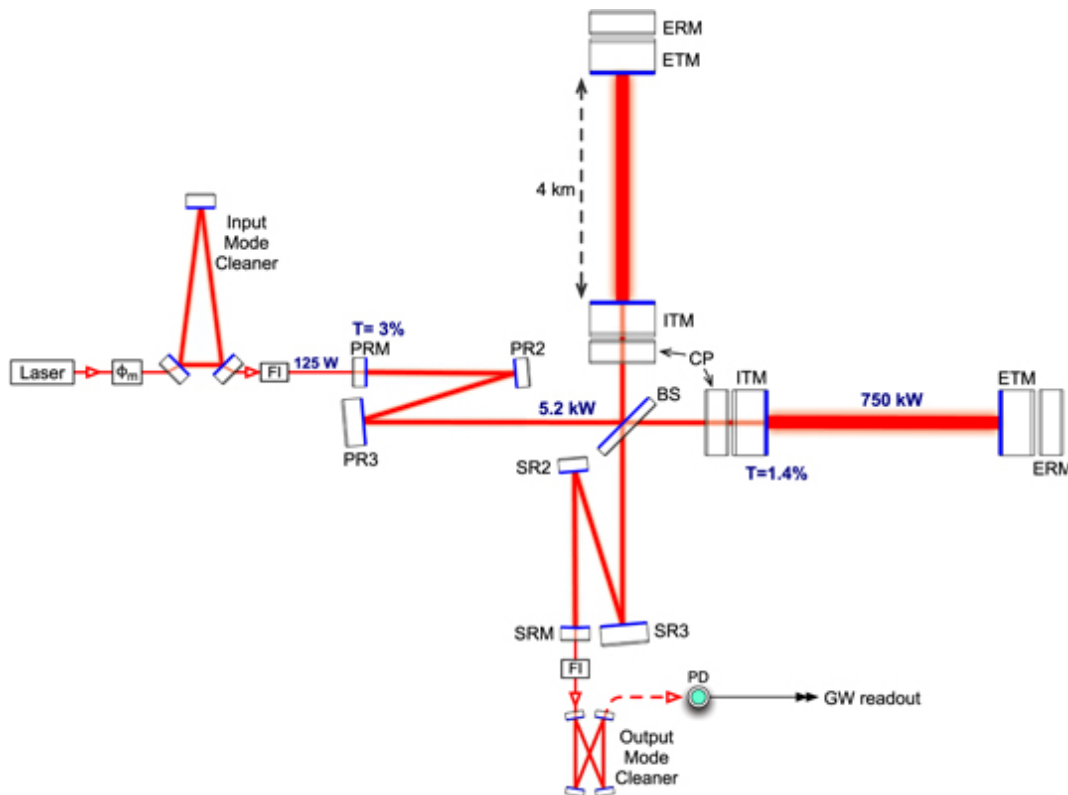


Figure 2.8: More accurate diagram of the aLIGO optical layout. (Image source: [23])

2.1.2.1 Control strategy

The aLIGO interferometer is controlled by an intricate network of control loops—many interconnected—each responsible for a particular **degree of freedom (DoF)**. The basic anatomy of a typical control loop is shown in Fig. 2.9. First, a sensor transduces some physical quantity (e.g., an optical power, an acceleration, a temperature, etc.) into a voltage. This voltage is then typically amplified and whitened¹⁴ at a remote rack near the sensor, and the amplified and whitened signal is carried over concentrated signal cables to a server room. Here, the signal is passed through an anti-aliasing filter to prevent folding of undersampled high-frequency content, then acquired into the digital system by an analog-to-digital converter (ADC).

Once in the digital domain, the signal can be manipulated by filtration and/or by combination with other signals (not shown in the example figure) before being re-converted into an analog voltage by a digital-to-analog converter (DAC). This output voltage is passed through an anti-imaging filter to remove any unwanted spectral

¹⁴Whitening is the process of compensating for the expected spectral shape of a signal to reduce the dynamic range required of the analog-to-digital converter.

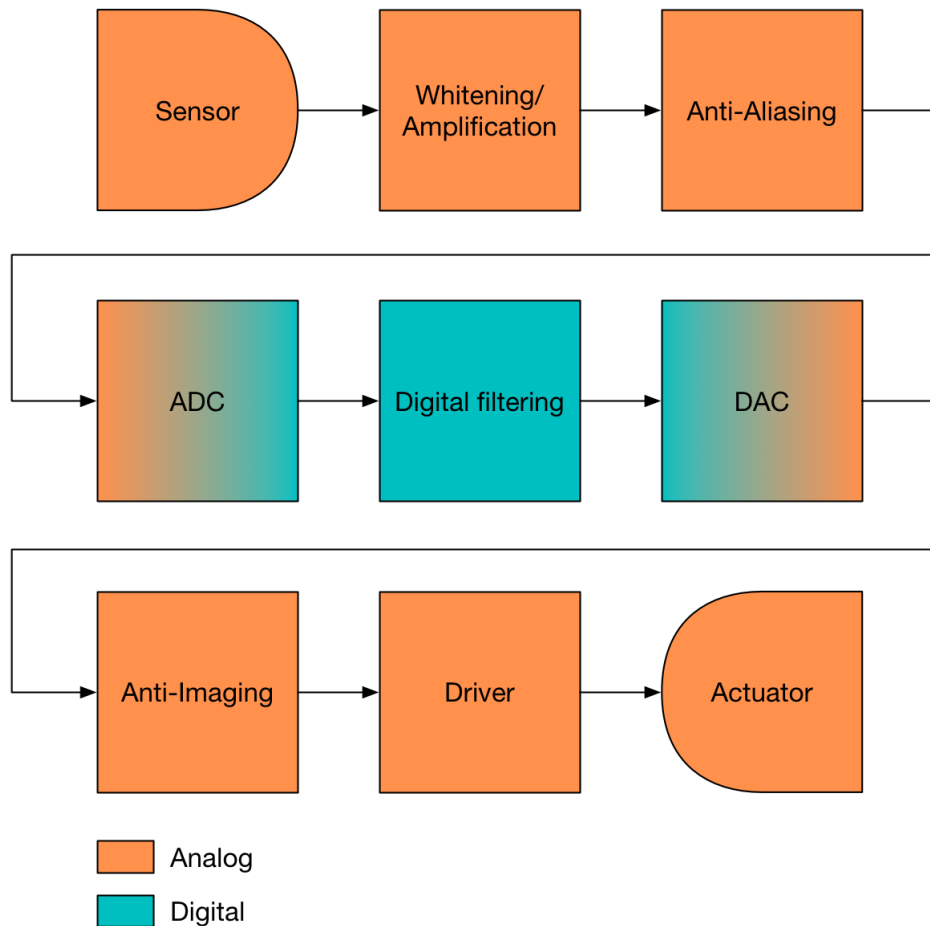


Figure 2.9: Simple block diagram of a typical digital control loop in aLIGO.

“image” components at high frequencies. Finally, the signal is piped out to a remote rack, where a driver performs analog amplification and filtering on it before applying it to an actuator, which converts a voltage to a physical quantity (e.g., a force via a PZT or magnet and wire loop, or an optical modulation via an applied electric field in an electro-optic crystal, etc.).

As digital loops can only operate up to frequencies a few times below the ADC/DAC conversion rates¹⁵, there are a few cases of purely analog control loops within LIGO, as discussed below.

¹⁵The presence of any analog components above the **Nyquist frequency** of $\frac{F_s}{2}$, where F_s is the sampling rate, leads to signal distortion from frequency folding. Therefore, the anti-aliasing filters must add a sharp roll-off at frequencies at and slightly below this frequency.

2.1.2.2 Pre-stabilized laser (PSL)

The system that generates and prepares the light at the input of the interferometer is called the **pre-stabilized laser (PSL)** [37]. As shown in the diagram in Fig. 2.10, the PSL system contains several interwoven loops.

At the very beginning of the system is a seed laser. This laser, an Nd:YAG non-planar ring oscillator (NPRO), is a solid-state source operating with a wavelength of 1064 nm. The seed light passes once through a traveling-wave amplifier that uses pumped gain elements to increase the beam's power. This pre-amplified laser is used to injection lock a high-power oscillator (HPO), which becomes modelocked to the seed laser via optical feedback. Over this multi-stage scheme, the seed laser's power of ~ 1 W is amplified to ~ 200 W.

The output of the HPO then goes into a small, rigid mode-cleaning cavity called the **pre mode cleaner (PMC)**, whose primary output is fed to the IMC and into the main interferometer. The other outputs of the PMC are used as parts of the intensity (ISS) and frequency (FSS) stabilization servo loops. In particular, one output is directed onto a photodetector (PD₂), whose signal is used to feed back to the intensity-modulating acousto-optic modulator AOM₁¹⁶. The other auxiliary output is shifted in frequency by the modulator AOM₂, and then this shifted beam is locked to a stable, rigid reference cavity in a thermally controlled vacuum chamber via feedback to the seed laser. At various points, electro-optic modulators (EOMs) are used to add RF phase modulation sidebands to the carrier field. These sidebands are used to derive error signals for keeping the various interferometer cavities locked, as discussed in the next section.

Whereas the power of the PSL must simply be stabilized to a constant value, its frequency—in addition to being stabilized to a very high degree—must also be *tunable*, so that it can be made to resonate in the main interferometer. First of all, the output of the PSL must be locked to the IMC. At low frequencies, this is accomplished by feeding back to the position of the IMC mirrors, pushing on them with magnetic actuators; at high frequencies, where it is harder to do this due to the rolloff of the mechanical response, the laser frequency can be adjusted by changing the frequency offset applied by AOM₂ (since, to the frequency stabilization servo, this looks as if the reference cavity has moved in frequency, and it will push the main laser frequency to compensate).

¹⁶This “inner loop” signal is combined with an “outer loop” signal derived from a pickoff of the beam exiting the IMC.

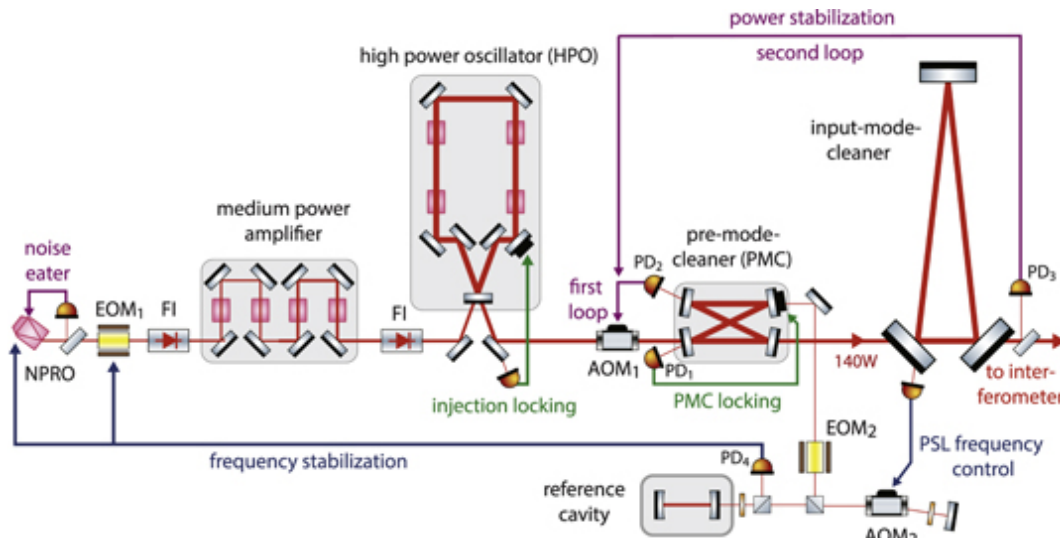


Figure 2.10: Advanced LIGO PSL. (Image source: [23])

Once resonant in—and therefore transmitted through—the IMC, the PSL must resonate in the main interferometer. This is accomplished partly from the laser side via two actuation points. At low frequencies, the IMC length can be adjusted intentionally, causing the laser to track its motion.¹⁷ At higher frequencies, IMC length feedback is not possible. In this regime, the relatively broad linewidth of the IMC (~ 20 kHz) is taken advantage of by adding a small offset into the IMC loop error point. As a result, the PSL frequency can be modulated minutely around the center of the IMC resonance without losing lock. The combination of these PSL actuation knobs with others in the main interferometer is discussed more in Sec. 2.1.2.4.

2.1.2.3 Seismic platforms and suspensions

Nearly every element in aLIGO is located within a vacuum chamber, bolted to a steel table—known as an internal seismic isolation (ISI) table—that provides passive and active electro-mechanical isolation. These tables are connected via bellows at the vacuum wall to hydraulic external pre-isolation (HEPI) platforms, which are rigidly connected to the earth. The seismic isolation system uses a network of acceleration and position sensors, as well as control inputs from the interferometer length control system (see Sec. 2.1.2.4), to isolate the instrument from external forces

¹⁷Note that, to use this actuation method, the low-frequency feedback of the FSS to the IMC length must be disabled, and the AOM feedback must be used to keep the IMC locked at all frequencies. Otherwise, the IMC loop will act to suppress this intentional pushing on the IMC in the force domain.

using feedforward and feedback. This isolation system as a whole is diagrammed in Fig. 2.11.

Some more general detail about this system, as well as a very detailed description of work that was performed as part of this thesis to address the issue of environmental tilt coupling within the aLIGO seismic isolation system, is given in Part II.




Figure 2.11: Advanced LIGO seismic isolation system, showing hydraulic external pre-isolator (HEPI) platforms and internal seismic isolation (ISI) tables. (Image source: [23])

Beginning shortly after the passage of the input laser beam from the PSL into the main interferometer vacuum system, every optical element it comes in contact with is held by a pendulum suspension with anywhere from one to four stages. To understand why, consider the simple pendulum shown in Fig. 2.12. Allowing the massless suspension point to move horizontally, we can calculate the equation of




Figure 2.12: A simple, single-stage pendulum.

motion of the mass at the end of the pendulum as:

$$m\ddot{x}_m = -mg \sin \theta \approx -mg\theta \approx -\frac{mg}{L}(x_m - x_p), \quad (2.59)$$

where g is the acceleration of gravity and we have assumed θ is small. Taking the Fourier transform and rearranging, we arrive at the transfer function

$$\frac{\tilde{x}_m}{\tilde{x}_p} = \frac{1}{1 - \left(\frac{\omega}{\omega_0}\right)^2}, \quad (2.60)$$

where we have defined the **resonance frequency**, $\omega_0 \equiv \sqrt{\frac{g}{L}}$. Thus, at low frequencies, the suspended mass follows very closely the motion of the platform, while at high frequencies, this motion is suppressed in the suspended mass as ω^{-2} . (This system appears to blow up at the resonance frequency, where the denominator above goes to zero. In reality, non-zero mechanical friction will provide damping that results in a finite amplification on resonance.) This pendulum effect can be compounded by adding a second pendulum suspended from the mass of the first, and another one to that, etc., such that the high-frequency isolation of the ultimate

mass relative to the top suspension point goes as f^{-2N} , where N is the number of pendulum stages.

Fig. 2.13 shows a diagram of the quadruple suspension used to isolate the main interferometer test masses. For reasons that will be discussed in Sec. 2.2.7, the test mass mirrors themselves, the “penultimate mass” dummy mirrors that sit one level up in the pendulum chain, and the four fibers that connect them to each other are all made out of fused silica; during installation, these components are all welded together as shown to form a single monolithic lower suspension assembly. The upper two stages are made of metal and connected to each other via more conventional steel wire. However, these metal stages have maraging steel blade springs above, between and below them for vertical isolation.

For global interferometer length control and/or to provide local damping of the high-quality coupled pendulum resonant modes, the quadruple pendulum system has actuators at every stage. On the upper three stages, integrated shadow sensors and magnetic actuators are coupled to magnetic flags for both displacement sensing and force actuation. On the very lowest stage, actuation is accomplished using an electrostatic drive. At all levels, forces are generated between the main pendulum and a parallel-hanging “reaction chain” to avoid a mechanical short-circuit to the ground via the actuators. Mirrors associated with less-sensitive length DoFs are held with all-metal suspensions (apart from the mirrors themselves) with one to three stages, and actuation is performed with respect to a rigid frame.

2.1.2.4 Length Sensing and Control (LSC)

The subsystem tasked with holding all the interferometer’s length degrees of freedom at their intended operating points is called, aptly, **Length Sensing and Control (LSC)**. Referring to the individual segment lengths called out in Fig. 2.7, the five length degrees of freedom of the main interferometer are defined in Table 2.1.

Four out of these five DoFs—CARM, MICH, PRCL and SRCL—are sensed using a heterodyne laser stabilization error signal generation scheme known as the **Pound-Drever-Hall (PDH)** technique [38], which is described in Appendix B. DARM, containing the GW signal, is sensed using a special form of homodyne readout known as **DC readout**, which is discussed further in Part III.

Of all the required resonance conditions, the most stringent is the common-mode arm length (CARM); this feedback loop, called the “common-mode servo” requires




Figure 2.13: Detail of the Advanced LIGO quadruple suspension. (Image source: [23])

DoF	Definition	Description
CARM	$L_+ \equiv \frac{L_x + L_y}{2}$	Common-mode arm length
DARM	$L_- \equiv \frac{L_x - L_y}{2}$	Differential-mode arm length
MICH	$l_- \equiv \frac{l_x - l_y}{2}$	Michelson differential length
PRCL	$l_p \equiv l'_p + \frac{l_x + l_y}{2}$	Power recycling cavity length
SRCL	$l_s \equiv l'_s + \frac{l_x + l_y}{2}$	Signal recycling cavity length

Table 2.1: Definitions of the five interferometer length degrees of freedom.

multiple actuators blended together at a sequence of crossover frequencies to achieve the required locking bandwidth (and, as a result, in-band gain). Using feedback to the laser frequency as described in Sec. 2.1.2.2, this loop achieves a UGF around 20-30 kHz. The other four loops, feeding back only to mirrors' positions, operate with UGFs between 20 and 200 Hz.

2.2 Noise

In this section, I will give an overview of the major noise sources that together set the sensitivity of aLIGO. For reference, the aLIGO design noise budget is plotted in Fig. 2.14¹⁸.




Figure 2.14: Advanced LIGO noise budget. (Image source: [23])

¹⁸Note that the ultimate expected sensitivity of aLIGO has evolved from this initial projection at the $\sim 30\%$ level.

2.2.1 Seismic and gravity-gradient (Newtonian) noise

Due to seismic activity, the ground is constantly moving at the level of $\sim 1 \mu\text{m}/\sqrt{\text{Hz}}$ around 0.1 Hz, decreasing somewhat at higher frequencies but not very quickly. Given the eventual target sensitivity of $\sim 2 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}}$ near 100 Hz, aggressive isolation is needed. This is accomplished using the sophisticated system described in Sec. 2.1.2.3.

Traditional seismic noise that couples to the test masses through the platforms can be mitigated in a number of ways:

- **Passive/active-passive isolation:** Using the mechanical transfer functions of the links between the ground and the mirrors (i.e., the HEPI platforms, the ISI tables, and the suspensions), the ground motion is naturally suppressed at high frequencies. In some cases, feedback loops are used to keep some of these mechanical links (e.g., the ISI stages) at a desired operating point, but they may still mimic a passive isolation transfer function.
- **Active feedback:** Blended signals from accelerometers, position sensors, and interferometric length measurements are filtered and fed back to the various actuators between the ground and the test masses.
- **Active feedforward:** Blended signals from seismometers and accelerometers sensing the ground motion are filtered and fed *forward* to the aforementioned actuators to cancel the ground motion before it passes further in towards the test masses.

In addition to coupling via mechanical linkages, seismic motion of the ground can also *directly* couple to the test masses by Newtonian gravitational attraction. As this coupling does not pass through the seismic platform and suspension, these do not afford any passive suppression of this **Newtonian noise** (or **gravity-gradient noise**). Therefore, to combat this noise, networks of ground sensors are used to estimate the forces on the test mass, and opposing forces are fed forward to the actuators in order to cancel them. Notably, this coupling also exists with any other moving massive object near the test masses (e.g., a fan, or even wildlife).

Much more information on traditional seismic noise and Newtonian noise can be found in [39, 40].

2.2.2 Laser noise

The laser that illuminates the interferometer has noise in both its intensity (power) and frequency. As discussed in Sec. 2.1.2.2, it is the job of the pre-stabilized laser (PSL) system to suppress this noise using hierarchical feedback loops [37]. The degree of suppression required is determined by computing the coupling of both kinds of laser noise to the GW channel.

The most straightforward coupling path of laser intensity noise is via the small offset from the dark fringe required for DC readout (see Sec. 5.1). Due to the interference condition, some amount of light is directed to the AS port after emerging from the arms and recombining on the beamsplitter. Thus—with the frequency shaping described in Sec. 2.1.1.6 due to the coupled cavity pole—the remaining intensity fluctuations entering the interferometer from the PSL are passed directly to the GW detection point, in proportion to the fraction of DC optical power that is allowed to pass there from the input.

Less directly, laser power fluctuations can couple to the GW channel in at least two additional ways [27]. For one, mismatches in circulating powers in the two arms (e.g., from mismatches between the two ITM reflectivities) can lead to differential force noise between the two arm cavities via unbalanced radiation pressure (see Sec. 2.1.1.7). Finally, the high optical power passing through the ITM substrates creates a considerable thermal lens [33]. As this effectively changes the supported spatial mode of the arm cavity as seen by the input laser¹⁹, the result is that a portion of the light does not enter the arms. Without resonating in the arms, the light does not experience the coupled-cavity spectral filtering, and this strongly increases the coupling of laser intensity noise to the dark port. In practice, this is the dominant effect at high frequencies.

As with intensity noise, direct coupling of frequency noise is minimal due to the operation of the interferometer very near to the dark port. However, there is an intentional macroscopic (~ 8 cm) mismatch added to the arms of the Michelson so that the upper f_2 control sidebands are passed to the dark port for control purposes, while the carrier and f_1 sidebands are not. This slight offset, called the **Schnupp asymmetry**, partially spoils the otherwise near-perfect cancellation of frequency noise coupling to the dark port.

As a result of the strong filtering provided by the PSL, the laser is not expected to be a limiting source of noise in aLIGO via the above coupling mechanisms. On the other

¹⁹See Sec. 5.2 for more about laser beam spatial modes and cavity coupling.

hand, some nonlinear coupling has been observed in the course of interferometer commissioning that threatens to undercut the design sensitivity in a narrow band; this is the subject of ongoing study as of the time of publication of this thesis.

2.2.3 Scattered light noise

The surfaces of the test masses, while extremely uniform and smooth, are not perfect. As a result, some fraction of the light incident on each mirror on each pass is **scattered** out of the cavity spatial mode and onto some other path. Some of this light can then reflect off of some other surface (e.g., the long, narrow beam tube), and some of *that* light can be scattered *back* into the cavity mode. When this happens, since the surface of the second scattering can in general be that of an object that is not isolated from environmental noise, this scattered light can contribute a large phase fluctuation to the main beam. To combat this effect, scattered light **baffles** are placed at strategic positions around the interferometer to trap scattered light in such a way that it cannot rejoin the main beam.

Another source of scattered light noise is an effect called **fringe wrapping**. All interferometer optics have an anti-reflective (AR) coating on their non-intentionally-reflective sides. Despite this fact, there is always some finite reflectivity of the beam at each interface. This is not a terrible problem when the object in question is not moving very much with respect to the main interferometer optics. However, in cases when the object is, the relative velocity can be so great that the parasitic reflection undergoes phase modulations of $\gg 2\pi$ peak-to-peak. In this case, this fringe wrapping leads to noise upconversion from relatively low-frequency swinging motion up into the GW detection band. This noise manifests as a shelf with an amplitude determined by the reflectivity of the offending surface and a high-frequency cutoff determined by its velocity.

Scattered light noise is usually hunted down and removed by painstakingly driving the velocity of all suspected objects and looking for resultant shelves in the interferometer noise spectrum.

2.2.4 Residual gas noise

The entire aLIGO vacuum system is maintained at a pressure of below 10^{-6} Pa, as the presence of gas molecules in the system can contribute to the noise floor in at least two ways.

First, the traversal of gas molecules across the beam path can directly impart an

impulse to the laser phase. The net effect of many such transits is an overall stationary phase noise contribution that can be computed via statistical analysis [27].

Second, the presence of a finite concentration of gas molecules between the surfaces of the test masses and their respective reaction masses—separated only by about 5 mm—contributes a damping force on the motion of the test mass motion. As we will discuss at length in Sec. 2.2.7, this damping necessarily results in an associated noise force on the test mass.

2.2.5 Electronic noise

By and large, electronics in aLIGO—be they photodiodes, filters, analog-to-digital or digital-to-analog converters, amplifiers, actuators, or anything else—are designed so as not to contribute noise at or near the levels of the other major sources described here. In most cases, electronic units have at least two modes: a high-gain/range, high-noise mode used for lock acquisition, and a low-gain/range, low-noise mode for GW detection.

One particularly strong coupling of electronic noise is due the electrostatic drive (ESD) used on the lowest stage of the quadruple suspensions (i.e., to drive the test masses directly). By design, a large DC voltage is applied gold electrodes patterned onto the surface of the reaction mass, causing a temporary charge to be induced on the nearby surface of the test mass. Then, a smaller AC voltage signal is summed onto the same electrodes, leading to a corresponding AC force between the reaction and test masses, enabling actuation.

Due to various effects, the test mass can be embedded with a net charge even in the absence of the ESD bias voltage. When this happens, it leads to gain offset and imbalance errors in actuation, which can cause calibration issues and instability. Additionally, the permanent charge on the surface of the mirrors can interact with potential variations on nearby grounded conductors, leading directly to force noise on the mirrors.

In practice, care is taken to discharge the test masses with ion guns before pumping down whenever they are exposed to air [41].

2.2.6 Quantum noise

It is well known that the quantum-mechanical nature of light imposes an ultimate limit to the sensitivity of any optical measurement. On the one hand, the fact that a light beam is composed of discrete photons, each passing any given point

randomly, independent of the ones before or after it, leads to an inescapable “photon counting” noise, or **shot noise**. This noise manifests itself as a fluctuation in power proportional to the square root of the average power, $\delta P_{\text{shot}} \propto \sqrt{P}$, such that the signal-to-noise ratio is also proportional to \sqrt{P} (since $\text{SNR} \propto \frac{P}{\delta P_{\text{shot}}}$). On the other hand, this shot noise power fluctuation leads to a corresponding radiation pressure force fluctuation on any object from which the light is reflected. This **quantum radiation pressure force** is an example of **quantum back-action**: by attempting to increase our measurement sensitivity with respect to shot noise by increasing the interrogating laser power, we inevitably act back on the measured system and get in our own way.

Historically, shot noise was believed to arise from randomness in the photodetection process. Caves [42] first demonstrated that the randomness was actually an intrinsic property of coherent light: **vacuum fluctuations**—the ground state of the electromagnetic field—enter the interferometer from the dark port, interact with the carrier entering from the bright port, and affect both the sensing of the test masses and the back-action on them. In doing so, he showed that the combined effect of shot noise and quantum radiation pressure noise can be seen as an enforcement of the Heisenberg uncertainty principle in the measurement of the positions of the test masses. Previously, Braginsky [43] had formulated the Heisenberg limit in the context of precision measurements of mechanical oscillators by defining the **standard quantum limit (SQL)**, which, when expressed as the spectral strain sensitivity of a GW interferometer, is

$$h_{\text{SQL}} \equiv \sqrt{S_h^{\text{SQL}}} = \sqrt{\frac{8\hbar}{m\Omega^2 L^2}} \left[\frac{1}{\sqrt{\text{Hz}}} \right], \quad (2.61)$$

where \hbar is Planck’s constant, m is the mass of each individual mirror, $\Omega = 2\pi f$ is the angular Fourier frequency, and L is the length of the interferometer’s arms. Traditionally, an interferometer’s sensitivity only reaches the SQL at one particular frequency—namely, the frequency at which shot noise and quantum radiation pressure noise intersect, as shown in Fig. 2.15.

The overall quantum noise spectrum can be written in terms of the SQL as [44]

$$h_{\text{QN}} = \sqrt{\frac{h_{\text{SQL}}^2}{2} \left(\frac{1}{\mathcal{K}} + \mathcal{K} \right)}, \quad (2.62)$$

where

$$\mathcal{K} \equiv \frac{8P\omega_0}{mL^2\Omega^2(\gamma^2 + \Omega^2)}, \quad (2.63)$$

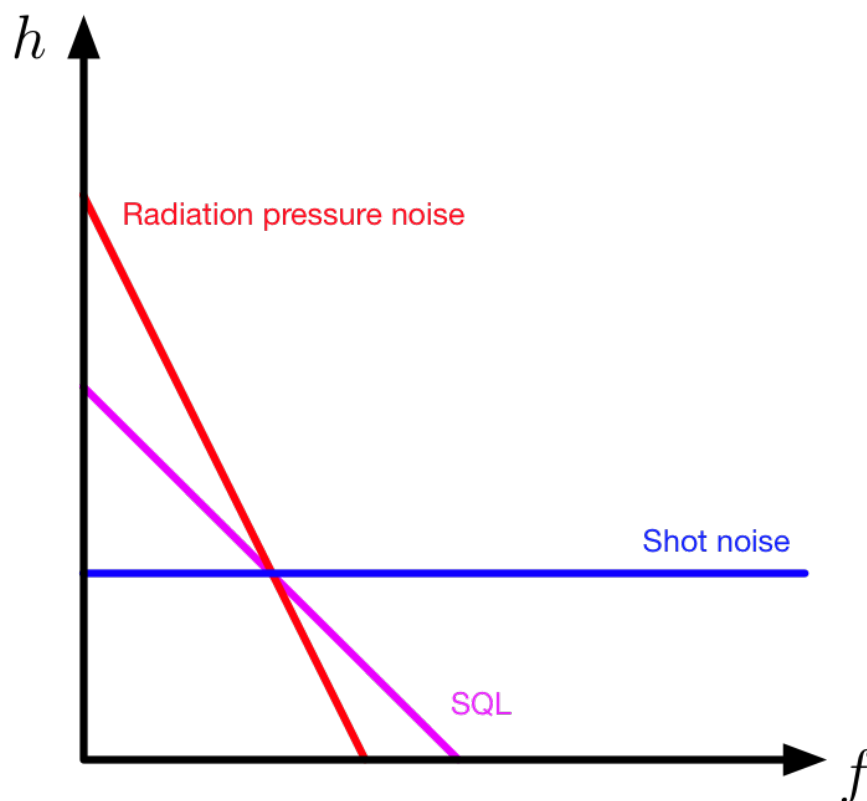


Figure 2.15: Quantum noise spectral density, showing the two components: quantum radiation pressure noise, which has a f^{-2} dependence for a free mass, and shot noise, which is white. The sensitivity at the intersection of these two lies on the standard quantum limit (SQL), which decreases as f^{-1} . (At higher frequencies, the shot noise curve will turn up and increase as f due to the limited detection bandwidth of the interferometer, but this is not shown here.)

with P being the laser input power, ω_0 its optical angular frequency, and γ being the detector half-bandwidth (i.e., the cavity pole angular frequency), is called the “optomechanical coupling constant”. The term proportional to \mathcal{K} in (2.62) is the radiation pressure, while the one proportional to \mathcal{K}^{-1} is the shot noise. Changing the laser power or wavelength can move these two contributions up and down, so that the SQL is met at a higher or lower frequency, but this still only happens at one frequency for the standard interferometer (i.e., when $\mathcal{K} = 1$).

As can be seen in 2.14, quantum noise limits the aLIGO design sensitivity nearly everywhere above the seismic wall.

2.2.7 Thermal noise

Systems at a finite temperature are, by definition, always in motion. As a result, any measurement of an object at a finite temperature will at some level be influenced by the microscopic motion of the molecules it comprises. As we will see, the amount of thermal motion exhibited by a system at a given temperature depends on internal material properties of that system.

In the early 19th century, botanist Robert Brown observed the random motion of pollen grains suspended in water [45]. He could not ascertain the origin of this motion, but he did conclude that it was not the result of any living organism. In 1905, Einstein showed that this **Brownian motion** was actually the result of the average interaction between the pollen grains and the molecules of the water in which they were suspended [46]. In particular, he showed that the diffusion constant of the grains, D , was proportional to the drag on them in the water:

$$D = \mu k_B T, \quad (2.64)$$

where k_B is Boltzmann's constant, T is temperature, and the mobility μ characterizes the drag, being the ratio of terminal drift velocity to applied force ($\mu = \frac{v_d}{F}$).

Some years later, Johnson [47] observed a new, temperature-dependent noise in resistors,

$$V_n^{\text{rms}} = \sqrt{4k_B T R \Delta f}, \quad (2.65)$$

where R is the resistance and Δf is the bandwidth of the measurement. Subsequently, Nyquist [48] was able to explain this noise by analogy with Einstein's result for Brownian motion, since resistance can be regarded as a sort of drag on an electron flowing in a conductor.

Eventually, Callen [49] developed a generalized model relating the noise of a system to dissipation within it. This **fluctuation-dissipation theorem (FDT)** can be written

$$S_F(\omega) = 4k_B T \text{Re}\{Z(\omega)\}, \quad (2.66)$$

where $S_F(\omega)$ is the force noise power spectral density, and $Z(\omega) \equiv \frac{F(\omega)}{v(\omega)}$ is the **mechanical impedance** of the system. Equivalently, the relation can be written in terms of the displacement noise power spectral density as

$$S_x(\omega) = \frac{4k_B T \text{Re}\{Y(\omega)\}}{\omega^2}, \quad (2.67)$$

where $Y(\omega) \equiv Z(\omega)^{-1}$ is known as the **mechanical admittance**.

This relationship can be understood via a symmetry argument as follows: when a system with dissipation is displaced, some amount of the energy stored in it is lost as heat to its environment; therefore, when the system exchanges heat with its environment, this must conversely induce a displacement.

Thus, the FDT provides a straightforward means of predicting the thermal noise of a system by measuring its dissipation. We will study this much further in Part V.

Thermal noise affects all mechanical components of the interferometer, however, it is most significant near and within the test masses themselves. For example, thermal motion of the fused silica fibers in the final stage of the suspension contributes significantly to the aLIGO noise floor in a narrow band just above the low-frequency seismic wall²⁰. More critically, thermal motion in the high-reflectivity optical coatings on the faces of the test masses—made of a thin $\text{SiO}_2/\text{Ta}_2\text{O}_5$ stack with much higher dissipation than that of the fused silica mirror substrate—limits aLIGO sensitivity in its most sensitive band around 100 Hz.

2.3 The path forward

There are several upgrades to Advanced LIGO planned for the future [50].

A short-term upgrade known as “A+” will be installed in the next few years. This will be an incremental upgrade to the aLIGO detectors, reusing most of the same optical and mechanical elements, but with some new, lower-loss coatings and a novel readout scheme.

A longer-term upgrade known as “LIGO Voyager” [51] will be installed near the end of the next decade. This will be a major upgrade, including all-new optical and mechanical elements, and pushing the existing facilities to their limits.

This section gives a brief overview of the path towards reducing the influence of the primary noise sources that limit aLIGO sensitivity: quantum noise and thermal noise.

2.3.1 Reducing noise in LIGO upgrades

2.3.1.1 Quantum noise: heavier mirrors, higher power and squeezed light

A quick glance at (2.61) reveals only two ways to improve the SQL strain sensitivity: 1) increase the length of the interferometer, or 2) increase the mass of the mirrors.

²⁰It is for this reason that the final two stages of the test mass quadruple suspension are constructed monolithically out of fused silica, as fused silica exhibits far less internal friction than metals do.

The former is not possible using the existing LIGO facilities, but the latter is. Note that, in order to reach the SQL at the same frequency after increasing the mass, the laser power must be increased commensurately, according to (2.62) and (2.63).

Additionally, shortly after describing the link between shot noise and quantum radiation pressure noise, Caves [52] also proposed injecting an exotic state of light known as a **squeezed state**—as opposed to the vacuum fluctuations—into the dark port, showing that this state could improve a detector’s sensitivity.

This squeezing concept is illustrated in Fig. 2.16. In each panel, a phasor diagram is shown, with the two axes representing the two orthogonal quadratures²¹ of the field exiting the interferometer. The blue phasor is an example signal, such as would be generated by a GW or displacement signal. The red ball is called the **noise ellipse**, and it quantifies the distribution of quantum noise²² over these two quadratures. Roughly speaking, the SNR of a measurement can be determined by taking the ratio of the signal phasor length to the radial extent of the noise ellipse along the direction of the signal phasor. For the vacuum state in panel (a), the noise ellipse has the same extent in the phase quadrature as it does in the amplitude quadrature—i.e., it is a circle. Using a variety of techniques, the vacuum state can be converted into either a phase-squeezed (b) or amplitude-squeezed (c) state. In doing so, the fluctuations in one quadrature can be reduced at the expense of the other (crucially, the *area* of this ellipse is constant for any minimum-noise state—reducing noise in the phase quadrature necessarily increases it in the amplitude quadrature, and vice versa). Note that this allows for “beating” the SQL over a narrow band, depending on the achievable input power and squeezing magnitude.

Now, field components at different sideband frequencies relative to the carrier can exhibit varying amounts of squeezing and a varying squeeze angle as a function of frequency. This can happen as a result of frequency dependence in the initial squeezing mechanism, *or*, it can happen via interaction of a field that starts with frequency-independent (FI) squeezing with a phase-rotating system. In particular, the aLIGO interferometer converts amplitude fluctuations entering the dark port into phase fluctuations re-emerging from it (along with the desired signal in the phase quadrature)²³ Since aLIGO operates at the SQL, this means that introducing

²¹These two quadratures are the phase and amplitude quadratures, corresponding to sidebands that generate phase or amplitude fluctuations in the sum field, respectively.

²²Noise ellipses can, in general, be used to quantify any kind of noise; here, they are used exclusively to represent quantum noise.

²³This is because amplitude-quadrature fluctuations are responsible for quantum radiation pres-




Figure 2.16: Squeezed states. The noise ellipse (red ball) is shown along with a phase-quadrature signal (blue phasor) for three cases: (a) vacuum state, (b) phase squeezed state, and (c) amplitude squeezed state.

phase-squeezed light will reduce quantum noise at high frequencies (shot noise) while *increasing* it at lower frequencies (radiation pressure); introducing amplitude-squeezed light will have the opposite effect.

Frequency-independent squeezing has already been demonstrated [53–55] on prototype and first-generation GW detectors, which operated above the SQL and could therefore benefit from FI phase squeezing. For aLIGO, other second-generation detectors, and their successors, a means of achieving frequency-dependent (FD) squeezing is necessary. Such a means was proposed in [44]: by reflecting an input field prepared with FI squeezing off of a detuned high-finesse **filter cavity**, high-frequency sidebands can be rotated due to interaction with the cavity so that their reflected counterparts are in the orthogonal phase. This technique was demonstrated in a prototype [56] and will be included in the A+ and LIGO Voyager designs to achieve quantum noise suppression across the active detector frequency range.

2.3.1.2 Thermal noise: materials, cryogenics and design optimization

The FDT asserts that there are two ways to reduce the thermal-noise-induced fluctuations in a system: 1) reduce the internal friction (dissipation) in the system, or 2) reduce its temperature. Upgrades to Advanced LIGO will make use of both of these strategies.

sure noise on the mirrors, which induces real motion in the mirrors, which then produces phase-quadrature signals that add noise to the main readout channel.

A+ will include only modest improvements in thermal noise. The primary modification will be to reduce the internal dissipation of the optical coatings on the test masses by a factor of ~ 4 , leading to a strain sensitivity improvement of a factor of ~ 2 . A+ will continue to operate at room temperature.

LIGO Voyager will address thermal noise much more substantially. For one thing, it will operate cryogenically at a temperature of around 120 K. Furthermore, because it will operate cryogenically and for other reasons that will be explained in the next section, the test masses and suspensions will be made of monocrystalline silicon, instead of fused silica, and this change of materials is accompanied by thermal noise benefits. Finally, the $\text{SiO}_2/\text{Ta}_2\text{O}_5$ optical coatings will be replaced by stacks made of SiO_2 and amorphous silicon or SiN, which have been demonstrated to have considerably lower dissipation [57].

Beyond improving materials and operating at lower temperatures, thermal noise can be further minimized by careful optimization of certain design parameters (e.g., the aspect ratios of suspension fibers/ribbons or of the cylindrical test masses, the distribution of layer thicknesses within the optical coating stacks, etc.). These optimizations are performed in multiple areas of the upgrade designs.

2.3.2 Putting it all together: the case for silicon in LIGO Voyager

Silicon was chosen as the LIGO Voyager test mass and suspension material as a result of a variety of considerations. The reasoning proceeds roughly as follows:

- In order to improve quantum noise for a given achieved level of squeezing, heavier test masses are desired.
- In order to reach the SQL at the same target frequency with heavier masses, a higher laser power is desired.
- Advanced LIGO, with fused silica mirrors and operating at ~ 1 -MW circulating powers, is strongly affected by thermally induced aberrations, and operating at higher powers with larger mirrors would likely be intractable. Monocrystalline silicon has a considerably higher thermal conductivity than fused silica does, and is therefore better suited for this level of thermal deposition.
- Additionally, the coefficient of thermal expansion (CTE) of silicon has a zero crossing at ~ 123 K (see Fig. 2.17). This results in two major benefits of

operating at this temperature, in addition to the direct reduction of thermal noise: 1) any small thermal gradients present in the test masses near this temperature will produce minimal thermal lensing, since such lensing is proportional to the CTE, and 2) the system will be free from a particular kind of thermal-noise-inducing dissipation known as **thermoelastic loss**, which will be discussed in detail in Part V.

- Last but not least, using silicon and operating at lower temperatures allows for the use of amorphous silicon coatings, which is not possible with fused silica substrates or at room temperature.




Figure 2.17: Coefficient of thermal expansion (CTE) of monocrystalline silicon as a function of temperature. Note the zero crossing at ~ 123 K. (Source: NIST/CODATA Swenson, 1983, via [51])

2.3.3 Outstanding questions about silicon

In order to make the transition away from fused silica, several outstanding questions must be answered about the material properties of silicon and about its availability in the dimensions and quality needed for use in a GW interferometer. A comprehensive

account of the major areas of research to this end can be found in [51]. Part V of this thesis concerns one aspect of this research: investigating the thermal noise in silicon ribbons similar to those that will be used in the LIGO Voyager suspensions.

Part II

Laser Gyroscope

This part is divided into two chapters. In Chapter 3, a short introduction to the issue of tilt coupling in GW interferometers is given, along with a summary of efforts to mitigate it in aLIGO. Chapter 4—adapted from [58]—contains a detailed description of the laser gyroscope built at Caltech, as well as a short history of laser gyroscopes in general.

Chapter 3

TILT SENSING IN ADVANCED LIGO

To suppress seismic noise, LIGO and other GW detectors use arrays of seismometers and accelerometers to sense the motion of the ground and then use intelligent filtering to feed this information forward to force actuators. In order for this scheme to work, the sensors must be truly measuring the motion of the ground. Unfortunately, such a measurement can be thwarted by spurious coupling to ground tilt.

To see this effect, consider a seismometer modeled by a mass-spring system in a box, as in Fig. 3.1. The device is supposed to measure linear ground acceleration of the box, which will produce a Hooke's law displacement of the inner mass relative to the box frame. However, if the ground tilts, the force of gravity can also cause a displacement, simulating a linear acceleration. If such a tilt occurs, the feedforward system will send the incorrect information to the force actuators.

Assuming a spring constant k , and calling the resting positions of the mass and the box x_m and x_b , respectively, the equation of motion of the mass is

$$m\ddot{x}_m = -k(x_b - x_m) + mg \sin \theta . \quad (3.1)$$

Taking the Fourier transform and assuming small θ , this becomes

$$-m\omega^2 \tilde{x}_m = -k(\tilde{x}_b - \tilde{x}_m) + mg\tilde{\theta} . \quad (3.2)$$

What we measure is actually the differential motion, $x_d \equiv x_m - x_b$. Substituting this in, we get

$$\tilde{x}_d = \frac{m\omega^2}{k - m\omega^2} \left(\tilde{x}_b + \frac{g}{\omega^2} \tilde{\theta} \right) . \quad (3.3)$$

From this, we can clearly see that the angular motion of the system can corrupt our inference of x_b . Judging by the weights of the \tilde{x}_b and $\tilde{\theta}$ terms, it is clear that the effect becomes more pronounced at lower frequencies. In particular, there is a response of $\tilde{x}_d(0) = \frac{mg}{k} \tilde{\theta}(0)$ at DC, while the response to x_b vanishes.

One solution to this problem is to sense tilt independently. Then, this tilt information can be implemented into the feedforward scheme to subtract the effect of tilt at low frequencies. Given the coupling calculated above and the required performance of the aLIGO seismic isolation system, Lantz, et al. [59] generated a tilt-sensing requirement, as shown in Fig. 3.2.

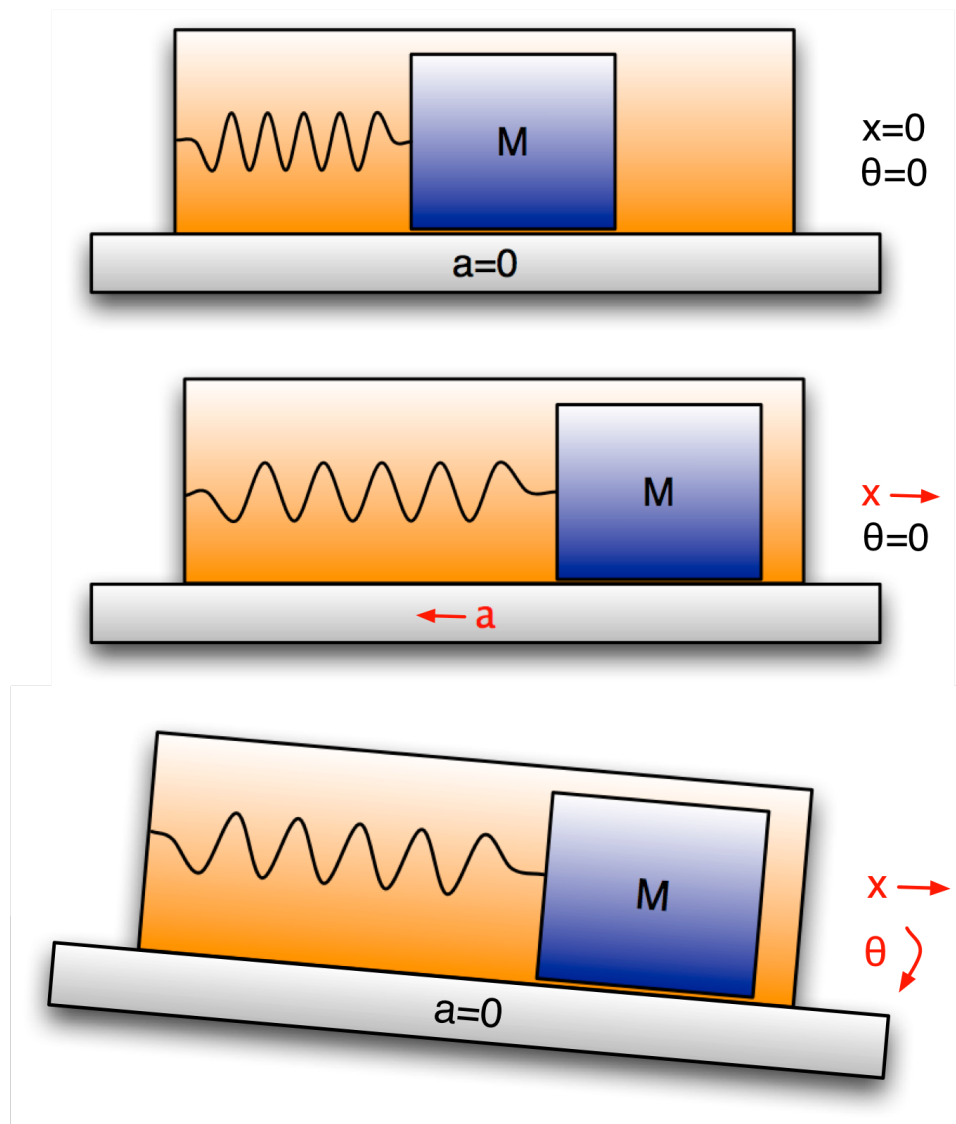


Figure 3.1: Simplified demonstration of the seismometer tilt-coupling problem.

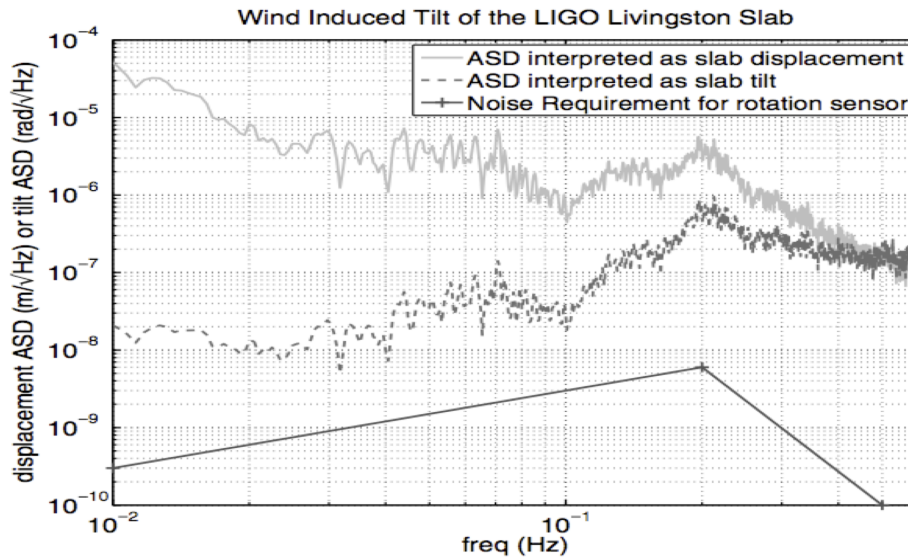


Figure 3.2: The aLIGO tilt sensing requirement definition, based on the seismic isolation requirements and the known seismometer tilt coupling. (Image source: [59])

Many avenues were pursued to find a suitable tilt sensor to serve this purpose in aLIGO. Ultimately, the winning solution was a mechanical balance beam sensor [60], which has performed well and which has also been used to make interesting seismological measurements [61]. The next chapter describes an effort to develop such a tilt sensor using a passive, free-space heterodyne laser gyroscope.

PASSIVE, FREE-SPACE HETERODYNE LASER GYROSCOPE

4.1 Introduction

4.1.1 Laser gyroscopes

The first use of a ring laser cavity to detect rotational motion was demonstrated by Macek and Davis [62] in 1963, and their design remains essentially unchanged in most current implementations. The operating principle for all optical gyroscopes is the *Sagnac effect*: in a ring geometry, if the system is rotating in the optical plane, the roundtrip optical path lengths traversed by two counter-propagating beams are unequal.

This earliest design was itself an improvement of a non-resonant, phase-sensitive interferometer introduced by Sagnac himself [63]. In that instrument, an interferometric fringe shift was produced at the output proportional to the rotation rate, and its sensitivity was therefore limited by the achievable fringe resolution¹. Conversion of the ring into a laser cavity created a bidirectional resonator, wherein the supported modes in each of the two directions—their frequencies being dependent on the respective roundtrip phases—are non-degenerate in the presence of rotation. This allowed Macek and Davis to use far more sensitive heterodyne techniques to measure the frequency splitting caused by rotation (at the time of their work, the achievable resolution was one part in 10^{12} , a significant improvement over the interferometric fringe readout). As reported in their original paper, the Sagnac-induced frequency shift is given by²

$$\Delta\nu = \frac{4}{\lambda S} \vec{A} \cdot \vec{\omega}, \quad (4.1)$$

where \vec{A} is the vector of area enclosed by the cavity, S is the cavity perimeter, λ is the laser wavelength, $\vec{\omega}$ is the angular velocity, and $\Delta\nu$ is the optical frequency splitting.

¹Another type of interferometric optical gyroscope, the *fiber optic gyroscope (FOG)*, is not discussed here. It is similar to Sagnac's original interferometer, but with the free-space system replaced by many windings of a fiber, increasing sensitivity. For an excellent contemporary review of all optical gyroscope technologies, see [64].

²The factor $\frac{4}{\lambda S} \vec{A} \cdot \hat{n}$, where \hat{n} is the unit vector associated with the angular velocity $\vec{\omega}$, is commonly referred to as the “scale factor.” The relation in Eq. 4.1 can easily be proven for a circular path and holds true for an arbitrary planar geometry.

The concept of an externally illuminated (“passive”) laser gyroscope was first presented by Ezekiel and Balsamo [65] in 1977. Previously, a major issue with the common active design had been discovered: at small rotation rates, backscatter-induced crosstalk effects caused the counter-propagating modes to lock to one another in frequency, leading to a null output³. It was believed that this effect was caused by the presence of the gain medium within the gyroscope cavity [65, 67]. Ezekiel and Balsamo sought therefore to circumvent this effect by locking an external laser to a passive optical ring cavity. In their setup, acousto-optic modulators (AOMs) were used to shift the laser frequency up macroscopically in common mode for the two counter-propagating beams. A primary loop locked the cavity length to one upshifted beam, and a secondary loop adjusted the frequency of the other beam’s AOM to lock it to the counter-propagating mode. Ultimately, it was found that even passive designs exhibit this “lock-in” effect [68, 69], which was determined to be the result of back-scattering from one beam to the other.

In this work, we report on a variant of the passive design by Ezekiel and Balsamo. The salient departure from that design is the operation of the two counter-propagating beams on adjacent axial modes of the ring cavity, such that the two fields are separated in frequency by one cavity free spectral range (100 MHz, in our case). This macroscopic frequency separation reduces the intracavity crosstalk, allowing for enhanced high-sensitivity, linear operation of the gyroscope down to zero frequency.

4.1.2 Current laser gyroscope sensitivities

Theoretically, due to the dimensional dependence in Eq. 4.1, larger-area gyroscopes are inherently more sensitive than smaller ones. In applications where size is not a great concern, such as in geophysical experiments, the paradigm has been to make instruments as large as is practical. Several large geophysical gyroscopes [70–75] have in recent years demonstrated a resolution of one part in 10^8 of the earth’s rotation rate over several-hour integration times. The largest of these gyroscopes is “UG-2” [75], a rectangle with $39.7 \text{ m} \times 21 \text{ m}$ sides.

On the other end of the spectrum, more compact designs have been used in aerospace for decades as important components of inertial guidance systems. Currently available models [76, 77], typically 10–20 cm on a side, exhibit best noise levels of $\sim 10^{-6} - 10^{-5} \text{ (rad/s)/}\sqrt{\text{Hz}}$ and DC stability of 0.001–0.01 °/hr.

³Upon the later construction of large enough units, it was found that the DC Sagnac shift afforded by the earth’s rotation was enough to prevent the “lock-in” effect. A calculation for the required gyroscope size in order to avoid lock-in, as a function of backscatter coefficient, can be found in [66].

In recent years, the use of laser gyroscopes has been investigated as a potential supplement to the active feedforward seismic isolation systems of the second-generation interferometric gravitational wave (GW) detectors Advanced LIGO [59, 78, 79] and Advanced VIRGO [80, 81]. In those systems, a network of seismometers is used to sense ground motion around the GW interferometers’ test masses, and their signals are used—via Wiener-filter based feedforward noise cancellation—to subtract the ground motion with force actuators. It is well known that seismometers exhibit a parasitic sensitivity to ground tilt at low frequencies [59], and the concept is therefore to use rotation sensors in parallel to remove the spurious tilt-induced component of the ground motion signal. In this application, cost and space constraints and the desire for localized tilt information near multiple interferometer components dictate that the size of the used gyroscope be on the order of a meter.

In the case of Advanced VIRGO, a prototype laser gyroscope [66, 82], “G-Pisa,” has been constructed. Using the conventional active design with a modest size of 1.4 m on a side, the G-Pisa sensor has been operating for some time at a sensitivity level of $\sim 10^{-8} - 10^{-9}$ (rad/s)/ $\sqrt{\text{Hz}}$. Based on a theoretical model for the active system, the Pisa group shows that this current noise floor is dominated by backscatter effects.

In this work, we will give a detailed description of a prototype passive, free-space laser gyroscope of 75-cm side length constructed with the aim of serving as a tilt sensor in the Advanced LIGO seismic isolation scheme. While not meeting these stringent requirements, we achieved a sensitivity of 10^{-8} (rad/s)/ $\sqrt{\text{Hz}}$ above 500 mHz.

4.1.3 Rotation sensing

In addition to the active and passive, free-space laser gyros described above, a number of other rotation sensor technologies exist. In Figure 4.1, we have compiled the angular sensitivity of a number of these in order to place the requirements for the gravitational-wave detectors in the proper context.

One of the earliest efforts to subtract tilt from the suspension points of a pendulum-based interferometer was proposed by Robertson, et al. [83] and utilized a rotational reference arm [84]. Recently a few groups have demonstrated low noise tilt sensing using balance beam sensors [85, 86]. These sensors are quite close to the sensitivity needed to reduce tilt at the upper stages of the LIGO seismic isolation system, but do not address the issue of tilts generated within the isolation system. For that, one would need to place a rotation sensor at the pendulum’s suspension point or to use

a tilt-free sensor [87] for the inertial isolation.

The fiber-optic ring gyro (FOG) [88] has seen rapid development in the last few decades and is much more sensitive than the MEMS gyros used in mobile devices, but is not yet competitive with the best active laser gyros. Further improvements in reducing coherent backscatter and polarization modulation may yield an order of magnitude improvement, but this would still fall far short of modest free space laser gyros as well as the tilt sensing requirements of gravitational wave detectors.

Atom interferometers [89–91] have also been used recently to make sensitive rotation measurements. In these systems, a laser is first used to prepare a beam of atoms into a known state, and then subsequently to interrogate the atomic state after some predetermined period of time. The rotation of the system can then be inferred through its well-understood coupling to the probability amplitudes of the output states. Another novel technique involves measuring the quantized flow of superfluid helium [92, 93].

The sensitivity and noise of the various types of rotation sensors are compared in Fig. 4.1.

4.2 The passive, free-space laser gyroscope

4.2.1 Overview

A schematic diagram of the experiment can be found in Fig. 4.2. Light from a commercial, 1064-nm Nd:YAG non-planar ring oscillator (NPRO) is locked to the counterclockwise mode of a square optical cavity via feedback to the laser frequency using the Pound-Drever-Hall (PDH) frontal-modulation heterodyne locking technique [101]. A pickoff of the input laser is upshifted using an acousto-optic modulator (AOM) by 100 MHz, the free spectral range (FSR) of the cavity. By using a separate PDH loop to feed back to the AOM frequency, this upshifted beam is locked to the clockwise mode of the same ring cavity. In this configuration, using Eq. 4.1, the frequency difference between these two beams is

$$\nu_{\text{cw}} - \nu_{\text{ccw}} = \frac{c}{S} + \frac{4}{\lambda S} \vec{A} \cdot \vec{\omega}, \quad (4.2)$$

where c is the speed of light and the area \vec{A} is defined as vertically oriented.

The rotation signal can be read out in either of two ways. Most simply, the signal is encoded directly into the AOM actuation signal. Alternatively, the frequency difference can be measured by recombining the transmitted beams and measuring




Figure 4.1: Shown here are the performance for a number of different rotation sensors: the MIT passive ring resonator [94], the G-Ring laser gyro [95, 96], tactical grade gyroscopes [97], the Laser FOG [98], the U. Wash. balance beam [85], the Caltech balance beam [86], the G-Pisa laser gyro [99], the Caltech optical gyro (this work), navigational grade gyros [100], the Stanford atom interferometer [91], and the UC-Berkeley superfluid He sensor [92].

the beat in the detected photocurrent using a phase-lock loop (PLL). In either case, the output signal is directly sensitive to phase noise in one of two RF voltage-controlled oscillators (VCOs): in the former case, it is VCO1 driving the AOM; in the latter, it is the VCO2 in the PLL. In each case, the small rotation-induced fluctuations are impressed on a 100 MHz carrier. Using Eq. 4.1, the target sensitivity of $\sim 10^{-9} \text{ (rad/s)/}\sqrt{\text{Hz}}$ corresponds to an optical frequency shift of approximately $1 \text{ mHz}/\sqrt{\text{Hz}}$, giving a required relative stability of $10^{-11} \text{ 1}/\sqrt{\text{Hz}}$ in the desired operational frequency band (typically between 10 mHz and 1 Hz for terrestrial GW detectors). Given the stability of currently available RF VCOs, this requirement puts a considerable constraint on the near-term improvement of the design we consider here⁴.

⁴It should be noted that in the optical beat and PLL readout case, one is free to use a wider

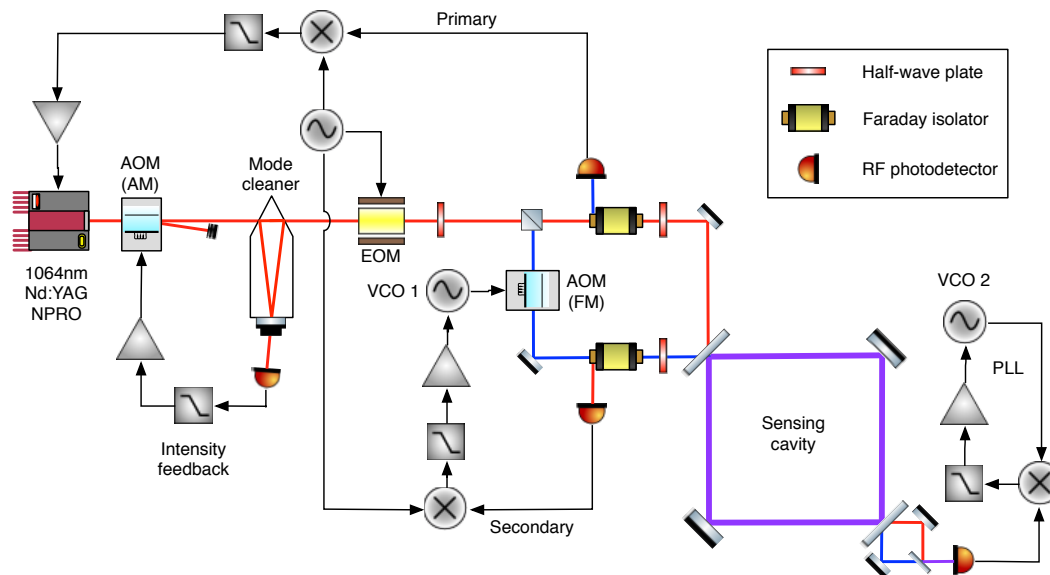


Figure 4.2: Simplified diagram of the laser gyroscope. After mode cleaning and intensity stabilization, the main laser is locked to the counterclockwise mode of the sensing cavity. A pickoff of the beam is upshifted macroscopically by 100 MHz—roughly the FSR of the cavity—and feedback is applied to the AOM to lock this upshifted beam to the clockwise mode of the cavity. The rotation signal is encoded in both the control signal to the AOM as well as the beat between the main and secondary beams in transmission of the cavity.

4.2.2 Detailed experimental design

4.2.2.1 Optical cavity

The core of the gyroscope is a square, free-space optical cavity of 75-cm side length. The cavity is critically coupled, with the input- and output-coupling mirrors each being 2-in diameter, flat, high-reflectivity (HR) optics positioned at opposing corners of the square. The cavity is geometrically stabilized using a single 3-m radius of curvature, 1-in HR mirror in another corner. The final optic is a flat, HR turning mirror. The coupling mirrors' power transmission of 200 ppm, along with the other mirrors' transmission and aggregate scatter and absorption losses, result in a finesse of approximately 12000.

The cavity is enclosed in a custom vacuum system composed of steel corner chambers connected by flexible metal braid tubes using KF flanges. The optical signals are injected and extracted from the vacuum envelope through optical-quality wedged, anti-reflection (AR) coated windows. The purpose of the vacuum system is to remove optical path fluctuations induced by index of refraction variations of the variety of frequency discriminators, some of which may in principle be inherently more stable.




Figure 4.3: Photo of the experiment.

air. As such, high vacuum is not required for low-noise operation; in practice, the chamber is evacuated to the mTorr level, sealed, and the pump removed to avoid excess vibration.

4.2.2.2 Beam preparation and cavity injection/extraction

At the output of the laser, the beam is first passed through a triangular cavity with a roundtrip length of 20 cm. The purpose of this mode cleaner cavity [102] is to suppress spatial jitter in the beam, as this can lead to errors in the frequency locking of the beams to the gyroscope cavity, coupling directly into the rotation signal. The cavity's length is adjustable with a PZT-mounted end mirror. To lock the cavity length to the input beam, a 1 MHz dither is applied to the PZT, and the photocurrent from the small leakage beam through the end mirror is mixed with the dither drive to provide a linear error signal. In essence, the mode cleaner converts beam jitter at its input into power fluctuations at the output. Completing the circle, the power fluctuations at the output of the mode cleaner are read out and fed back to an acousto-optic modulator and thereby suppressed. The length control loop for the mode cleaner has a bandwidth of a few kHz, which is sufficient to keep the RMS error well below its cavity linewidth.

After the mode cleaner, the beam passes through an electro-optic modulator (EOM) crystal that is resonant at 29.489 MHz, which imparts phase modulation sidebands at this frequency for the PDH locking scheme. To minimize the effects of residual

amplitude modulation (RAM) [103–105] from the crystal, the EOM is passively thermally isolated with insulating foam, and its metal housing is actively temperature stabilized above room temperature.

Following the EOM, the beam is split. The primary (CCW) beam is sent directly towards the cavity, while the secondary (CW) beam is upshifted by 100 MHz using an AOM. To avoid strong beam jitter due to the first-order dependence of the Bragg scattering angle on the modulation frequency (which is varied minutely by the feedback), a double-pass scheme is used: on a first pass, the beam is upshifted by 50 MHz; then, the beam is retro-reflected using a spherical mirror and passes through the AOM once more, acquiring a total roundtrip shift of 100 MHz before being directed towards the cavity.

The geometry of the cavity dictates that the one beam's reflected path is co-spatial with the input path of the other. In order to separate the respective reflected beams, a polarization-isolation scheme is used: each beam passes through a Faraday isolator (FI) on its way to the cavity, and then half-wave plates are used to ensure that each reflected beam is rejected by the conjugate FI on return. Here, the two reflected fields are acquired on custom-built, low-noise RF photodetectors (RFPDs) to generate the locking error signals (these RFPDs are described in more detail below).

The photodetectors used for PDH locking of the cavity, as well as for the transmitted-beam beat readout, are custom-built versions of the design described by Grote [106]. This RFPD topology is well suited for low-noise detection of narrowband RF signals in the presence of unwanted harmonics and with high DC power levels. In addition—unlike conventional resonant designs—it is not susceptible to photodiode bias modulation in the presence of large signals, which is particularly important in the case of the transmitted-beam readout, since the heterodyne beat signal is not suppressed.

The primary loop, which locks the CCW beam to the cavity via feedback to the laser, can be considered to control the “common mode” of the system since actuation on the laser adjusts the frequency of both beams. As discussed in the noise analysis section below, this common degree of freedom must be controlled with high accuracy to prevent pollution of the differential mode that contains the rotation signal. Therefore, the primary loop is designed to be a high-performance system. After mixing and low-pass filtering using the PDH method, the derived error signal is passed through a custom-built multi-stage servo filter that allows for the very high low-frequency gains necessary to sufficiently suppress the common-mode noise. The output of the

servo filter drives the NPRO PZT actuator to adjust the laser frequency. The primary loop achieves a unity-gain frequency (UGF) of ~ 15 kHz, limited by resonances in the NPRO PZT, with a resulting low-frequency loop gain of $> 10^{12}$ below ~ 1 Hz.

The requirements on the secondary loop are not so stringent, since environmental noise is all but eliminated by the primary loop. In reality, the rotation signal can be read out from the secondary error signal while the loop is open. This mode of operation was tested, and in most cases a similar performance was achieved. However, this can lead to non-linearities at large signal levels and in general requires a more complicated calibration procedure. To lock this loop, the CW error signal is fed to another, lower-performance servo filter, whose output drives the external modulation input of the VCO driving the AOM. The actuation gain of the VCO is controllable using the external modulation input deviation range. The relative frequency noise of the VCO is proportional to this range, and therefore the minimal acceptable range is chosen based on the observed signal level. The secondary loop is chosen to have a UGF of ~ 20 kHz, and it provides loop gains of $> 10^8$ below ~ 1 Hz.

4.2.3 Realized performance and noise analysis




Figure 4.4: Noise budget of the laser gyroscope. The measured sensitivity is plotted alongside the individual noise contributions. Mechanical noise is predicted to be the largest contributor, but the gyroscope exhibits excess noise above the expected level.

A noise budget for the passive laser gyroscope is in Fig. 4.4, showing the realized

sensitivity of 10^{-8} (rad/s)/ $\sqrt{\text{Hz}}$ above 500 mHz and increasing as $1/f$ at lower frequencies. We acknowledge that this realized noise floor is higher than the incoherent sum of all known noise contributions. While we cannot fully account for this, we offer the following two observations as possible explanations.

First, the low-frequency excess was observed to be non-stationary on several-day time scales. A careful realignment of the injected and reflected beams was found to minimize the noise in all cases. Therefore, while we were unable to fully mitigate the effect, it is quite possible that alignment issues caused this excess noise (through, e.g., diffractive losses from apertures and/or scattered light interference, etc.). It is for this reason that, in Sec. 4.3, we propose a possible future design that is less susceptible to alignment effects.

Second, we cannot rule out the possibility that the excess “noise” was caused by a true rotation signal. Detailed seismic data for the GW detector sites that set our sensitivity goal suggest this level is too high to be accounted for by ambient seismic activity *at these locations*. However, such data is not available for our laboratory, and it is feasible that our setting is more seismically active. In an attempt to confirm or deny this effect, we performed correlation studies with seismometers placed on the optical table. The seismometers’ signals were combined to isolate the degree of freedom to which the gyroscope is sensitive, and offline coherent subtraction was attempted. While strong coherence and subtraction could be seen when the rotation noise was high (either due to abnormally high environmental rotation noise, or as a result of an intentional noise injection using a mechanical shaker), we found that the internal noise of the seismometers was too high for the scheme to be effective when the gyroscope exhibited its lowest noise levels⁵. Therefore, we were unable to exclude this effect as a potential limiting contributor.

The remainder of this section offers a detailed analysis of the various well-understood noise contributions. When analyzing the performance of the instrument, it is helpful to first consider the ideal scenario. In that case, the primary loop has infinite gain at all frequencies, making the laser frequency perfectly follow the CCW mode of the cavity. Then, the secondary loop must only correct for differences between the CW and CCW modes. Were it not for the one-axial-mode shift between the beams’ frequencies, any length fluctuations in the cavity would be completely common

⁵This is somewhat expected, since seismometer noise and degree-of-freedom cross-coupling are the reason that a rotation sensor is desired in the first place.

mode, and would therefore not couple to the rotation signal to first order⁶. The fact that these two modes are separated in frequency leads to a residual length noise contribution with a common-mode rejection of $\nu_{\text{FSR}}/\nu \approx 3.5 \times 10^{-7}$ (see Sec. 4.2.3.1 below).

Returning to the realistic limit of finite primary-loop gain, another source of noise emerges: since the primary loop cannot perfectly follow the CCW mode, any common-mode residuals must be corrected for by the secondary loop. This residual noise appears directly in the rotation signal, and it is the reason for the stringent primary-loop gain requirement.

These are the only two noise sources somewhat unique to the gyroscope; as seen below, the remainder are more standard contributors.

4.2.3.1 Mechanical noise

As described briefly above, cavity length fluctuations driven by mechanical noise are largely common mode. However, the macroscopic frequency shift between the CW and CCW beams spoils the otherwise perfect cancellation. Allowing for some small fluctuation δS in the cavity perimeter, we find that the supported eigenfrequencies in both directions are

$$\nu_{\text{ccw}} = n \frac{c}{S + \delta S} \quad (4.3)$$

$$\nu_{\text{cw}} = (n + 1) \frac{c}{S + \delta S}, \quad (4.4)$$

where n is an integer. From the raw beat signal, we will subtract the known offset of $c/S \approx 100$ MHz, giving a fluctuating frequency signal

$$\Delta \nu_{\text{sig}} = -\frac{c \delta S}{S^2 + S \delta S} \approx -c \frac{\delta S}{S^2}, \quad (4.5)$$

where the approximation $\delta S \ll S$ has been made. The linear term above is the first-order modulation of the FSR offset signal due to the cavity length fluctuation. Comparing it with the common-mode frequency shift from the same length fluctuation (i.e., $\Delta \nu_{\text{sig}}^0 = -\nu \delta S/S$), we obtain the common-mode rejection ratio

$$\text{CMRR} \equiv \frac{\Delta \nu_{\text{sig}}}{\Delta \nu_{\text{sig}}^0} \approx \frac{\nu_{\text{FSR}}}{\nu} = \frac{\lambda}{S} \approx 3.5 \times 10^{-7}. \quad (4.6)$$

⁶Since the scale factor (Eq. 4.1) is dependent on the cavity length, there would still be a small, higher-order coupling.

This noise can be measured using auxiliary channels and subtracted from the rotation signal, either online or in post-processing. To estimate the contribution, the actuation signal for the primary loop is monitored. Where that signal is dominated by the length fluctuations in the cavity (as opposed to, e.g., laser frequency noise), it is a faithful monitor of the common-mode noise. This signal is then multiplied by the CMRR, converted to rotation noise via Eq. 4.1, and subtracted from the rotation signal.

In practice, below ~ 1 Hz, the free-running laser frequency noise is larger than the frequency equivalent of the cavity motion. Since laser frequency noise *is* completely common mode, this can lead to an overestimation of the mechanical noise, and therefore to pollution of the rotation signal upon subtraction. To combat this, we have installed a laser frequency monitor by beating a pickoff of the gyroscope laser output with light that has been stabilized by locking to a quiet reference cavity. When the gyroscope is locked, this in-loop signal contains information only about the cavity motion, which allows for faithful subtraction.

4.2.3.2 Residual common-mode noise

The secondary loop simply acts to adjust its beam's frequency to match that mode's eigenfrequency. Therefore, any residual failure of the primary loop to lock the laser frequency to *its* mode results in an injection of common-mode noise into the secondary loop (and hence directly into the rotation signal). Mathematically, the open-loop frequency error signal seen by the secondary loop is

$$v_{\text{err}} = \left(\frac{1}{1 + G_p} + \text{CMRR} \right) \nu \frac{\delta S}{S} + \frac{4}{\lambda S} \vec{A} \cdot \vec{\omega}. \quad (4.7)$$

The first term in the parentheses—where G_p is the primary loop gain—is the residual frequency error that remains after the action of the primary loop, while the second term is the differential mechanical noise described in the previous section. The last term is the Sagnac rotation signal.

In principle, this noise could be subtracted as well, but it is more effective to simply increase the primary loop gain until it is suppressed below the desired level. Since the target operational band is in the 10 mHz - 1 Hz range, it is easy to shape the primary servo filter to have an acceptable level of gain. For budgeting purposes, this noise contribution is calculated by measuring the in-loop primary error signal and referring it first to optical frequency (by dividing by the primary-loop optical gain in [V/Hz]) and then to rotation noise using Eq. 4.1.

4.2.3.3 Oscillator frequency noise

As described in Sec. 4.2.1, the frequency stability of the RF VCOs used in the experiment is an important factor in the ultimate sensitivity of the gyroscope⁷. Both oscillators in this experiment are IFR/Marconi 2023A [107] RF sources, operated in external-input frequency modulation mode. For improved low-frequency stability, both sources are locked to a Stanford Research Systems FS725 [108] rubidium frequency standard via a 10-MHz reference signal. In the external-FM mode, the 2023A's frequency noise is proportional to the FM actuation range.

The VCO noise enters the rotation signal in different ways depending on the mode of operation. In the AOM actuation readout mode (see Sec. 4.2.1), the rotation signal is taken at the external FM input of VCO1. The secondary loop acts to cancel this noise at the optical error point, and therefore this signal contains directly the frequency noise of the VCO.

In the transmission beat note readout, on the other hand, the CW optical signal contains the much smaller, loop-suppressed contribution of the noise from VCO1. However, the noise from VCO2 is imposed directly on the PLL control signal readout in a similar fashion to above.

From one case to another, the required FM deviation range and the carrier frequency are only different by a factor of 2 (VCO1 is at $f_{c1} = 50$ MHz, due to the AOM double-pass setup, while VCO2 is at the full $f_{c2} = 100$ MHz), and so the ultimate contribution to the rotation noise is roughly the same.

Since VCO1 is used as an actuator, there is little room to reduce the oscillator noise contribution in the first case. In the second case, however, VCO2 is only used as an actuator due to the PLL topology. Here, all that is necessary is a low-noise frequency discriminator, and so one is free to choose an alternate design with lower internal frequency noise. Due to the demanding relative frequency noise requirement of $10^{-11}/\sqrt{\text{Hz}}$ at and below 1 Hz, several candidates (e.g., delay-line mixer frequency discriminator, LC detector, etc.) seem impractical [109]. Finding a suitably stable frequency discriminator is an important step to further improving the sensitivity of our design⁸.

⁷The stability of the fixed RF oscillator used to provide the PDH sidebands is comparatively unimportant, as the common-frequency, balanced-phase modulation/demodulation scheme gives first-order insensitivity to this oscillator noise.

⁸While the 2023A is a top-end production VCO, certain specialty and custom-made oscillators may exhibit better stability. For a future gyroscope reaching the oscillator phase noise limit, a thorough cost/benefit analysis of all available units may be beneficial.

4.2.3.4 Electronic noise

The electronic noise of each loop is measured independently and calibrated to units of rotation signal as appropriate. In order to measure this noise, the laser is blocked and the actuation signals from both loops are measured in this “dark” state. Both measured noise spectra are divided by their respective servo filter and optical gains to refer them to their inputs, and these form the effective sensing noise levels.

For the primary loop, this sensing noise sets a limit to the achievable common-mode noise suppression. In effect, it adds a fixed term $\delta\nu_1^{\text{sens}}$ to Eq. 4.7, such that increasing G_p beyond a threshold level of

$$G_p^{\text{SNL}} \equiv \frac{\nu \delta S}{S \delta\nu_1^{\text{sens}}} - 1 \quad (4.8)$$

no longer reduces the contribution of the primary loop noise to the rotation signal. If $G_p > G_p^{\text{SNL}}$, the primary loop is *sensing noise limited*. This threshold depends on the environmental noise present (through δS), but in practice G_p is high enough to meet this criterion at all times in the frequency range of interest.

Since the secondary loop reads the rotation signal out directly, sensing noise there constitutes an absolute limit for the rotational sensitivity in the most basic way, using Eq. 4.1.

Using ultra-low-noise front-end electronics and preferentially distributing gain upstream in the signal chain to reduce the effect of noisier later stages, the total contribution of electronic noise can be made negligible with respect to other sources. This is one motivation for using a high-finesse optical cavity, since a given voltage fluctuation in the front-end electronics corresponds to a smaller optical frequency fluctuation than with a lower-finesse cavity.

4.2.3.5 Residual amplitude modulation (RAM)

As with any FM spectroscopy scheme, the PDH technique is susceptible to error from *residual amplitude modulation (RAM)* in the EOM used for control sideband generation. The EOM operates by applying an electric field across a crystal that exhibits the Pockels effect (i.e., it has a birefringence linearly proportional to the electric field applied). If the input beam polarization and electrode axis are exactly aligned with the appropriate crystal axis, the result is a pure phase modulation (PM) of the output beam. However, any slight misalignment can result in oscillatory rotation of the output beam polarization at the same frequency as the phase modulation.

Upon interaction with any polarization-sensitive optics (e.g., polarizing beamsplitters, birefringent mirrors, etc.), this polarization rotation is converted directly into amplitude modulation. As a result, the beams incident on the optical cavity have both PM and AM at the same frequency.

In the PDH scheme, a slight frequency offset from resonance converts PM into AM, which can be detected in the reflected beam by a photodetector, and the phase of this AM signal encodes the sign information about the offset. Therefore, any AM present in the input beam leads directly to an unwanted offset in the locking loop (i.e., the loop acts to cancel the RAM-induced offset by creating equal and opposite AM via frequency offset from the cavity eigenmode).

A fixed AM level generates a static offset, which can be corrected for in principle by adding an electronic offset at the front end. A time-varying offset, however, is indistinguishable from a true signal, and in this way RAM directly produces noise in the rotation signal.

The dominant driving source for low-frequency RAM fluctuations is temperature noise. For this reason, the EOM is temperature stabilized both passively and actively. For passive isolation, the EOM is covered by shape-fitting thermally insulating foam, which suppresses temperature-driven noise above the resultant thermal pole around 100 mHz. To suppress noise below the thermal pole, the temperature of the EOM enclosure is actively stabilized using a custom controller.

This active/passive stabilization scheme is observed to suppress the RAM contribution considerably in broadband. However, even with this frontal suppression, the low-frequency rotation noise was often discovered to be coherent with RAM in at least one of the two beams. To further suppress this noise, an out-of-loop RAM witness photodetector was placed at a pickoff of each input path near the cavity injection (see Fig. 4.5). Since the beam sampled here has not yet interacted with the cavity, any AM present at the modulation frequency is RAM-induced. Each monitor photocurrent is mixed down with the PDH local oscillator signal—with the appropriate phase shift—in order to obtain the spurious RAM-induced component of the signal.

To choose the appropriate demodulation phase, the input polarization to the EOM is misaligned to introduce intentionally large RAM. Then, the demodulation phase is adjusted to maximize the DC output of the RAM monitors. Once this phase is set, the polarization is realigned to zero the RAM monitor signals, which are then

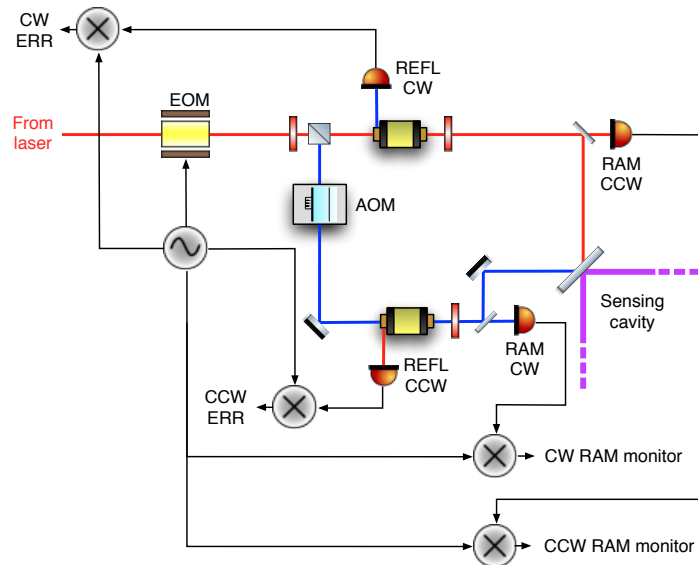


Figure 4.5: Diagram of the out-of-loop RAM monitor setup.

amplified and digitally acquired for subsequent post-processing.

With the RAM monitor signals recorded along with the rotation signal, etc., Wiener filtering can subsequently be used to subtract any components of the rotation noise coherent with the RAM monitor signals.

4.3 Possible future designs




Figure 4.6: Proposed fiber-distributed gyroscope array scheme.

Given the information gleaned from this initial experiment, we have considered one possible future design. In doing so, we have paid particular attention to robustness,

cost, and ability to be scaled up for use in demanding scientific environments.

In the foregoing analysis, it is clear that mechanical disturbances are a major source of noise. One obvious improvement, therefore, is to increase the mechanical stability of the ring cavity. This could be done by using a monolithic construction, where the cavity is of a single piece of a suitable low-expansion material. A superior solution involves changing the locking topology: by adding a length actuator to one cavity mirror, one can lock the cavity length to the laser, rather than the other way around. If the laser source is externally stabilized, it becomes a quiet reference against which the cavity length fluctuations are measured and strongly suppressed via feedback.

As an added advantage, using a quiet, fixed-frequency laser allows for multiple gyroscopes to be illuminated by the same source. By taking advantage of this fact, one could create a relatively inexpensive array of gyroscopes. Each individual unit would consist of its own cavity, two photodetectors, and an acousto-optic modulator to internally generate its secondary beam. The central laser would be distributed to each gyroscope using a robust system of optical fibers, and each unit would be pre-aligned internally, allowing for easy removal and reconnection of the light source for unit relocation.

Since this scheme would strongly suppress the mechanical noise that dominates in our current system, the ring cavity could also be reduced in size by a factor of two or more while maintaining improved sensitivity. As per Eq. 4.1, the rotation sensitivity for a given readout frequency noise level scales linearly with the cavity length. Therefore, for example, the area may be scaled down by a factor of four, and the resulting projected noise floor would be only a factor of two higher than the magenta trace in Fig. 4.4 (since this is the next dominant noise term after mechanical noise). In addition to making the design more compact, this may also facilitate making the monolithic cavity structure highly mechanically stable.

4.4 Conclusion

We have described a novel laser gyroscope design employing a passive, free-space optical cavity. By separating the counter-propagating sensing optical fields macroscopically in frequency, this design is free from the lock-in effects observed in previous systems (active and passive). We have demonstrated a rotation sensitivity of 10^{-7} radians/ $\sqrt{\text{Hz}}$ below 1 Hz. Our instrument noise floor is compared with those of other contemporary rotation sensors in Fig. 4.1. A convenient feature of this design is the ability to use a commercial laser to illuminate the system, rather

than building a custom laser resonator.

We believe our design lends itself to integration in complex systems requiring accurate rotation sensing. In particular, the sensing cavity need not be an ad-hoc mirror-and-mount structure; we envision that a future version could benefit from a monolithic design, which could make the system more compact as well as more immune to environmental disturbances.

Part III

**Advanced LIGO Output Mode
Cleaner**

Chapter 5

BACKGROUND

5.1 DC Readout

In Initial LIGO [18], all length degrees of freedom were sensed using the RF heterodyne Pound-Drever-Hall [38] described in Sec. 2.1.2.4 and Appendix B. Beginning in Enhanced LIGO [22] and continuing into Advanced LIGO [23], sensing of the DARM degree of freedom—containing the GW signal—was changed into a special case of homodyne detection known as **DC readout**.

The reasons for this are manifold, and a thorough treatment can be found in [110, 111]. However, the main reasons are highlighted below:

- **RF oscillator noise coupling:** As described in Sec. 2.1.1.6, laser intensity and phase noise around the carrier frequency are suppressed above the coupled cavity pole before reaching the dark port. As the control sidebands used in the PDH readout scheme are not resonant in the arms, they are not likewise suppressed. This leads to a strong noise coupling from the oscillators used to generate these sidebands to the GW readout.
- **Junk light:** Since the spatial distribution of the resonant modes in the arm cavities are in general slightly different from each other (and from the input laser) due to aberrations, some amount of light—called **junk light**—does not resonate and/or destructively interfere upon recombination at the BS. This provides a short-circuit from laser noise at the input of the interferometer to the GW readout. Junk light also adds additional shot noise to the detection process.
- **Mode overlap:** In the RF readout, the error signal is generated by the beating of the carrier audio sidebands and the RF control sidebands. The former fields are resonant in the arms, whereas the latter are not. Therefore, these modes were not always well overlapped in general, leading to a reduction in the optical readout gain, in addition to adding shot noise as with junk light.

DC readout addresses these issues by using a mode cleaning cavity at the output—the **output mode cleaner (OMC)**—to filter the light leaving the interferometer

before it is detected, both spectrally (to remove the control sidebands) and spatially (to remove the junk light).

In order to sense the motion of the DARM DoF, a small offset is introduced between the arms¹, giving a linear sensitivity of the output power to DARM. This can be thought of as a form of homodyne detection: instead of the traditional method, where a phase-controlled pickoff of the laser—the “local oscillator (LO)” —is sent around the interferometer and mixed with the output light on a beamsplitter, in this case, the small amount of DC power emerging from the dark port itself serves as the LO.

5.2 Mode cleaners in detail

In Sec. 2.1.1.8, we briefly discussed how the selective transmission of an optical cavity can be used as a sort of filter, allowing light with certain characteristics to pass, while rejecting light with others. Until now, we have primarily been considering a mode cleaner’s *spectral* response: since a cavity is only resonant—and therefore only transmissive—at or near a certain set of supported frequencies, we can use mode cleaners to filter out unwanted frequency components of an optical beam. In this section, we will focus on the *spatial* filtration capabilities of mode cleaners.

5.2.1 Transverse modes and Gaussian beams

For simplicity, we often discuss the laser phase front as if it were a plane wave. This is because, for the purposes of longitudinal propagation calculations, we typically only consider the on-axis behavior of the light. Because any real optical beam must have a finite transverse extent (i.e., since it cannot extend to infinity as a plane wave does), it must also spread as it travels due to **diffraction**. This effect determines both the shape of a beam’s phase surfaces and also its intensity profile as a function of transverse distance.

A laser must be generated by a laser cavity. Two important criteria for laser operation are 1) that the cavity has a gain above unity in one roundtrip, so that there is net amplification, and 2) that the cavity is geometrically stable (i.e., that there exists some propagation mode or set of modes such that an optical phase front returns to its original shape after one roundtrip around the cavity). One simple and ubiquitous geometry is a combination of spherical mirrors, which produces a very special mode

¹This is the $\delta\phi$ introduced in (2.45).

known as a **Gaussian beam**. Such a beam traveling in the $+\hat{z}$ direction has a field²

$$E(r, z) = E_0 \frac{w_0}{w(z)} e^{-\frac{r^2}{w^2(z)}} e^{-i\left(kz + k\frac{r^2}{2R(z)} - \psi(z)\right)}, \quad (5.1)$$

where r is the radial distance from the central beam propagation axis, z is the distance along the propagation axis, and $k = \frac{2\pi}{\lambda}$ is the wave number. The beam **waist size**, w_0 , is an important quantity that describes the radius at which the electric field falls off by $1/e$ at its narrowest point.

The waist size essentially specifies the shape of the Gaussian beam, which is shown in Fig. 5.1. From it, we can calculate the **Rayleigh range**,

$$z_R \equiv \frac{\pi w_0^2}{\lambda}, \quad (5.2)$$

which describes the length of the so-called **near field** region—straddling the beam waist—where the beam reaches its most focused. Also, the **divergence angle** that outlines the conical shape the beam approximates at high z is given by

$$\theta = \frac{\lambda}{\pi w_0}. \quad (5.3)$$




Figure 5.1: Diagram of a Gaussian beam.

Using w_0 and z_R , we can calculate the remaining unknown elements in (5.1). For example, the beam width as a function of distance is

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}. \quad (5.4)$$

Here we see that, in the **far field** where $z \gg z_R$, the beam size increases linearly with the propagation distance.

²Here and many other places in this thesis we leave out the polarization of the electric field, which can be linear or elliptical in the plane perpendicular to the direction of propagation.

The beam phase front is spherical in shape³, with a radius of curvature

$$R(z) = z \left[1 + \left(\frac{z_R}{z} \right)^2 \right]. \quad (5.5)$$

Note that in the deep far field, $R(z) \approx z$. On the other hand, $R(z) \rightarrow \infty$ as $z \rightarrow 0$. This transition of the phase front from spherical to planar and back to spherical with the opposite inflection as the beam passes through a focus leads to an overall phase *advance* known as the **Gouy phase**,

$$\psi(z) = \arctan \left(\frac{z}{z_R} \right). \quad (5.6)$$

The Gaussian mode is not the only one supported by a given spherical-mirror cavity. In fact, it is simply the lowest-order mode in a family known as the **Hermite-Gaussian** modes, which describe solutions with varying number of nodes in both transverse directions (i.e., perpendicular to the beam path). For a beam traveling in the \hat{z} direction (as above), a mode with m nodes in the \hat{x} direction and n in the \hat{y} direction is typically referred to as the **TEM_{mn}** mode⁴. This mode has a field

$$E_{mn}(x, y, z) = E_0 \frac{w_0}{w(z)} H_m \left(\frac{\sqrt{2}x}{w(z)} \right) H_n \left(\frac{\sqrt{2}y}{w(z)} \right) e^{-\frac{x^2+y^2}{w^2(z)}} e^{-i \left(kz + k \frac{x^2+y^2}{2R(z)} - (m+n+1)\psi(z) \right)}, \quad (5.7)$$

where $H_m(x)$ is the **Hermite polynomial** of order m , which follows the recursion relation

$$H_{m+1}(x) = 2xH_m(x) - 2mH_{m-1}(x). \quad (5.8)$$

The first few Hermite polynomials are

$$H_0 = 1 \quad (5.9)$$

$$H_1(x) = 2x \quad (5.10)$$

$$H_2(x) = 4x^2 - 2 \quad (5.11)$$

$$H_3(x) = 8x^3 - 12x. \quad (5.12)$$

As we can see, (5.7) reduces to (5.1) when $m, n = 0$ and we identify $r^2 = x^2 + y^2$.

In addition to describing the supported modes of a resonator cavity, the TEM_{mn} modes comprise a basis set for a generic optical beam in the paraxial approximation, with respect to a particular propagation direction, waist position and transverse axis

³This spherical shape matches that of the mirrors in the cavity used to generate the beam.

⁴TEM stands for “transverse electro-magnetic”.

orientation. In other words, any beam can be described as a sum of TEM_{mn} modes in any chosen basis that is paraxial to it:

$$E(x, y, z) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} c_{mn} E_{mn}(x, y, z). \quad (5.13)$$

For example, a beam that is slightly displaced in position or misaligned in angle from the z axis can be described as a sum of primarily the associated TEM_{00} mode and the next-lowest TEM_{10} and TEM_{01} modes. Conversely, a beam that is centered and aligned, but with a different waist size (and, therefore, a different radius of curvature at all points) is described as a sum of primarily the TEM_{00} mode along with the TEM_{20} and TEM_{02} modes.

5.2.2 Cavity transverse mode spacing and higher-order mode transmission spectrum

Note that the Gouy phase is preceded by the factor $(m+n+1)$ in (5.7). This indicates that the Hermite-Gaussian higher-order modes (HOMs) experience an extra phase advance when passing through a focus, in proportion to the mode order. This means that, in order to calculate the resonance frequencies for the higher order modes—relative to those for the TEM_{00} mode—we must include the effect of this additional Gouy phase accumulation in each roundtrip around the cavity⁵. Considering a cavity of length L , the difference in phase from the input side to the end side is

$$\phi(z_e) - \phi(z_i) = kL - (m+n+1) (\psi(z_e) - \psi(z_i)). \quad (5.14)$$

To simplify calculations for the case of a two-mirror cavity, we can define the cavity **g parameters** as

$$g_1 \equiv 1 - \frac{L}{R_1} \quad (5.15)$$

$$g_2 \equiv 1 - \frac{L}{R_2}, \quad (5.16)$$

where R_1 and R_2 refer to the radii of curvature of the input and end mirrors, respectively. Using these as shorthand, it can be shown [28] that the single-pass Gouy phase accumulation is

$$\psi(z_e) - \psi(z_i) = \arccos(\pm\sqrt{g_1 g_2}), \quad (5.17)$$

⁵Of course, the Gouy phase must also be considered even for the TEM_{00} mode, leading to a shift in the resonant frequencies we found in 2.1.1.3 for the simple plane wave. In practice, this slight shift is imperceptible in comparison to the large optical frequency; by contrast, the *difference* between the resonance frequencies of modes of differing order discussed here is a very significant effect.

where the sign of the argument is positive for $g_1, g_2 > 0$ and negative for $g_1, g_2 < 0$. The resonance condition therefore becomes

$$2 [kL - (m + n + 1) \arccos (\pm\sqrt{g_1 g_2})] = l \cdot 2\pi, \quad (5.18)$$

where l is an integer. Recalling that $k = \frac{2\pi\nu}{c}$, this leads to a set of supported mode frequencies:

$$\nu_{lmn} = \left(l + \frac{(m + n + 1) \arccos (\pm\sqrt{g_1 g_2})}{\pi} \right) \frac{c}{2L}. \quad (5.19)$$

In the above, note that the supported modes are defined in terms of both the *axial* mode number, l —i.e., the number of wave cycles that fit into the cavity on one round trip—in addition to the transverse mode numbers. The axial modes of a given transverse order are spaced one free spectral range ($\nu_{\text{FSR}} = \frac{c}{2L}$) apart, as expected. On the other hand, transverse modes of the same axial mode order but differing $m + n$ are spaced by the smaller interval

$$\nu_{\text{TMS}} = \frac{\arccos (\pm\sqrt{g_1 g_2})}{\pi} \nu_{\text{FSR}}, \quad (5.20)$$

which is known as the **transverse mode spacing (TMS)**.

The non-degeneracy of transverse modes of differing order leads to a particular HOM spectrum for a cavity with a given set of g parameters. An example of such a spectrum for a cavity with $g_1, g_2 = 0.98$ is shown in Fig. 5.2. As described in the caption, using this HOM spectrum along with the cavity transmission as a function of frequency, one can calculate how HOM content in a beam will be transmitted by the cavity.

Therefore, for a suitable choice of geometry and a high enough finesse, a cavity can be used to filter out light that is not in the desired spatial mode, even if it is at the desired frequency⁶. This effect can be used to create an output beam with very stable and clean spatial characteristics, while the input beam may be spatially fluctuating or contaminated, whether due to jitter-induced misalignment or to the presence of junk light from optical aberrations. Any such unwanted components in the input beam are represented as HOM content in the cavity basis, and are therefore rejected⁷.

⁶This complements a mode cleaning cavity's other job of filtering out light that *is* in the desired spatial mode, but *not* at the desired frequency.

⁷Note that this rejection of unwanted HOM content can lead in general to conversion of spatial fluctuations to *intensity noise* at the cavity output, since the power from the HOMs is not transmitted.




Figure 5.2: Higher-order mode spectrum for a cavity with $g_1, g_2 = 0.98$. The smooth curve is the transmission of the cavity assuming a finesse of 1000. This plot can be used to determine the transmissivity of the cavity to HOMs as follows. The assumption is that the input beam is a sum of the TEM_{00} mode plus some HOM content. If the cavity is locked to the TEM_{00} , how much of the HOM content is transmitted? For each HOM transmission line, there is actually a smooth transmission function as has been plotted for the fundamental mode. The fractional amount of the HOM transmitted is equal to value of this transmission function at zero frequency. Due to symmetry, this is equal to the value at the intersection of the given HOM transmission line with the single transmission function curve. For example, the $n = 3$ mode (i.e., TEM_{30} , TEM_{21} , TEM_{12} or TEM_{03}) is shifted from the fundamental mode resonance by ≈ 0.19 FSR, and this results in a fractional transmission of just below 10^{-5} .

To combat this, intensity stabilization loops are often used to lock the power of the output beam to a fixed level. Thus, a two-step process is used to create the final, clean and quiet beam: one to stabilize the spatial mode and another to correct the power fluctuations imposed by the first.

Chapter 6

THE ADVANCED LIGO OMC

6.1 Optical design

The purpose of the OMC is to prepare the output beam from the interferometer for photodetection by 1) removing the RF control sidebands, and 2) removing junk light outside the desired spatial mode (at any frequency). These two jobs can be treated under the same framework using the cavity basis formalism described in the previous section. Specifically, this amounts to designing the cavity such that, while the carrier TEM_{00} mode is resonant, there is minimal transmission of any HOMs *or* of the RF sideband fields in the TEM_{00} spatial mode.

The macroscopic design parameters of the OMC are chosen by balancing several factors, including [112]:

- **Suppression of undesired modes:** The cavity should be effective at suppressing the undesired elements of the optical field. This favors a higher-finesse, longer cavity such that linewidth is minimized.
- **Transmission of desired mode:** The cavity should be strongly transmissive to the desired fundamental carrier mode and its audio (signal) sidebands. This is easier to accomplish in the presence of unavoidable optical losses with a *lower-finesse* cavity.
- **Size:** The cavity should be small enough to be built on a monolithic bench for geometric stability. This imposes a maximum roundtrip length on the order of 1 meter.
- **Shape:** The cavity should be a ring cavity—rather than a linear cavity—for easy extraction of the reflected beam. Also, to simplify the HOM spectrum, the cavity should have an even number of mirrors¹.

Given the shape and length considerations, the OMC was chosen to be a “bowtie” cavity (as depicted earlier in Fig. 2.8), and the approximate radius of curvature for

¹Due to the odd number of reflections on a single trip around a cavity with an odd number of mirrors, HOMs with odd horizontal order m incur an extra effective phase shift of π on each circuit. This splitting of even and odd modes of the same overall order can complicate the HOM structure.

the non-flat mirrors was chosen to be 2.5 m, to place the cavity in a stable regime. The finesse was chosen to achieve a cavity transmission of $> 98\%$ for the carrier TEM_{00} . Given expected scatter loss of 40 ppm per mirror, along with a design leakage transmission of 50 ppm for the two non-coupling mirrors—i.e., a total loss of 140 ppm roundtrip—this specifies $\mathcal{F} \leq 450$. For margin, the target finesse was chosen to be $\mathcal{F} \approx 400$. A cavity with such a finesse and a roundtrip length on the order of a meter gives a linewidth near a few hundred kHz. The aLIGO RF modulation frequencies are $f_1 \approx 9.1$ MHz and $f_2 \approx 45.5$ MHz. Therefore, the rejection factor for the f_1 sidebands will be on the order of 1000, while that for the f_2 sidebands will be around 20000.




Figure 6.1: HOM spectrum for a cavity with the coarse parameters chosen for the OMC, taking into account the aLIGO control sideband fields at $f_1 \approx 9.1$ MHz and $f_2 \approx 45.5$ MHz.

The fine choice of cavity length and mirror radius of curvature must be made in consideration of the exact HOM structure. Following the same procedure as in Fig. 5.2, but for a cavity with the approximate parameters described here, and taking into account the upper and lower f_1 and f_2 sideband fields, we get the HOM

spectrum shown in Fig. 6.1².

Now, this chart only shows how much of a given field *present at the input of the OMC* is transmitted to its output; the actual amount of unwanted light added to the photodetection process—the real figure of merit—is this transmission factor times the power of that component incident on the OMC. Therefore, we must also have this information if we want to design the OMC to minimize the unwanted light transmitted along with the signal. To this end, mode scans of the optical field at the dark port of the eLIGO interferometers were used to create power-law models of the optical modes expected to be incident on the aLIGO OMC [113], as shown in Fig. 6.2.

Equipped with this power-law model, one can compute the HOM spectrum at a given choice of roundtrip length and mirror radius of curvature, then sum over all modes—for the carrier and the upper and lower f_1 and f_2 sidebands—to get the total unwanted light power in transmission. As either or both of these parameters are varied, this total power will change, depending on which modes are closest to resonance and how much of each mode is present at the input. The optimal choice of roundtrip length and mirror radius of curvature is that for which this total transmitted power is minimized. The results of this minimization procedure for the length and radius of curvature of the L1 OMC are shown in Figs. 6.3 and 6.4, respectively.

6.2 Mechanical design, fabrication and benchtop testing

6.2.1 Overall construction and components

The low tolerance of the OMC’s performance to deviations from the design optical parameters require it to be rigid and mechanically stable. Additionally, though the transmitted power—and therefore the inferred GW signal—has no linear coupling from fluctuations in the OMC cavity length on resonance, there can be quadratic and/or bilinear coupling if the resonance condition is not kept well enough or if there are strong enough fluctuations. For these reasons, the OMC is monolithic in construction, with discrete fused silica cavity mirrors fixed permanently to a large “breadboard,” also made of fused silica. A photo diagram of the L1 OMC is shown in Fig. 6.5.

²This calculation and most of those made in this chapter were done using software written by the author for rapid computation of cavity parameters from a simple user definition. This program was designed to accept nearly arbitrary cavity geometries, which it accomplishes by explicit propagation of the Gouy phase around the legs of a cavity, rather than converting to an equivalent two-mirror cavity and using the associated g parameters.

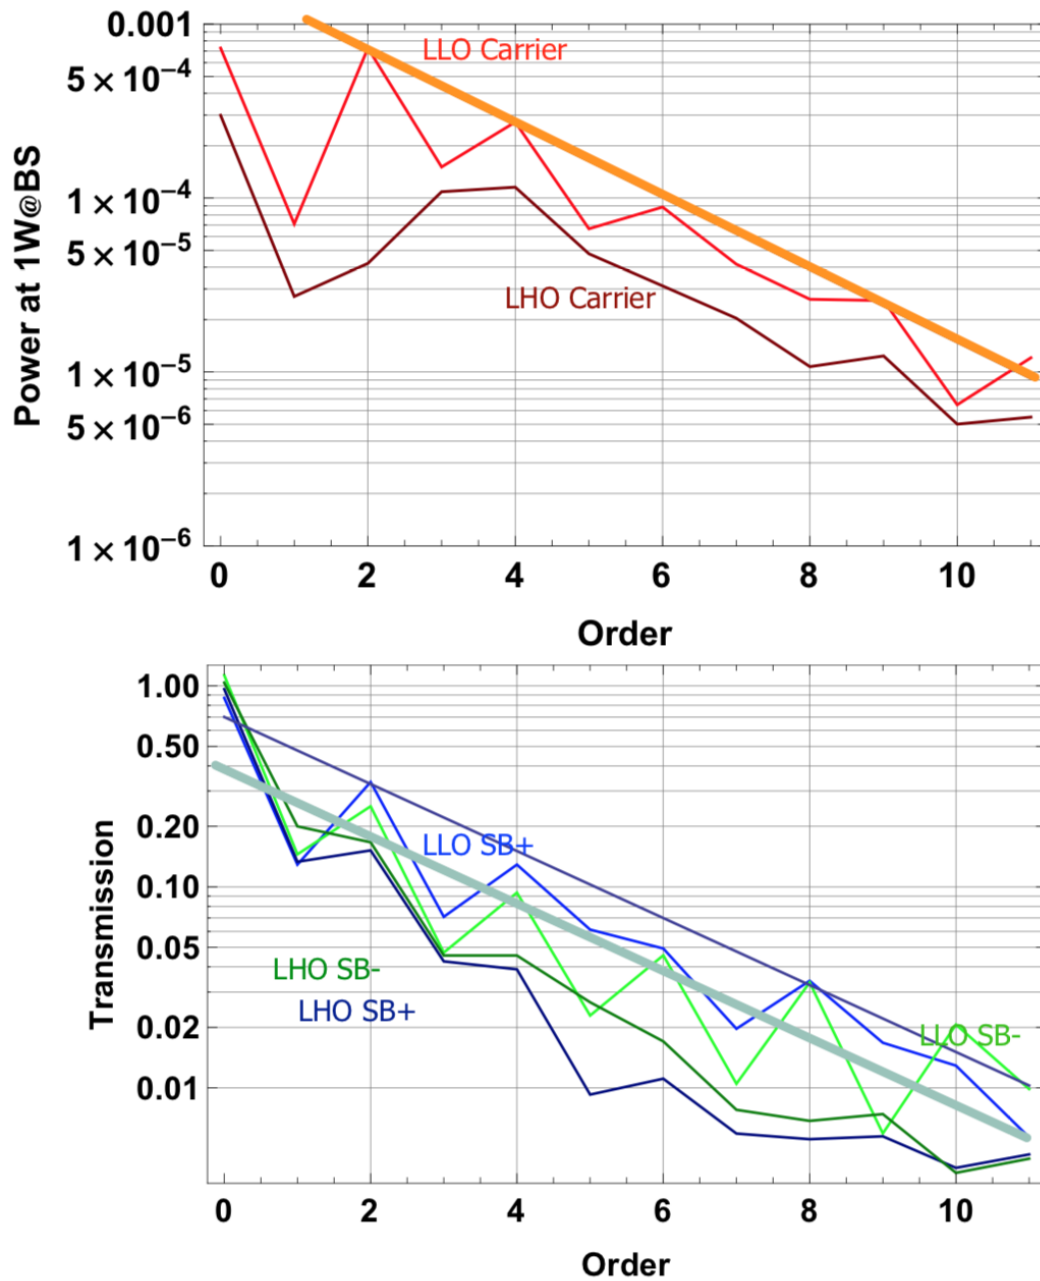


Figure 6.2: Power-law models of the eLIGO dark port field spatial structure. These power laws were used as weighting estimates when optimizing the aLIGO OMC parameters.

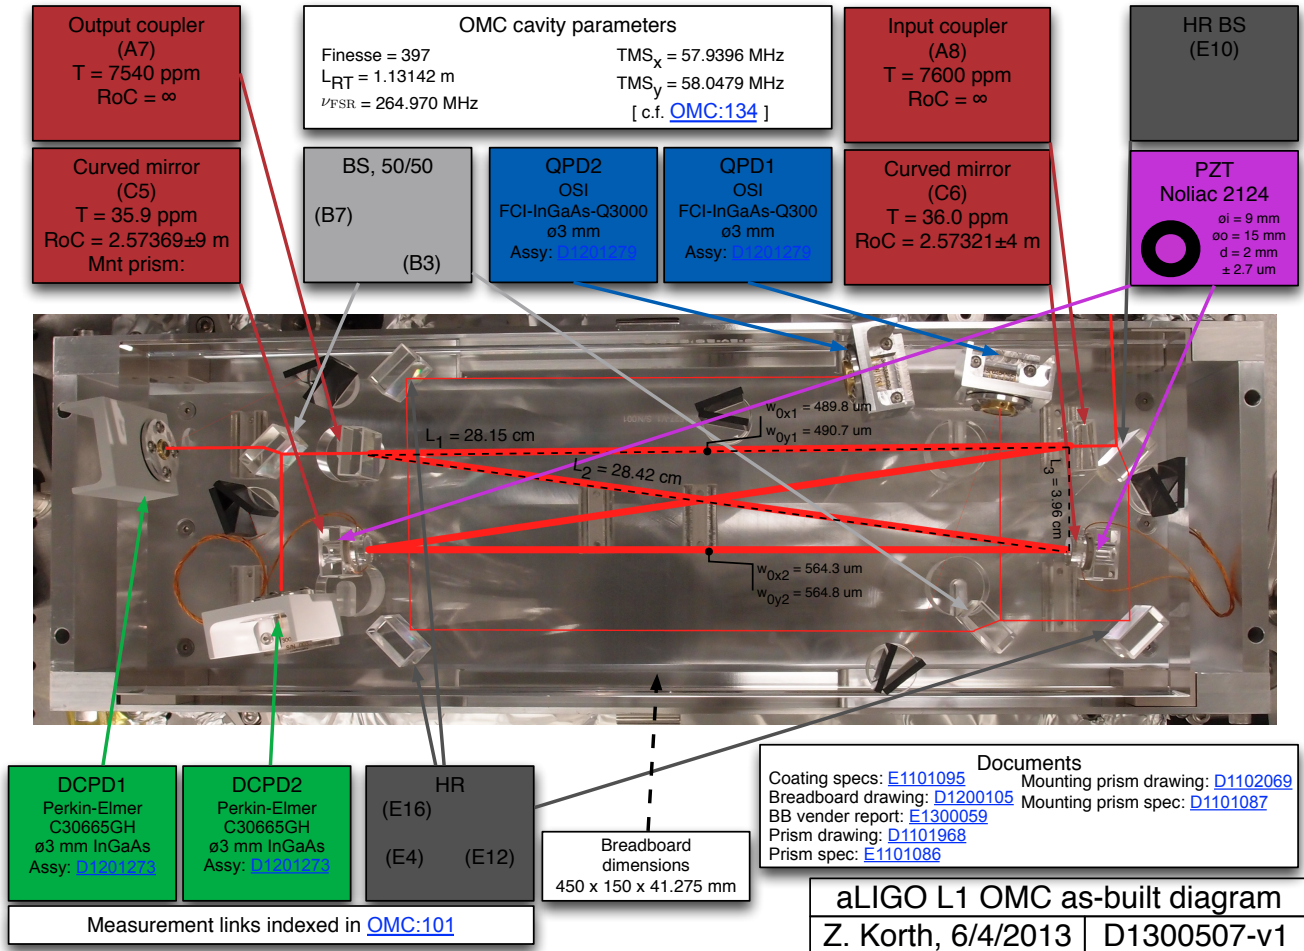


Figure 6.5: As-built diagram of the L1 OMC, highlighting some important dimensions and parameters and indicating for posterity the serial numbers of the various components used on this unit.

The cavity itself consists of two flat prisms coated to the coupling mirror transmissivity ($T_{\text{cpl}} \approx 7500$ ppm) and two curved mirror assemblies, each with a 1/2-inch curved mirror coated to a much lower transmissivity ($T_{\text{HR}} \approx 50$ ppm³). These curved mirrors are attached via PZT actuators to glass prisms that are similar to the aforementioned couplers, but that have holes in them to allow a beam to pass through them to the curved mirrors. These PZTs allow the cavity length to be minutely adjusted for the purposes of locking it to the interferometer beam. The prisms and the curved mirrors each have anti-reflective (AR) coatings on their non-cavity-facing sides.

³This was the average transmissivity of all curved mirror units; the ones chosen for the L1 OMC had slightly lower transmissivity, as indicated in Fig. 6.5.

In addition to the cavity mirrors, there are several other glass prisms affixed to the breadboard that serve as steering mirrors and beamsplitters between the input beam, the cavity, and various photodetectors described below. These optics have varying transmissivities depending on their function.

There are two pairs of photodetectors onboard the OMC. The most important is the set of two sensitive DC-coupled, single-area detectors that sit in transmission of the cavity. These units, aptly known as the **DCPDs**, are responsible for transducing the optical power transmitted through the OMC into an electrical signal that can be fed into the aLIGO data acquisition system. The other pair of photodetectors is a set of two quadrant photodetectors (QPDs) used to assist in the alignment of the input beam to the OMC cavity. In order to sense all four alignment degrees of freedom (i.e., displacement and angle in both the horizontal and vertical directions), these two QPDs must each be placed at a different point along the beam propagation direction⁴.

Each photodiode for each DCPD and QPD is held in its own metal enclosure equipped with a heatsink and is connected via cabling to its associated biasing and readout electronics. Due to their critical role in the GW detection process, the DCPD electronics are located in hermetically sealed pods near the OMC assembly in the vacuum system; the QPD electronics are located just outside the vacuum.

The last optical components on the breadboard are a set of beam dumps used to control stray light. These dumps, made of “V”-shaped assemblies of black glass, are positioned at specific acceptance angles to expected ghost beams (e.g., in reflection from a photodiode) in such a way that the beam reflects internally multiple times before reemerging from the dump, thereby ensuring the beam is all but absorbed.

Finally, the other face of the breadboard is populated with a set of suspension attachment fixtures, as well as standoff structures allowing for strain relief and mechanical isolation of the various cables that carry signals to and from the OMC.

⁴More precisely, one must be placed near the location of the beam waist (where a local displacement at the QPD corresponds to a displacement in the cavity basis), while the other must be placed far from it (where a local displacement at the QPD corresponds to an angular misalignment in the cavity basis). Note that the QPDs are fed beams that have been picked off of the main beam path, so the locations described here are to be understood as *equidistant from the relevant beamsplitter* to the indicated locations of the main beam.

locked to the cavity using PDH feedback (see Appendix B) to the laser frequency. Then, phase modulation was applied with a broadband EOM, and an RF network analyzer was used to measure a transfer function from the EOM input to a high-speed photodetector placed in transmission of the cavity. This transfer function was measured with the laser intentionally misaligned from the cavity horizontally, and then again vertically. In this configuration, peaks are observed at intervals of the transverse mode spacing for the misaligned orientation, since the modulation sideband is partially transmitted at frequencies corresponding to these offsets from the carrier⁵. A diagram of the experimental setup can be found in Fig. 6.6.

Once all cavity mirror units were measured, sets of two flat mirrors and two curved mirrors were selected to form cavities. This was done to match transmissivities within each cavity for optimum transmission. After choosing all components for an OMC, the unit was built as follows:

1. Each curved mirror was bonded to an annular PZT, and these mirror-PZT stacks were in turn bonded to prisms to form two curved mirror assemblies. These components were bonded with Master Bond EP30-2 epoxy.
2. The beamsplitter that steers the input beam onto the input coupler was glued in place on the breadboard, followed by the input and output couplers themselves.
3. The two curved mirror assemblies were placed in their initial trial positions, with their fronts butted up against 2-actuator micrometer stages.
4. The two actuators on each micrometer stage were adjusted systematically to steer the cavity eigenmode to the centers of all mirrors.
5. Once the mode was centered on all mirrors, the cavity was locked with the curved mirrors still just resting on the breadboard.
6. While resonant—though typically requiring lock re-acquisition on each iteration—the cavity length was adjusted using all micrometer actuators until the desired cavity roundtrip length and TMS frequency were achieved.
7. Once the cavity exhibited the desired design parameters, the curved mirrors were also illuminated to cure them in place.

This in-situ metrology allowed for accurate tuning and freezing-in of the cavity parameters. A diagram of the optical test setup is shown in Fig. 6.7. Following

⁵In both cases, the transmitted beam was partially clipped between the test cavity and the high-speed photodetector. Otherwise, no signal would be observed, due to the orthogonality of all TEM modes within a given basis.

fabrication of the main cavity, the remaining auxiliary components and fixtures were aligned and bonded onto the breadboard according to the as-built locations of the beams.

After assembly and operational testing at Caltech, the OMCs were loaded into protective cases and sent to their respective future homes in Livingston, LA and Hanford, WA.

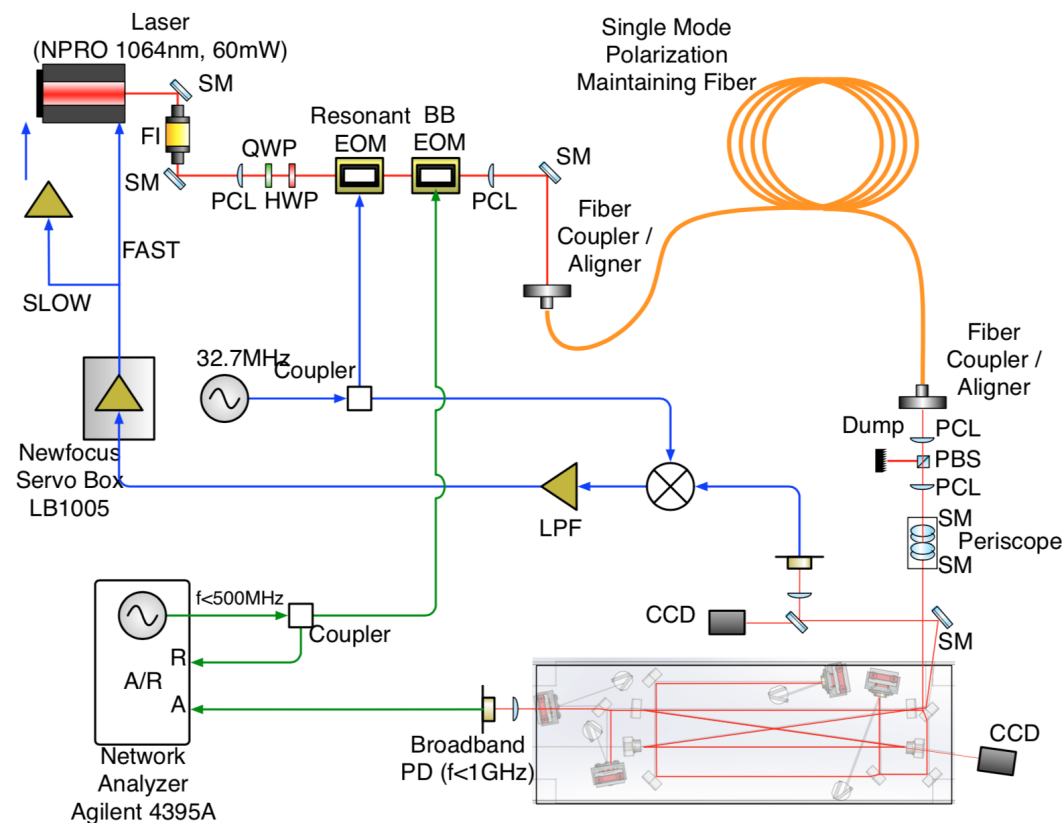


Figure 6.7: Diagram of the OMC test setup. This setup was used during the gluing of the OMC cavity mirrors, as well as to characterize the cavity after gluing was complete.

6.3 Installation

6.3.1 OMC suspension assembly

As with most of the major optical components in aLIGO, the OMC is suspended from a multi-stage suspension. The OMC suspension is housed in an aluminum frame that is bolted to the top of an active seismic isolation platform within the Horizontal Access Module 6 (HAM6) chamber. Two large maraging steel blade springs are attached to the top of the frame, from which an upper stage mass hangs via steel wire. The upper mass, in turn, has four smaller steel blade springs attached

to it, from which the OMC itself hangs, also via steel wires. The rigid-body motion of the OMC is therefore suppressed with respect to the motion of the seismic platform as f^{-4} above the pendulum resonances. For the purposes of damping the motion at these resonances and to allow for active control of the OMC's position and orientation, a set of six combination shadow sensor and magnetic actuator units are coupled to magnetic flags attached to the upper mass.

Prior to the arrival of the OMC, the OMC suspension was constructed and balanced using a dummy metal breadboard at the bottom stage. This procedure was performed in a clean environment as follows:

1. The overall frame was assembled, including the subframe and shroud that surround the upper mass and hold the sensor/actuator units.
2. The two upper and four lower suspension wires were pulled to design stress using weights, then clamped and cut.
3. The blade springs were clamped at their bases, and the clamped wires were attached to the tips of the blade springs.
4. The double suspension was hanged using first-guess balancing supplement masses, and then these masses were adjusted to adjust the height and balance of the full system.

Once the suspension was built and balanced, it was installed into the HAM6 chamber, awaiting the installation of the OMC optical bench. At the same time, three small, single-stage-suspended mirrors known as "tip-tilts" were also installed in the chamber. These mirrors form a telescope that focuses the beam exiting the main interferometer into the OMC mode. Additionally, they serve as actuators for maintaining alignment into the OMC.

Fig. 6.8 contains a photo of HAM6, showing the OMC in its suspension along with the tip-tilt mirrors. A bird's-eye diagram of the HAM6 layout, showing these components and some others that were added subsequently, can be found in Fig. 6.9.

6.3.2 Detector electronics

Initially, the aLIGO OMC was to reuse the readout electronics used in eLIGO. As such, prior to the OMC's arrival to the observatories, the twin photodiodes and hermetically sealed in-vacuum preamp electronics were removed from the eLIGO assemblies and characterized prior to re-integration [114].




Figure 6.8: A photo of the interior of HAM6, showing the OMC in its suspension as well as the three tip-tilts used for beam mode matching and alignment. All components sit atop a seismic isolation platform.

The in-vacuum DCPD preamps [115] consist of 1) a switchable wire-wound transimpedance resistor and 2) a two-stage active whitening filter, with each stage adding a zero at 0.8 Hz and poles at 80 Hz and 15.9 kHz. This whitening is necessary to provide enough gain at higher frequencies—where the interferometer signal is small—to overcome noise in the analog-to-digital conversion (ADC) electronics, while not saturating them at lower frequencies, where the interferometer signal is relatively quite large. The transimpedance is switchable to allow for variation in the sensed DC optical power.

Being at the crucial interface of the optical and electronic domains, the DCPDs must exhibit ultra low noise and have a very stable and well-modeled response. Thus, the two major characterization tests performed when repurposing these electronics

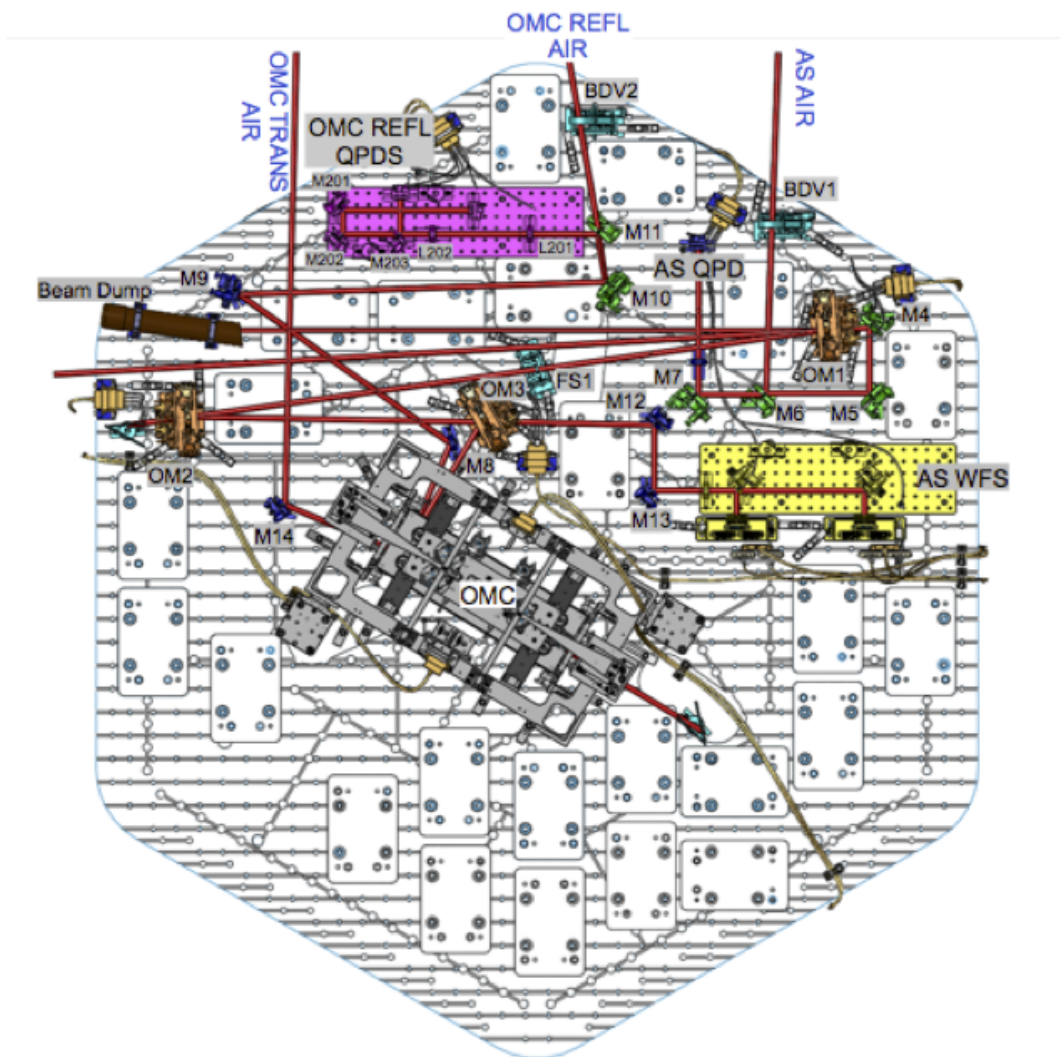


Figure 6.9: Bird’s-eye diagram of the HAM6 layout, showing the OMC and tip-tilts, as well as other interferometer components. (The tip-tilts are labeled “OM1-3”.)

from eLIGO were:

- **Input-referred noise:** The output noise of the preamps was measured with the biased diodes connected to the input, in a dark state. This output was then divided by the modeled transimpedance transfer function to give the *input-referred* noise of the system, in $A/\sqrt{\text{Hz}}$, and the result was compared with an analytic model for the noise of the system. The results of this test can be found in Fig. 6.10.
- **Transfer function:** The voltage transfer function was measured from the input

to the output of the preamps and compared with an analytic model. For this test, a voltage source was connected to the input node of the circuit (where the photodiode anode is usually connected), so the effect of the transimpedance resistor was negligible. However, this resistance was also measured for use with the compensation model described below. The results of this test can be found in Fig. 6.11.




Figure 6.10: Evaluation of the input-referred noise of the in-vacuum DCPD preamps. This measurement was made with both transimpedance values ($Z_{\text{high}} = 400 \Omega$ and $Z_{\text{low}} = 100 \Omega$).




Figure 6.11: Evaluation of the voltage transfer function of the two-stage whitening filter in the in-vacuum DCPD preamps, compared with an analytic model of the circuit. Higher-order models of the measured transfer functions were used in the digital compensation filters used to balance the response of these analog electronics.

6.4 Integration

6.4.1 Bring-up

After installing the OMC into its suspension and connecting the onboard electronics (as well as those for the tip-tilts) via vacuum feedthroughs to their supply and readout units in racks outside the chamber, the first task was to steer the interferometer output beam into rough alignment with the cavity. This was accomplished by configuring the main interferometer so that a stable beam of reasonable power (a few milliwatts) was delivered to HAM6, then manually adjusting the tip-tilt orientations in pitch and yaw (electronically) while tracing the faint multi-pass propagation of the beam within the OMC cavity with a sensitive IR viewer. During this process, the OMC length

actuator was ramped, and the alignment procedure continued until faint resonant transmission flashes were observed on the DCPDs. At this time, the locations of the pickoff beams on the onboard QPDs were recorded electronically as a rough reference of good alignment.

After obtaining approximate OMC alignment, further auxiliary optics were installed on the optical table to direct reflected and transmitted beam pickoffs to in-air cameras and photodetectors on tables adjacent to the HAM6 chamber. From this point forward, OMC integration work was largely performed remotely from the control room, though some iteration of auxiliary optic alignment was necessary once robust OMC locking was achieved. Ultimately, the chamber was closed and pumped down before any low-noise testing or interferometer operation.

6.4.2 Control system

6.4.2.1 Overview

The OMC is controlled with a handful of digital control loops (see Sec. 2.1.2). A simple diagram of the digital part of the control system—i.e., everything on the middle panel of Fig. 2.9—is shown in Fig. 6.12. In it, sensor signals are shown being received from ADC units, after which they are manipulated by the OMC digital control model before being sent back out into the analog domain via DACs. The length (top) and angle (bottom) sensing and control sections will be discussed more in the next two sections. Between these two is the “DARM normalization/linearization” block, whose purpose is to take the signal representing the power transmitted through the OMC and convert it to a differential arm (i.e., GW) signal as a function of some known and/or measured parameters (e.g., the DARM offset, circulating power, etc.).

6.4.2.2 Length sensing and control

Fig. 6.13 shows the OMC control screen, used by an interferometer operator to run the OMC subsystem. This screen shows a more detailed graphical representation of the OMC length and angular control loops.

On the top left are the DCPDs, which are used for both subsystems, in general. The small block at the front end takes the two independent PD signals, applies anti-whitening and transimpedance amplifier gain correction, then applies a weighting before generating difference (NULL), sum (SUM) and normalized sum (NORM) signals. The difference signal is used in some noise characterization methods, as

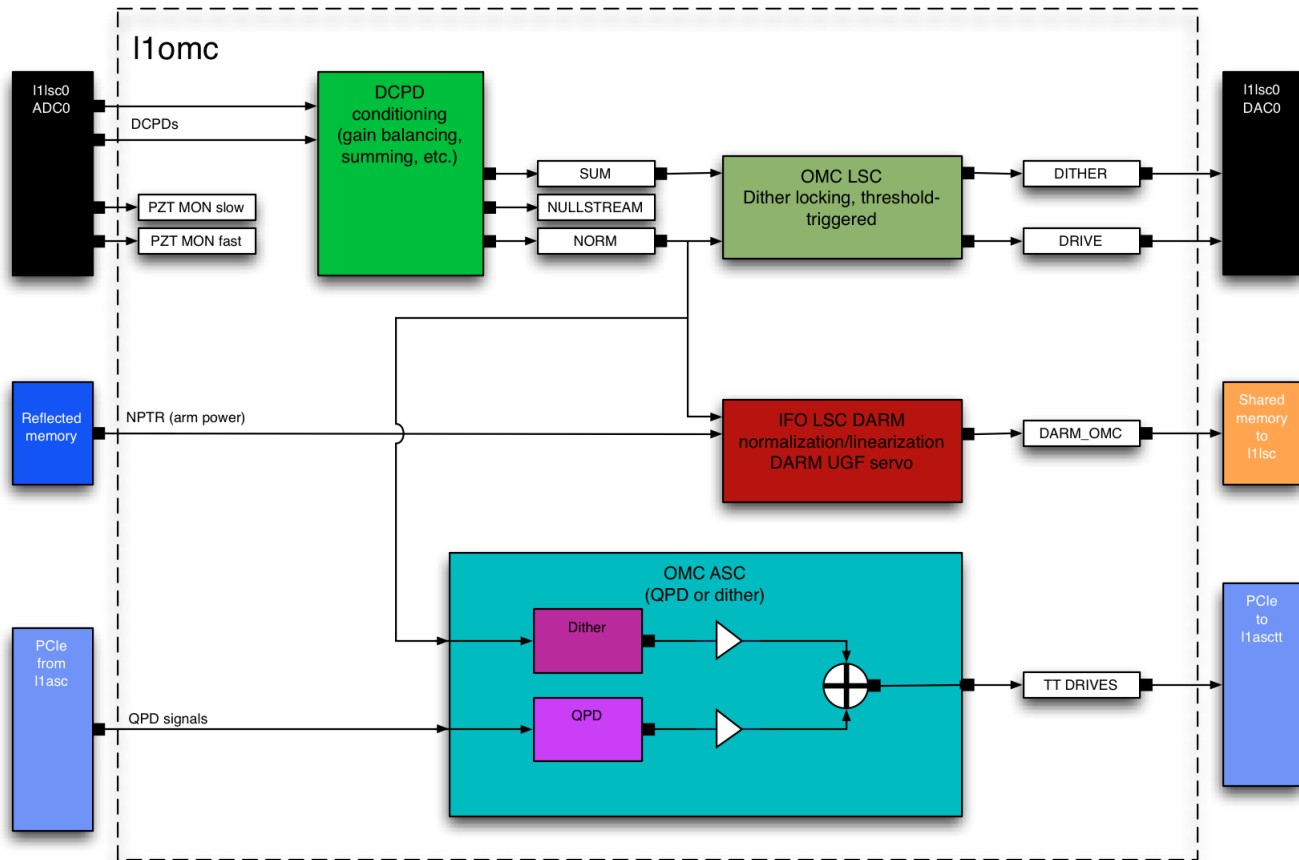


Figure 6.12: Simple diagram of the digital model for the OMC length and angular control.

well as to verify correct spectral reconstruction of the PD signals. The sum signal is used as a trigger to lock the length control loop, and more importantly as the input to the GW calibration block. The normalized sum is used for length and angular control⁶.

Length control of the OMC is accomplished via a technique known as **dither locking**, which is demonstrated in the diagram in Fig. 6.14. Essentially, this involves modulating the cavity length and then demodulating the measured output power at the same frequency. When the cavity is near but not quite on resonance, there will be a linear response proportional to the detuning (with the sign encoding the direction of the detuning). On resonance, by contrast, there is only a quadratic response, and the output power modulation will be at twice the modulation frequency. (This time-

⁶The normalized DCPD sum is used so that the overall gain of the OMC length control loop is not affected by the input power to the OMC (e.g., if the input power to the main interferometer changes, or if the DARM offset is changed).

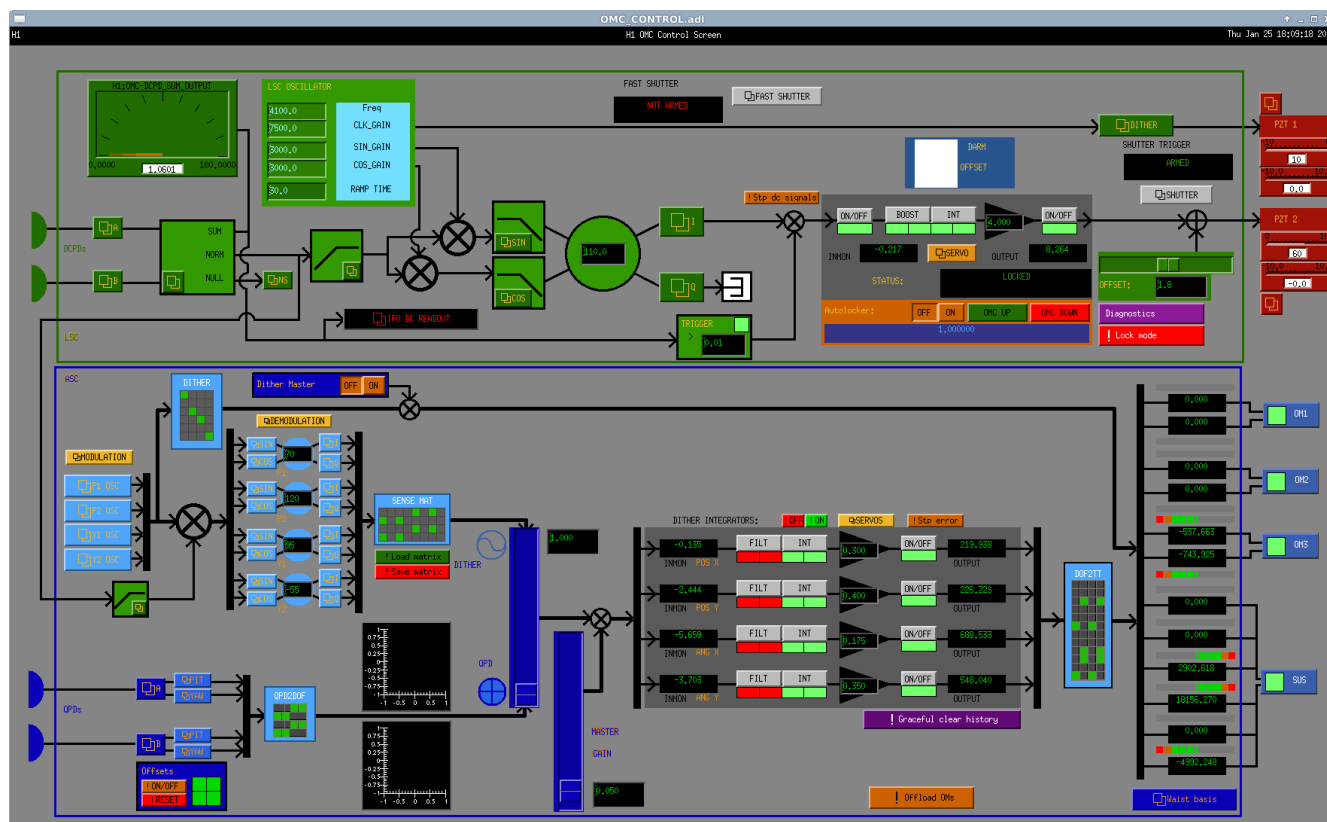


Figure 6.13: OMC control screen, showing a more detailed schematic view of the system.

domain picture is sufficient to describe the operation of a dither lock loop, but such a loop can also be described in the frequency domain as was done for PDH locking in Appendix B. In this case, the sidebands are generated via the length modulation within the cavity and resonate along with the carrier.)

A modulation signal at f_m is generated in the digital domain, then pumped through a DAC to one of the two PZT actuators on board the OMC. The same digital signal is then used to demodulate the normalized DCPD sum signal, which is first high-passed to reduce any DC or low-frequency information⁷. To account for potential phase lag incurred by the dither signal between where it is generated and when it returns in the digitized photocurrent, it is demodulated with both sine and cosine versions of the digital dither excitation, and an arbitrary phase can be derived from a combination of the two.

The demodulated and phase-corrected signal, now a linear error signal indicating a

⁷Due to the dither locking process, all the required cavity length information is upshifted to a band centered on f_m .

signed offset of the OMC length from the laser frequency, is passed through a digital servo filter, then out through the DAC and to the second of the two onboard PZTs.

The OMC length dither frequency f_m is typically a few kHz, and the length loop is set to be a simple integrator with a UGF around 100 Hz.




Figure 6.14: Dither locking. A sinusoidal modulation is applied to the cavity length. When the cavity is not tuned to the laser frequency, a signal can be seen in the output power at the drive frequency, with the phase of the signal encoding whether the cavity is too long or too short. On resonance, a modulation only produces a signal at twice the modulation frequency.

6.4.2.3 Angular sensing and control

The OMC angular control system can also be seen in Fig. 6.13. Towards the bottom right, the QPD signals can be seen entering the control model. From here, geometric relations are used to convert the QPD X and Y signals (corresponding to local displacements at their respective planes) into generalized displacements of the four cavity angular degrees of freedom: displacement in X, displacement in Y, pitch about the cavity waist, and yaw about the cavity waist. These “waist basis”

displacements are then used as one available set of error signals for four independent angular control loops.

The angular control loops act on a combination of the tip-tilt pitch and yaw, as well as the OMC suspension pitch, yaw and roll. These actuators—twelve in all—are clearly more than the four that are necessary to control the angular DoFs. However, this redundancy allows for independent control of the beam position and angle at other locations in HAM6. To determine the output matrix that maps the four waist basis control loops to actuators, the following algorithm is used: first, a usable choice of actuator or actuators is made for each degree of freedom. Then, these actuators are stepped by a fixed amount, and the displacements are measured in the waist basis. Finally, this measurement matrix is inverted to give the mapping from DoF to actuator.

For various reasons, using the DC-coupled QPDs as actuators is not the most effective method. For example, junk light on the input beam to the OMC might create a situation where, even though the total power is balanced across quadrants of the QPDs, the TEM_{00} component itself is misaligned. A better approach is to use the dither technique as used in the length control. Since the power transmission of an optical cavity is quadratically dependent to small transverse displacement or angle misalignments, the same exact principle can be used. In this case, the tip-tilts are dithered in angle at four different frequencies (one each for pitch and yaw on each of the two dithered mirrors). A “sensing matrix” is constructed by displacing the input beam by known amounts in each DoF, then observing the resultant output in each demodulated signal.

Still more advanced techniques for OMC alignment have been proposed and tested in eLIGO and aLIGO. An excellent description of the overall problem of OMC angular alignment and these various techniques is given in [111]. In practice, QPD signals are used to acquire rough alignment of the OMC to achieve lock, whereafter more sensitive sensors are swapped in for low-noise operation.

6.4.3 Advanced LIGO commissioning with the OMC

The OMC is an essential component of the aLIGO detector. In addition to its main role in transducing optical signals from the main interferometer into electrical signals for storage in computers—including all GW signals recorded to date—it has been and continues to be used as an important diagnostic tool. All these detections and much of the useful diagnostic work the OMC has seen have come

after I completed my fellowship and returned to Caltech. However, the following two sections highlight nice results that were achieved with the newly commissioned OMC prior to my departure.

6.4.3.1 AS port mode evaluation

During initial integration testing of the OMC, one of the first tests was to use it to evaluate the spatial mode of the beam incident on it from the main interferometer. This test was performed by aligning the input beam and the OMC as best as possible, then slowly sweeping the cavity length over a full FSR and observing the transmitted light. The output beam was also monitored on a CCD camera at the same time, so that peaks could be identified with specific spatial modes.

As seen in Fig. 6.15, the presence of a large component in the “donut” mode was observed in the scan, indicating a relatively large mismatch on the order of 25%. This measurement and subsequent measurements were combined with thermal lens modeling of the interferometer [116] to determine what optic(s) had to be moved to correct the mismatch.

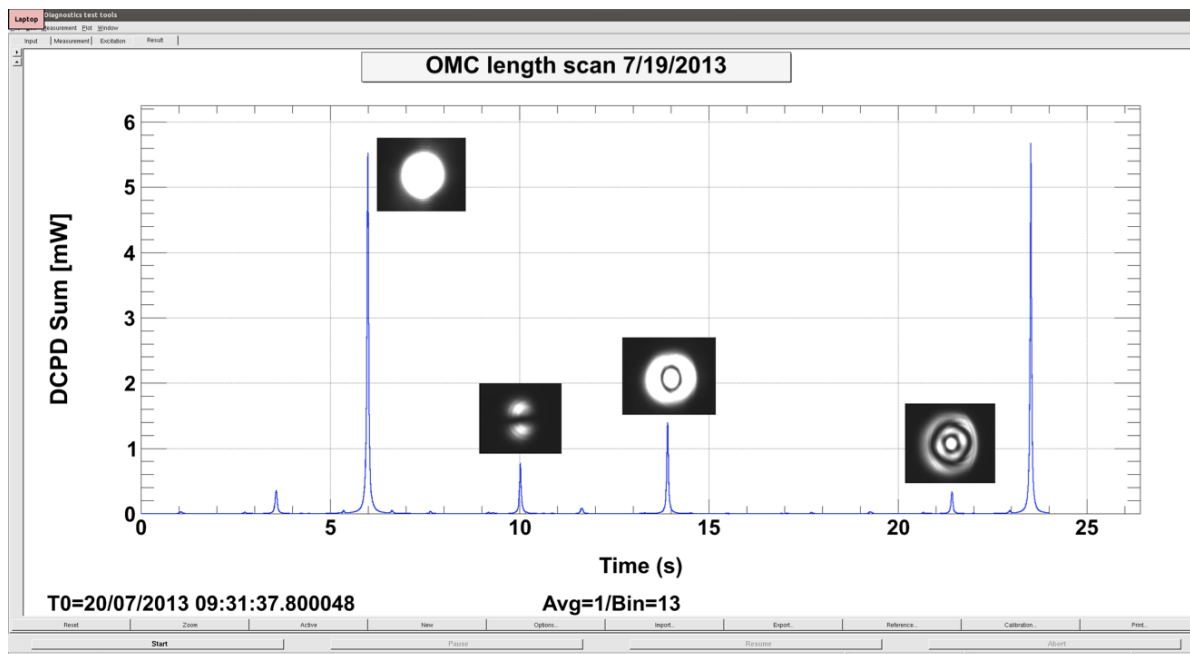


Figure 6.15: Length scan of the OMC while closely aligned to the interferometer output beam. The presence of the donut mode with appreciable amplitude indicates a relatively large mode mismatch error between the OMC and the main interferometer.

6.4.3.2 Power-recycled Michelson interferometer (PRMI) testing

Prior to commissioning of the arm cavities, the most sensitive interferometer configuration was the power-recycled Michelson interferometer (PRMI), operated with the Michelson length DoF (MICH) sensed via DC readout. Following success with using DC readout to lock the simple (non-power-recycled) Michelson, power recycling was configured and the interferometer was locked on PDH (RF) readout for all DoFs. Then, a handoff procedure was initiated whereby the MICH sensor was transitioned from the RF signal to the OMC DCPD-derived signal.

Upon completing the transition to DC readout, the noise of the PRMI was measured, and it was found to be in agreement with the expected noise generated by modeling over wide bands. In particular, the MICH noise settled on the expected contribution from input laser power fluctuations (“intensity noise”), which confirmed the proper operation of the OMC.




Figure 6.16: Noise in the low-noise power-recycled Michelson interferometer (PRMI) during aLIGO commissioning. The convergence of the measured displacement noise with the expected level from measured intensity noise at high frequencies confirmed the operation of the OMC in its design DC readout function.

Part IV

Suppression of Quantum Radiation Pressure Noise in an Optical Spring

This part of the thesis is adapted from the published work in [117], written by the author, with a large contribution from H. Miao. Some of the thermal noise concepts discussed here are explained in further detail in Part V. A detailed set of accompanying calculations—published along with the material in this part as an appendix to [117]—can be found in Appendix C.

SUPPRESSION OF QUANTUM RADIATION PRESSURE NOISE IN AN OPTICAL SPRING

7.1 Introduction

Catalyzed by vast improvements in micro- and nanofabrication processes, the field of cavity optomechanics has seen a recent boom in interest [118–120]. In addition to providing a means for quantum-limited force measurements [121], e.g., in gravitational-wave detection [122] and scanning probe microscopy, optomechanical devices can also be used to probe the quantum behavior of mechanical systems. Recently, several experiments have demonstrated the cooling of a resonator down to its quantum ground state via cryogenics or optomechanical interaction [123–125]. In addition, more than one group [126, 127] has demonstrated the so-called optomechanically induced transparency (OMIT) effect, an analog of the electromagnetically induced transparency (EIT) [128, 129] effect observed in atomic systems. This effect can be used to make narrow-band quantum filters, e.g., to effect the frequency-dependent phase rotation of squeezed light injected to enhance the sensitivity of quantum noise-limited interferometric gravitational wave detectors [130, 131]. This also opens up the possibility for the processing and storing of nonclassical states of light through coherent transfer of quantum states between light and a mechanical oscillator, a technique that would find much use in the emergent field of quantum information processing.

The ubiquitous bath of thermal energy presents a major obstacle to these efforts, randomly exciting a system and masking its underlying quantum nature. A characteristic figure of merit for quantifying this thermal decoherence effect is given by the ratio of the thermal occupation number \bar{n}_{th} and the mechanical quality factor Q :

$$\frac{\bar{n}_{\text{th}}}{Q} = \frac{k_B T}{\hbar \omega_m Q} \propto (Qf)^{-1}, \quad (7.1)$$

where k_B is the Boltzmann's constant and $f = \omega_m/2\pi$ is the mechanical frequency. When this ratio becomes smaller than one, the oscillator quantum state will survive longer than one oscillation period before the thermal effect destroys it.

As is apparent from Eq. (7.1), the quantum state lifetime is ultimately limited by the product of the quality factor Q and mechanical frequency f . A significant

body of research has focused on increasing this “ Qf product” for a wide range of mechanical systems. If Q were truly a frequency-independent quantity—as in the “structural damping” model as described by Saulson [132]—then moving to higher eigenfrequencies would lead to an immediate improvement. In the opposite direction, there are many experiments that would benefit from the use of low-frequency (sub-kHz) resonators. A number of bulk structures have been found to exhibit extremely high Q in this frequency range [133, 134]; unfortunately, such systems tend to have relatively large (gram- to kg-scale) effective masses, making them unsuitable for typical optomechanics experiments. The realization of sub-microgram effective masses requires the use of nanofabricated resonators. In practice, excess damping from surface effects [135], phonon tunneling loss [136] or intrinsic mechanisms such as thermoelastic [137] and Akhiezer damping [138] limits the achievable Q and thus the Qf product in these devices. In addition, we add the further requirement that the desired system exhibit excellent optical quality (i.e., high reflectivity owing to low scatter loss and absorption), which limits the resonator options considerably, especially in light of the fact that typical dielectric materials used to create multi-layer optical coatings (e.g., $\text{SiO}_2/\text{Ta}_2\text{O}_5$) exhibit low mechanical quality factors [139]. Here, we propose a method for using the optical spring effect in linear optomechanical devices [140–145] to increase the effective Q of a given mechanical resonator, while simultaneously suppressing the quantum radiation pressure noise that would normally be imparted by the optical fields. This technique should facilitate the creation of an oscillator with a Qf product considerably higher than those available today, enabling useful applications in quantum metrology and also creation of long-lived quantum states at lower frequencies than were previously practical.

The concept in this paper makes use of the fact that when a strong optical spring is linearly coupled to a mechanical resonator, the resonator’s Hamiltonian becomes augmented or even dominated by contributions from the radiation pressure forces of the optical fields. In this way, the bare resonator’s thermal noise is “diluted” by the ratio of the intrinsic elastic energy to that stored in the optical field [144]. Typically, the modification of a resonator’s dynamics via linear coupling is accompanied by excess noise from quantum back-action—the quantum fluctuations of the radiation pressure, in our case. This has been identified as a serious issue in the strong dilution regime by Chang *et al.* [146] and Ni *et al.* [147], who instead propose to achieve optical dilution by using a nonlinear quadratic optical potential to trap a partially reflective membrane [148], which would be immune to linear quantum back-action.

The device we propose evades such parasitic quantum back-action by detecting it in the outgoing field and actively feeding back to the system, resulting in a nearly noise-free optical spring. Since this method allows for straightforward coupling of the diluted mechanical resonator to an external optical system from the reverse side of the resonator, it can be used as a “black-box” effective mechanical resonator of exceptionally high quality.

7.2 Optical spring

The canonical optomechanical system is shown in the dashed box in Fig. 7.1. In such a system, the “optical spring” effect arises from dynamical back-action of the optical cavity field on the mechanical oscillator forming one cavity boundary. The mechanical oscillator displacement \hat{x} is coupled to the cavity field \hat{a} via radiation pressure, as described by the following interaction Hamiltonian [149]:

$$\hat{\mathcal{H}}_{\text{int}} = \hbar G_0 \hat{x} (\bar{a}^* \hat{a} + \bar{a} \hat{a}^\dagger) \equiv -\hat{x} \hat{F}_{\text{rad}}. \quad (7.2)$$

The coupling constant is $G_0 = \omega_c/L$; \bar{a} is the classical mean amplitude of \hat{a} due to coherent driving of an external laser; ω_c is the cavity resonant frequency; L is the cavity length. When the frequency of the external laser ω_0 that drives the cavity field is detuned from ω_c , \hat{F}_{rad} depends on the oscillator displacement, creating a mechanical response that mimics a spring. More specifically, \hat{F}_{rad} in the frequency domain can be written as (the details of the derivation are found in Appendix C):

$$\hat{F}_{\text{rad}}(\omega) = -K_{\text{os}}(\omega) \hat{x}(\omega) + \hat{F}_{\text{noise}}(\omega), \quad (7.3)$$

where the optical spring coefficient K_{os} is approximately given by

$$K_{\text{os}} \approx -\frac{2\hbar G_0^2 |\bar{a}|^2 \Delta}{\Delta^2 + \gamma^2} - \frac{4i\hbar G_0^2 |\bar{a}|^2 \gamma \Delta \omega}{(\Delta^2 + \gamma^2)^2} \equiv m\omega_{\text{os}}^2 - im\Gamma_{\text{os}}\omega, \quad (7.4)$$

with the cavity detuning $\Delta \equiv \omega_c - \omega_0$ and γ being the cavity bandwidth. Here, the approximation is taken for the case of large detuning and cavity bandwidth, which we will show to be the relevant parameter regime for realization of this idea. In addition, we have introduced the optical spring frequency ω_{os} and the optical damping Γ_{os} . As we can see, when the detuning is negative, i.e., $\Delta < 0$, the optical rigidity is real and positive, and the optical damping is negative Γ_{os} (heating), and vice versa. By introducing an additional driving field with a different detuning frequency, one can create the so-called stable double optical spring [143], which exhibits both positive rigidity and positive damping (we will elaborate on this issue later). The optical




Figure 7.1: Simplified experimental layout, with the canonical optomechanical system shown within the dashed box. Input vacuum fluctuations drive the cavity mode, which in turn exerts radiation pressure forces on the mechanical resonator forming the cavity boundary. The output mode of the cavity is sensed with a photodetector, and—in the relevant parameter regime—the measured power contains the radiation pressure fluctuations that drive the resonator with minimal sensitivity to the resonator position. This signal can therefore be used as an error signal for feeding back to the laser amplitude to suppress the radiation pressure noise on the resonator.

spring modifies the mechanical susceptibility $\chi_0(\omega)$, defined through $\chi_0(\omega) \equiv \hat{x}(\omega)/\hat{F}(\omega)$, from its original value $\chi_0^{-1}(\omega) = -m(\omega^2 + i\gamma_m\omega - \omega_m^2)$ to an effective one:

$$\chi_{\text{eff}}^{-1}(\omega) = -m[\omega^2 + i(\gamma_m + \Gamma_{\text{os}})\omega - (\omega_m^2 + \omega_{\text{os}}^2)]. \quad (7.5)$$

For a strong optical spring $\omega_{\text{os}} \gg \omega_m$, we can significantly stiffen the mechanical oscillator with the restoring energy from the optical field.

One immediate issue with this approach comes from the quantum radiation pressure noise $\hat{F}_{\text{noise}}(\omega)$ in Eq. (7.3), which arises from quantum fluctuation of the optical field:

$$\hat{F}_{\text{noise}}(\omega) \equiv \frac{2\hbar G_0 |\bar{a}| \sqrt{\gamma}}{\sqrt{\gamma^2 + \Delta^2}} \left[\frac{(\gamma^2 + \Delta^2 - i\gamma\omega)\hat{v}_1 + i\Delta\omega\hat{v}_2}{(\omega - \Delta + i\gamma)(\omega + \Delta + i\gamma)} \right], \quad (7.6)$$

where $\hat{v}_1 \equiv (\hat{a}_{\text{in}} + \hat{a}_{\text{in}}^\dagger)/\sqrt{2}$ and $\hat{v}_2 \equiv (\hat{a}_{\text{in}} - \hat{a}_{\text{in}}^\dagger)/\sqrt{2}i$ are the amplitude and phase quadratures of the input optical field. This additional noise term will increase the effective temperature of the thermal bath and drive the mechanical oscillator away from the quantum regime, as pointed out by Chang et al. [146]. In the large bandwidth and detuning regime, this reduces to

$$\hat{F}_{\text{noise}}(\omega) \approx -\frac{2\hbar G_0 |\bar{a}| \sqrt{\gamma}}{\sqrt{\gamma^2 + \Delta^2}} \hat{v}_1(\omega), \quad (7.7)$$

indicating that the radiation pressure noise is dominated by fluctuations in the amplitude quadrature of the input field. The strength of this noise can be quantified by its spectral density:

$$S_F(\omega) \approx \frac{4\hbar^2 G_0^2 |\bar{a}|^2 \gamma}{\gamma^2 + \Delta^2}. \quad (7.8)$$

From the above expression and Eq. (7.4), we learn the optical rigidity (real part of K_{os}) scales with the optomechanical coupling strength in the same way as the radiation pressure noise:

$$K_{\text{os}}, S_F \propto G_0^2 |\bar{a}|^2. \quad (7.9)$$

Essentially, this means that an increase in the optical spring frequency is accompanied by an increase in the radiation pressure noise when we scale up the optical power.

7.3 Evading quantum radiation pressure noise

To solve the aforementioned issue, we make use of the fact that the output field emerging from the cavity contains information about the quantum radiation pressure noise that has been imposed onto the mechanical oscillator. In particular, as we will show, for the large bandwidth and detuning limit, the power fluctuations in the output field originate from the same quadrature that is responsible for the radiation pressure noise. The photodetector measures these fluctuations, and by feeding this signal back to the mechanical oscillator with the correct filter, we can therefore evade the quantum radiation pressure noise and achieve a nearly noiseless optical spring. Note that this does not violate the fundamental principle of quantum measurement—any linear continuous measurement of a dynamical variable that does not commute at different times (non-conservative) is associated with quantum back-action on that variable [121]; here, we only sense the quantum radiation pressure noise and have almost no sensitivity to the mechanical displacement, and that is why we can evade such back-action noise.

To elaborate on this idea, we use the standard input-output relation for this system $\hat{a}_{\text{out}}(\omega) = -\hat{a}_{\text{in}}(\omega) + \sqrt{2\gamma} \hat{a}(\omega)$, and it, for high bandwidth and detuning, gives [refer to Appendix C for more detail]:

$$\hat{a}_{\text{out}}(\omega) \approx -\frac{\Delta + i\gamma}{\Delta - i\gamma} \hat{a}_{\text{in}}(\omega) - \frac{\sqrt{2\gamma} G_0 \hat{a}}{\Delta - i\gamma} \hat{x}(\omega). \quad (7.10)$$

Accompanying these input fluctuations is a classical mean amplitude, \bar{a}_{in} , and we can define a phase reference for the system by setting this field to be real and positive.

This field also receives a phase shift upon interaction with the cavity:

$$\bar{a}_{\text{out}} = -\bar{a}_{\text{in}} + \sqrt{2\gamma}\bar{a} = -\frac{\Delta + i\gamma}{\Delta - i\gamma}\bar{a}_{\text{in}}. \quad (7.11)$$

The power fluctuation measured by a photodetector placed at the cavity output reads: $\delta\hat{P} \equiv \bar{a}_{\text{out}}^*\hat{a}_{\text{out}} + \bar{a}_{\text{out}}\hat{a}_{\text{out}}^\dagger$. In our stated limit and in the frequency domain, this fluctuating piece is given by

$$\delta\hat{P}(\omega) = \bar{a}_{\text{out}}^*\hat{a}_{\text{out}}(\omega) + \bar{a}_{\text{out}}\hat{a}_{\text{out}}^\dagger(\omega) \approx \sqrt{2}\bar{a}_{\text{in}}\hat{v}_1(\omega). \quad (7.12)$$

Therefore, due to the common phase rotation experienced by the DC and fluctuating components [c.f. Eqs. (7.10) and (7.11)], the output power is still a measure of the amplitude fluctuations of the input field. As shown in Eq. (7.7), it is this quadrature responsible for the radiation pressure back-action on the resonator, and so the noise can be suppressed by feeding this signal back to the amplitude of the pump laser.

7.4 Residual radiation pressure noise

While strong radiation pressure noise cancellation can be achieved using this technique, a small fraction cannot be canceled owing to two effects: (i) *optical loss* due to imperfection of the cavity and non-unity quantum efficiency in photodetection, which will introduce vacuum noise that is uncorrelated with \hat{v}_1 and \hat{v}_2 ; (ii) *finite cavity bandwidth and detuning* that modifies the input-output relation to give residual parasitic sensitivity to the oscillator displacement \hat{x} , which we have thus far ignored by assuming very large bandwidth and detuning. In actual experimental setups, there is always certain amount of optical loss, and the bandwidth and detuning are both finite.

By taking these effects into account, we see that the total measured power fluctuation is

$$\delta\hat{P}(\omega) = \sqrt{2}\bar{a}_{\text{in}}\hat{v}_1(\omega) + \delta\hat{P}_\epsilon(\omega) + \delta\hat{P}_\eta(\omega) + \delta\hat{P}_x(\omega). \quad (7.13)$$

Here, the second term

$$\delta\hat{P}_\epsilon(\omega) = \frac{2\sqrt{2\gamma\gamma_\epsilon}\bar{a}_{\text{in}}}{\gamma^2 + \Delta^2}(\gamma\hat{v}'_1 - \Delta\hat{v}'_2) \quad (7.14)$$

arises from the vacuum fluctuation $\hat{v}'_{1,2}$ (uncorrelated with $\hat{v}_{1,2}$) due to optical loss in the cavity, and $\gamma_\epsilon \equiv c\epsilon/(4L)$ with ϵ being the round-trip optical loss in the cavity; the third term,

$$\delta\hat{P}_\eta(\omega) \approx \sqrt{2}\bar{a}_{\text{in}}\sqrt{1-\eta}\hat{n}, \quad (7.15)$$

comes from the non-unity quantum efficiency, η , of the photodetector (here, \hat{n} is the vacuum fluctuation associated with this loss port); the last term represents the parasitic position sensitivity:

$$\delta\hat{P}_x(\omega) = -\frac{2G_0|\bar{a}|^2\Delta(2\gamma_\epsilon - i\omega)}{\gamma^2 + \Delta^2}\hat{x}(\omega), \quad (7.16)$$

which arises both from the intracavity loss and from the first-order correction to the frequency dependence due to finite bandwidth and detuning, and the associated quantum back-action (radiation pressure) noise reads:

$$\hat{F}_{\delta\hat{P}_x} = -\frac{2\hbar G_0\sqrt{\gamma}|\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}} \left[\sqrt{\frac{\gamma_\epsilon}{\gamma}} \hat{v}'_1 + \frac{i\omega\Delta}{\gamma^2 + \Delta^2} \hat{v}'_2 \right]. \quad (7.17)$$

Using this modified photodetector output, we compute a full, closed-loop noise model of the system (details in the Appendix). In the ideal feedback limit—i.e., for infinite open-loop gain—the residual force noise

$$\hat{F}^{\text{res}}|_{\text{gain} \rightarrow \infty} = \frac{2\hbar G_0\sqrt{\gamma}|\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}} \left\{ \sqrt{\frac{\gamma_\epsilon}{\gamma}} \frac{(\gamma^2 - \Delta^2)\hat{v}'_1 + 2\gamma\Delta\hat{v}'_2}{\gamma^2 + \Delta^2} - \sqrt{1 - \eta} \hat{n} - \frac{i\omega\Delta\hat{v}'_2}{\gamma^2 + \Delta^2} \right\}, \quad (7.18)$$

of which the spectral density reads:

$$S_F^{\text{res}} = \frac{4\hbar^2 G_0^2 \gamma |\bar{a}|^2}{\gamma^2 + \Delta^2} \left[\frac{\gamma_\epsilon}{\gamma} + (1 - \eta) + \frac{\omega^2 \Delta^2}{(\gamma^2 + \Delta^2)^2} \right]. \quad (7.19)$$

By comparison with the thermal force spectrum from a viscous damping model, $S_F^{\text{th}} = 4m\gamma_m k_B T$, we can assign an effective temperature to this residual force noise as

$$T_{\text{eff}}^{\text{res}} \equiv \frac{S_F^{\text{res}}}{4m\gamma_m k_B}. \quad (7.20)$$

In order not to dominate, this residual temperature must be kept below the environmental temperature.

Another interesting result of this closed-loop analysis is that, again for an infinite loop gain, the effective mechanical susceptibility of the resonator becomes

$$\chi_{\text{eff}}^{-1} \rightarrow \chi_{\text{eff}'}^{-1} = -m[\omega^2 + i\gamma_m\omega - (\omega_m^2 + \omega_{\text{os}}^2)]. \quad (7.21)$$

Comparing this with Eq. (7.5), we see that the damping contribution from the optical spring, Γ_{os} , is removed. Recall that, for an optical spring with a positive restoring

force, we have negative damping: $\Gamma_{\text{os}} < 0$ [cf. Eq. (7.4)], and it is for this reason that a second optical spring field is usually necessary to make the system stable—the double optical spring scheme [143], discussed in the next section. In our case, if the loop gain is high enough (i.e., if $G_{\text{OL}} \gg |\Gamma_{\text{os}}|/|\gamma_m|$), the negative damping will be removed due to the finite response to the mechanical displacement, indicated by Eq. (7.16), and therefore the system can be stabilized by the positive internal damping of the mechanical system. A practical issue for implementing this is that the required gain could be high in certain applications, and a double optical spring can therefore be used to ease the requirement.

7.5 Experimental realization with double optical spring

In the following, we will detail a proposed experiment, in which a mechanical oscillator with a bare resonance frequency of $\omega_m/2\pi = 100$ Hz is optically stiffened to a new, optomechanical resonance of $\omega_{\text{os}}/2\pi \approx 100$ kHz, leading to a commensurate increase in its effective Q factor. This parameter regime is chosen because—due to the low natural loss rate of the resonator—it highlights the long thermal decoherence timescales achievable with such a technique.

Despite the active stabilization effect discussed above, it may be impractical to use a single optical spring due to the very high feedback gains required¹. Instead, we consider a novel approach proposed in Ref. [143] that uses a second optical spring field to create a passively stable system. The linear combination of two K_{os} s, with one red-detuned and the other blue, can be made to exhibit both positive restoring and damping, resulting in a passively stable spring. The sum of the two optical spring contributions is thus:

$$K_{\text{os}}^{\text{tot}} \approx -im\omega \left[\frac{\gamma_B \omega_{\text{osB}}^2}{(\gamma_B^2 + \Delta_B^2)} - \frac{\gamma_R \omega_{\text{osR}}^2}{(\gamma_R^2 + \Delta_R^2)} \right] + m\omega_{\text{osB}}^2 - m\omega_{\text{osR}}^2 \quad (7.22)$$

where γ_B, γ_R and Δ_B, Δ_R are the cavity bandwidth and detuning as seen by the blue and red fields, respectively (note that $\Delta_B < 0$). For a proper choice of these parameters as a function of the ratio $|\omega_{\text{osB}}/\omega_{\text{osA}}| > 1$, the expression in the brackets can be made to vanish, and the effective resonator is stiffened without instability

¹In our example below, using a single optical spring would dictate optical damping Γ_{os} on the order of $2\pi \times 1$ kHz. The mechanical damping is $\gamma_m \approx 2\pi \times 10^{-4}$ Hz, and therefore the required gain at the optical spring frequency of 100 kHz is $G_{\text{OL}}^{\text{req}} \approx 10^5$. In practice, obtaining laser amplitude actuation bandwidths above ~ 1 MHz is quite challenging, and so it would be difficult to implement a stable loop in this case.

parameter	m	L	$\omega_m/2\pi$	Q	$\gamma_B/2\pi$	$\Delta_B/2\pi$	P_B	$\gamma_R/2\pi$	$\Delta_R/2\pi$	P_R
value	250 ng	1 mm	100 Hz	10^6	20 MHz	-20 MHz	390 mW	4 MHz	4 MHz	16 mW

Table 7.1: A sample set of parameters. These values generate an optical spring with $\omega_{os}/2\pi \approx 100$ kHz and $Q_{\text{eff}} \approx 10^9$. The laser powers P_B and P_R refer to the circulating powers, and Q refers to the quality of the bare mechanical system. For the specified geometry, the required finesses are of order $\mathcal{F} \approx 10,000$, compatible with the optical quality of resonators in production today.

or excess damping². Additionally, the effect of the feedback discussed above is to suppress the damping contribution from both springs, causing any mismatching of the damping cancellation to be further suppressed. In practice, it may not be trivial to set different bandwidths for two optical fields of macroscopically similar frequency. In this case (i.e., $\gamma_B = \gamma_R \equiv \gamma$), one can still cancel the imaginary terms by choosing the appropriate detunings. In particular, if $|\omega_{osB}/\omega_{osA}| = \kappa$, cancellation is obtained for $\Delta_B^2 = (\kappa^2 - 1)\gamma^2 + \kappa^2\Delta_R^2$.

A set of sample parameters is given in Table 7.1. Under these conditions, an oscillator with a resonant frequency of $\omega_{os}/2\pi \approx 100$ kHz and an effective Q of 10^9 is formed³. Such a device can in principle be cooled to its ground state from an environmental temperature of $T \approx 4800$ K (clearly, this should not be attempted, but it serves to illustrate what this technique implies in the context of quantum experiments)! From Eq. (7.20), we can also calculate the effective temperatures of the residual quantum radiation pressure noise from the two optical spring fields as $T_{\text{eff}}^{\text{res,B}} = T_{\text{eff}}^{\text{res,R}} \approx 23$ mK, in the lossless case, or $T_{\text{eff}}^{\text{res,B}} \approx 33$ K and $T_{\text{eff}}^{\text{res,R}} \approx 21$ K, for 99.5% quantum efficiency and $\epsilon = 10$ ppm. Even in the lossy case, the residual noise temperatures are considerably lower than most target environment temperatures.

7.6 Conclusion

We have proposed a method for creating a tunable effective mechanical resonator with extremely high Qf product. In addition, these resonators can be made to

²Note that the expression need not vanish, but only be positive for the resultant resonator to be stable. Furthermore, any positive damping from the optical fields is cold, and therefore does not contribute noise or degrade SNR. We specifically consider the case of zero additional damping, however, since it leads to an effective resonator whose Q is determined solely by the intrinsic damping of the bare mechanical system.

³This Q_{eff} value is calculated assuming a viscous damping model; the mechanical damping, γ_m , is fixed, and so, since the optical spring adds no damping, the improvement is given by $Q_{\text{eff}} = (\omega_{os}/\omega_m)Q$. Several candidate mechanical resonators are predicted to be better approximated by a structural damping model, in which case the improvement in Q is potentially much greater.

operate in lower frequency bands than current ones of competitive quality, allowing for exceptionally long rethermalization timescales. While the use of optical dilution to mitigate thermal noise has been proposed and demonstrated in the past, we have considered a parameter regime in which the deleterious effects of quantum radiation pressure noise from the strong optical spring fields can be all but eliminated, allowing for greatly unhindered dilution. We feel that the application of this technique holds great promise for any field requiring very-high-Q resonators, including, but not limited to, those of quantum optomechanics and sensitive force measurement.

Part V

Thermal Noise in Silicon Ribbons

The LIGO Voyager suspensions and test masses will be made of monocrystalline silicon [51]. Open questions regarding this design choice include how silicon suspensions can be fabricated, and whether they can be made to exhibit low enough thermal noise. This part of this thesis details the development of an experiment intended to answer these questions.

Much research has been done into the mechanical quality of silicon [150–152] within the context of its future use in GW detectors. However, this research has been limited to measurements of the dissipation of samples' resonant modes, and the off-resonance noise behavior is then inferred from these measurements using the fluctuation-dissipation theorem (FDT). The goal of this experiment is to demonstrate the sufficiently low off-resonance thermal noise of silicon ribbon structures by direct interferometric interrogation.

The organization of this part is as follows. Chapter 8 contains a more detailed treatment of thermal noise, followed by a motivation of the resonator design based on thermal noise estimates for LIGO Voyager, and finally a description of a numerical model used to compute thermal noise and other mechanical properties of the resonators. In Chapter 9, an overview of the fabrication process used to create the devices is given. Chapter 10 concerns the apparatus and techniques used to do the on-resonance quality factor measurements similar to those performed in the previous work cited above (this “ringdown” analysis was used in order to corroborate the off-resonance measurements, and as a straightforward method of validating the fabrication methods). Finally, Chapter 11 details the progress towards an off-resonance interferometric measurement of thermal noise in the fabricated samples.

THERMAL NOISE THEORY, RESONATOR DESIGN, AND NUMERICAL MODELING

This chapter is divided into three sections. In Sec. 8.1, a more detailed description of thermal noise is given, drawing from and expanding the simple description found in Sec. 2.2.7. In Sec. 8.2, these principles are applied to the proposed LIGO Voyager design to determine the required thermal noise performance of the suspension, and these results motivate a resonator design for the prototype experiment described here. Finally, in Sec. 8.3, a numerical model used to compute mechanical transfer functions and thermal noise estimates for the prototype experiment resonators is presented.

8.1 Thermal noise in detail

8.1.1 The loss angle




Figure 8.1: A simple mass-spring system with damping.

Consider a simple mass-spring system like the one shown in Fig. 8.1. If the spring is massless, with spring constant k and velocity damping factor b , the equation of motion of the system is

$$m\ddot{x} = -kx - b\dot{x} + F, \quad (8.1)$$

where F is an externally applied force in the $+\hat{x}$ direction. In many cases, the

internal damping of materials has been found to be of the form [132]

$$b = \frac{k\phi}{\omega}, \quad (8.2)$$

so that the equation of motion can be written in the form

$$m\ddot{x} = -k(1 + i\phi)x + F. \quad (8.3)$$

That is, the damping is expressed as an imaginary component to the spring constant. If $\phi \ll 1$, the homogeneous solution describing the free evolution of this system (i.e., for $F = 0$) is

$$x(t) = e^{-\frac{\phi\omega_0 t}{2}} (A \sin \omega_0 t + B \cos \omega_0 t), \quad (8.4)$$

where $\omega_0 \equiv \sqrt{\frac{k}{m}}$ is system's resonant frequency. This describes sinusoidal motion with an envelope that decays exponentially with a lifetime $\tau = \frac{2}{\omega_0\phi} = (\pi f_0\phi)^{-1}$, where $\omega_0 = 2\pi f_0$. Thus, ϕ quantifies the rate of decay of the system.

Now, if the system starts at rest at $t = 0$ with an amplitude x_0 , it has an initial energy

$$U_0 = \frac{1}{2}kx_0^2. \quad (8.5)$$

One cycle later, at $t = T \equiv f_0^{-1}$, its amplitude is

$$x(T) = x_0 e^{-\pi\phi} \approx x_0(1 - \pi\phi), \quad (8.6)$$

giving a new stored energy

$$U_1 = \frac{1}{2}kx_0^2(1 - \pi\phi)^2 \approx \frac{1}{2}kx_0^2(1 - 2\pi\phi). \quad (8.7)$$

If we measure the fractional energy lost in this cycle, we find

$$\frac{\Delta U}{U_0} = \frac{U_0 - U_1}{U_0} = 2\pi\phi. \quad (8.8)$$

So, ϕ is the fractional energy lost by the system per radian. For this reason, ϕ is known as the **loss factor** or **loss angle**. Another useful quantity is its reciprocal,

$$Q \equiv \frac{1}{\phi} = \pi f_0 \tau, \quad (8.9)$$

which is called the **quality factor**.

8.1.2 The fluctuation-dissipation theorem (again)

In Sec. 2.2.7, we introduced the fluctuation-dissipation theorem (FDT) and found that the thermal noise force power spectral density of a system was related to its mechanical resistance:

$$S_F(\omega) = 4k_B T \operatorname{Re}\{Z(\omega)\}, \quad (8.10)$$

where $Z(\omega) \equiv \frac{F(\omega)}{v(\omega)}$. Equivalently, we found that we could compute the system's thermal displacement noise power spectral density as

$$S_x(\omega) = \frac{4k_B T \operatorname{Re}\{Y(\omega)\}}{\omega^2}, \quad (8.11)$$

with $Y(\omega) \equiv Z(\omega)^{-1}$.

Taking the Fourier transform of (8.2) and rearranging gives

$$Z(\omega) = \frac{\tilde{F}}{i\omega\tilde{x}} = im\omega - \frac{ik}{\omega} + \frac{k\phi}{\omega}. \quad (8.12)$$

Taking the reciprocal,

$$Y(\omega) = \frac{\omega k\phi + i(\omega k - m\omega^3)}{(k - m\omega^2)^2 + k^2\phi^2}. \quad (8.13)$$

Plugging the above into (8.11), we get

$$S_x(\omega) = \frac{4k_B T k\phi}{\omega [(k - m\omega^2)^2 + k^2\phi^2]}. \quad (8.14)$$

8.1.3 Sources of loss

While it is possible for the loss angle of a material to be relatively constant over a wide range of situations, it typically varies considerably with frequency and temperature:

$$\phi = \phi(\omega, T). \quad (8.15)$$

Also, the total loss angle is in reality the sum of contributions from a number of independent loss mechanisms. The well-known loss mechanisms that apply to monocrystalline silicon resonators are outlined below.

8.1.3.1 Thermoelastic (Zener) loss

Zener [137, 153] described a loss mechanism in flexures due to the effect illustrated in Fig. 8.2. On the left, two adjacent domains of a volume of material are in equilibrium at a temperature T . On the right, the system has been bent¹ so that the

¹Imagine that the system here is a short section of a long bar extending far to the left and right off the page. The right panel is what that section looks like if the ends of this bar are bent far upward.




Figure 8.2: Illustration of the thermoelastic loss mechanism.

upper volume is compressed, while the lower volume is expanded. This causes an increase in the temperature of the upper section and a decrease in the temperature of the lower section. This temperature gradient leads to a heat flow and an increase in entropy, which is always associated with dissipation.

The loss factor contribution from thermoelastic loss is

$$\phi_{\text{TE}}(\omega, T) = \frac{\alpha^2 Y T}{\rho c_p} \frac{\omega \tau}{1 + \omega^2 \tau^2}, \quad (8.16)$$

where α is the coefficient of thermal expansion (CTE), Y is the Young's modulus, ρ is the density, c_p is the specific heat and

$$\tau = \frac{\rho c_p d^2}{\pi^2 \kappa} \quad (8.17)$$

is a characteristic time, with d the sample thickness and κ its thermal conductivity.

Thermoelastic loss is the dominant contributor in silicon at room temperature, and the α^2 dependence is the reason why LIGO Voyager will operate near 123 K, where $\alpha = 0$. However, in the case of a suspension under stress due to the mass it supports, an additional term to (8.16) must be considered. Specifically, in this case, α is no longer the only mechanism by which a change in temperature can induce a change in the volume of the system: if the Young's modulus, Y , is a function of temperature, then a change in temperature will induce a change in strain, and thereby a change in volume. The loss thus becomes

$$\phi_{\text{TE}}(\omega, T) = \frac{Y T}{\rho c_p} \frac{\omega \tau}{1 + \omega^2 \tau^2} \left(\alpha - \beta \frac{mg}{YA} \right)^2, \quad (8.18)$$

where m is the suspended mass, A is the cross-sectional area of the suspension fiber/ribbon, and

$$\beta \equiv \frac{1}{Y} \frac{dY}{dT} \quad (8.19)$$

is the fractional change in the Young's modulus per unit temperature.

In aLIGO, the area of the fused silica suspension fibers was chosen such that the β term in the equation above exactly cancelled the α term (possible since β is positive for fused silica). In LIGO Voyager, the temperature will have to be biased slightly away from the $\alpha = 0$ point to achieve a similar cancellation. This does not spoil Voyager's relative immunity to thermal distortions when compared to aLIGO because 1) the magnitude of this α bias in the case of the proposed silicon suspensions is lower than the α of fused silica, and 2) the much larger thermal conductivity of silicon compared to fused silica results in much smaller thermal gradients to begin with, for a given level of absorption.

8.1.3.2 Phonon-phonon (Akhiezer) loss

Another source of dissipation, attributed to Akhiezer [138, 154], involves the distribution of phonons within a solid. When a solid in thermal equilibrium has a strain applied with a wavelength much longer than the phonon wavelength, the equilibrium is disturbed, causing a reorganization of the phonon distribution to a *new* equilibrium. This reorganization requires an increase in entropy, and therefore dissipation.

The contribution from this effect is [151]

$$\phi_{\text{ph-ph}}(\omega, T) = \frac{c_p T \gamma^2}{v^2} \frac{\omega \tau_{\text{ph}}}{1 + \omega^2 \tau_{\text{ph}}^2}, \quad (8.20)$$

where γ is the ‘‘Grüneisen parameter’’ (of order unity), v is the speed of sound, and

$$\tau_{\text{ph}} = \frac{3\kappa}{c_p v^2} \quad (8.21)$$

is the phonon relaxation time.

Phonon-phonon loss is the dominant intrinsic mechanism in regimes where thermoelastic loss is minimal, such as near silicon's CTE zeros and in thick samples. In particular, it has been a consideration for bar GW detectors [155]. More recently, phonon-phonon loss has been measured directly in MEMS devices [156], where it is believed to be the ultimately limiting contributor.

8.1.3.3 Surface loss

The preceding two mechanisms are bulk effects and are dominant when the sample has a relatively low surface-to-volume ratio. On the other hand, losses at and near the surface of a sample can be far higher than in the bulk due to defects, undesired film growth, etc. As the surface-to-volume ratio of a sample increases, the relative contribution of these surface effects to the overall dissipation increases, as well.

If the loss factor in these surface regions—as well as their physical extent—is known, the loss of the overall system can be computed by integrating the volume average over all domains, weighted by the strain energy in the corresponding regions:

$$\phi_{\text{tot}} = \frac{\phi_0 \int_{\text{bulk}} W_s(\mathbf{r}) dV + \phi_s \int_{\text{surface}} W_s(\mathbf{r}) dV}{\int_{\text{all}} W_s(\mathbf{r}) dV}, \quad (8.22)$$

where ϕ_0 is the bulk loss, ϕ_s is the loss in the surface regions, and W_s is the strain energy density in J/m^3 .

8.1.3.4 Clamping/recoil loss

A resonator must be attached to *something*. Typically, if resonance with low dissipation (“high-Q”) is desired, the resonator is attached to something very stiff and heavy, so that the attached-to object does not react appreciably to the motion of the resonator. In practice, this scheme never works perfectly, and the object always moves somewhat along with the resonator. This leads to two qualitatively different but mathematically identical sources of loss: 1) the motion of the resonator at the attachment point may lead to local elastic strains in the object (**clamping loss**), and 2) the motion of the resonator may cause rigid-body reactive motion in the object (**recoil loss**).

These effects lead to extra dissipation because the loss angle of the attached-to object—or of whatever *it* is attached to, in the case of recoil—is typically higher than that of the resonator. Therefore, whatever energy is stored in the elastic motion of anything but the resonator is lost more rapidly.

The effect of these losses can be computed by integrating over domains as in (8.22) for surface loss.

8.2 Suspension thermal noise in LIGO Voyager

In this section, we will use the result in Sec. 8.1.2 to estimate the loss requirement for the LIGO Voyager suspensions. The analysis here is adapted from an earlier

technical note [157], which itself draws on estimates from an early LIGO Voyager suspension design study [158]².

8.2.1 Background

As discussed previously, whereas the Advanced LIGO suspensions were formed of heat-drawn fused silica fibers, the LIGO Voyager suspensions will be “ribbons” made of monocryalline silicon. This presents challenges for fabrication, but also affords some potential benefits. In particular, while the aLIGO suspensions were constrained to having a circular cross section, the LIGO Voyager suspensions can be made to be rectangular in cross section with an arbitrary aspect ratio. This aspect ratio tuning can have strong thermal noise advantages [159].

8.2.2 The ribbon cross section

The required cross-sectional area of the fiber can be determined by dividing the weight of the test mass by an appropriate stress for the fiber material:

$$A_{\text{fiber}} = \frac{Mg}{\sigma_{\text{fiber}}}. \quad (8.23)$$

The suspension optimization note gives $M = 143 \text{ kg}$ and $\sigma_{\text{fiber}} = 1.4 \text{ GPa}$ ³, giving $A_{\text{fiber}} = 1.00 \times 10^{-6} \text{ m}^2$. Assuming that—as in aLIGO—there will be four ribbons connecting the two stages, we can create a structure with cross-sectional area $A_{\text{fiber}}/4$, but with the aspect ratio of our choosing. Three sample cross-sections are shown in Fig. 8.3. We will use these three examples in calculating the thermal noise requirement.




Figure 8.3: Possible ribbon cross sections, each with area $A_{\text{fiber}}/4$.

8.2.3 Suspension thermal noise

As we saw found above, the thermal noise displacement power spectrum is

$$S_x = \frac{4k_B T}{\omega^2} \text{Re}\{Y\}, \quad (8.24)$$

²Certain design parameters have changed slightly since the time of this initial calculation, but not drastically enough that the result is invalidated.

³It is noted in the optimization note that there is still some uncertainty as to what the stress should be.

where Y is the admittance function of the system. In the case of a pendulum with some elasticity, we must modify the admittance from (8.13) to include the gravitational restoring force in addition to the elastic one [160]:

$$\operatorname{Re}\{Y\} = \frac{\omega (K_g \phi_g + K_e \phi_e)}{\left[(K_g + K_e) - M\omega^2 \right]^2 + (K_g + K_e)^2 \phi^2}, \quad (8.25)$$

where M is the suspended mass, ϕ is the total loss factor, and $K_{g,e}$, $\phi_{g,e}$ are the effective spring constant and loss factor contributions from gravity and elasticity, respectively. In the limit that gravity dominates the restoring force ($K_g \gg K_e$) and that gravity is lossless ($\phi_g = 0$), this reduces to

$$\operatorname{Re}\{Y\} \approx \frac{\omega K_e \phi}{(K_g - M\omega^2)^2 + K_g^2 \phi^2}, \quad (8.26)$$

where we have identified $\phi = \phi_e$.

As a short aside: by plugging the high-frequency limit of Eq. 8.26 into Eq. 8.24, we obtain

$$S_x \approx \frac{4k_B T K_e \phi}{M^2 \omega^5} = \frac{4k_B T \omega_0^2 (K_e/K_g) \phi}{M \omega^5}, \quad (8.27)$$

where $\omega_0 \equiv \sqrt{K_g/M}$ is the resonant frequency of the corresponding ideal pendulum. This result is equivalent to that for an oscillator with complex spring constant $K = K_g(1 + i(K_e/K_g)\phi)$, and so the factor $K_e/K_g \ll 1$ —often called the “dilution factor”—quantifies how the influence of the loss in the suspension has been reduced by the lossless gravitational restoring force.

Returning to Eq. 8.26, in order to evaluate the thermal noise spectrum, we must find the effective spring constants. Limiting our consideration to the low-frequency region below the suspension violin modes, and in the regime of strong dilution [160], these can be separated as

$$K_g = \frac{Mg}{L} \quad (8.28)$$

$$K_e = \frac{Mg\Delta}{L^2}, \quad (8.29)$$

where the characteristic bending length of the fiber, Δ , is

$$\Delta = \sqrt{\frac{EI}{F_T}}, \quad (8.30)$$

with E the suspension material Young’s modulus, I the area moment of inertia of the fiber cross section⁴, and F_T the tension in a single fiber. In our case with four fibers, $F_T = Mg/4$, and we arrive at the elastic spring constant $K_e = \sqrt{4EIM}g/L^2$.

⁴ $I = bh^3/12$, with b the thicker and h the thinner dimension of the cross section.

8.2.3.1 Bulk loss

Parameter	Value	Unit
T	124	K
M	143	kg
L	1.0	m
E	167.4	GPa
F_T	350.3	N
$I_{1:1}$	5.22×10^{-15}	m^4
$I_{5:1}$	1.74×10^{-15}	m^4
$I_{10:1}$	5.22×10^{-16}	m^4
ϕ_0	10^{-8}	1

Table 8.1: Parameters used for the thermal noise calculation.




Figure 8.4: Projection of bulk suspension thermal noise for each of the three aspect ratios, alongside other projected LIGO Voyager noise contributions and the expected total.

Using the expressions above and the parameters listed in Table 8.1, it is straightforward to evaluate the thermal noise that results from the bulk loss, ϕ_0 , of the

suspension material alone⁵. A projection of this noise for each aspect ratio case is plotted alongside other expected LIGO Voyager noise contributions in Fig. 8.4. As can be seen, the bulk loss contribution is well below the Voyager noise floor throughout.

8.2.3.2 Surface loss

As described in Sec. 8.1.3.3, the overall loss of the suspension—and therefore the noise—will be dominated by a thin but much-higher-loss surface layer surrounding the bulk. Using Fig. 8.4 and noting that the noise far from resonance scales with $\sqrt{\phi}$, we can determine an acceptable total loss in each aspect ratio case. The loss required to bring suspension thermal noise to 1/2 the level of quantum noise and seismic noise at 10 Hz in each case is given in Table 8.2.

Aspect ratio	1:1	5:1	10:1
ϕ	4.3×10^{-7}	7.5×10^{-7}	1.37×10^{-6}

Table 8.2: Acceptable total loss for each aspect ratio.

For various reasons, it will be useful to parameterize this acceptable loss in terms of two more easily characterized quantities: the loss ϕ_s and thickness d of the surface layer itself. To do this, we must construct the appropriate map from these quantities to the total loss. Specifically, the volume average of the two contributing losses (i.e., ϕ_0 and ϕ_s) must be weighted by the strain energy in the corresponding domains as in (8.22).

We can very easily look for the correct dependence of W_s on position using a finite-element analysis tool like COMSOL⁶. The results of such an analysis can be found in Fig. 8.5. Here, a ribbon of each aspect ratio is clamped at one end, while the other end is free. A modal analysis is run, and $W_s(\mathbf{r})$ is computed in three dimensions for the fundamental mode. Finally, an evaluation curve is defined along the bending dimension of the ribbon at three different locations down the length of the ribbon (namely, at 10%, 20% and 30% of the total length). As can be seen, the dependence is quadratic. The energy density also varies vertically along the fiber, as shown—it is highest at the clamped section—but this does not affect the surface loss consideration.

⁵Bulk loss in silicon has been measured to be as low as below 10^{-9} in some low-temperature experiments [161]. We use 10^{-8} here for margin.

⁶This tool will be used more extensively in the next section.




Figure 8.5: Normalized strain energy density along the bending direction of the ribbon, at three vertical locations along its length (10%, 20% and 30% of the total length) for each aspect ratio. To demonstrate the dependence, a quadratic fit to a single trace is also shown.

Using the quadratic dependence confirmed above, it is now straightforward to calculate the total loss of the suspension as a function of ϕ_s and d . This is plotted for the extremal (1:1 and 10:1) cases in Fig. 8.6, with the acceptable losses from Table 8.2 highlighted in green.

8.2.4 Conclusion

It appears from the foregoing analysis that, using an aspect ratio in the range of 10:1, tolerable suspension thermal noise in the LIGO Voyager interferometer requires the total loss to be on the order of 10^{-6} . Therefore, it is desired to create a prototype system with reasonably similar dimensions that has a loss this low or better. The remainder of this part of the thesis will detail the design and measurement of such a device.




Figure 8.6: Total loss factor as a function of surface layer loss and thickness, for the two extremal aspect ratios. In each case, the acceptable total loss contour is highlighted in green.

8.3 Ribbon cantilever design and numerical modeling

This section describes the design of the silicon cantilever resonator designed to simulate a LIGO Voyager-type suspension ribbon, then details the numerical model used to compute mechanical transfer functions and thermal noise.

8.3.1 Design

The final design of the silicon cantilever resonators used in this experiment is shown in Fig. 8.7. The device comprises two pieces, both made of silicon: 1) the main cantilever section, and 2) a small mirror attached to the end to serve as one mirror in a Fabry-Pérot cavity used to make an interferometric measurement of the cantilever's displacement noise. The mirrors will be discussed further in Chapter 11.

The main cantilever is monolithic, fabricated out of a silicon wafer as detailed in Chapter 9. At each end of the cantilever is a 1 cm x 1 cm section that retains the full 500- μm thickness of the wafer. One of these sections has a hole etched through it so that a laser beam can pass uninhibited to the mirror. The other is solid, and it is this section of the cantilever that is clamped in place. The central region—5 cm long in the final design—is thinned to around 250 μm . This thinning is done to reduce the fundamental resonance frequency of the cantilever, as well as to reduce clamping losses by keeping the bending from the cantilever motion away from the actual clamping regions.




Figure 8.7: Design of the silicon cantilever.

It should be noted that, contrary to the foregoing case in the previous section for long ribbons supporting a heavy mass, the restoring force on these devices is dominated by the elasticity, not gravity. In fact, the main experiment is carried out with the cantilever clamped at the *bottom*, extending vertically with the mirror ends at the *top*. This means that there is no reduction in thermal noise from gravitational dilution, but that is not important in this case: since we are interested in characterizing this noise, maximal coupling is actually preferred. Additionally, these devices are considerably shorter (7 cm vs 1 m) and wider (1 cm vs 1.6 mm) than the proposed LIGO Voyager ribbons above.

On the other hand, due to the similar thickness dimensions, the surface-to-volume ratio is also similar in both cases. To see this, consider that for a rectangular prism of length L , width W , and thickness T , the surface-to-volume ratio is

$$\frac{S}{V} = \frac{2 \cdot (LW + LT + WT)}{LWT} \approx \frac{2}{T} \quad (8.31)$$

when $L \gg W \gg T$. Additionally, though the cantilever's resonant frequency is higher than the ~ 1 -Hz pendulum resonance frequency above, it is still low enough that we can investigate the off-resonance thermal noise in the region from 100 Hz - 1 kHz that is pertinent to GW detectors.

8.3.2 Finite-element analysis (FEA) numerical model




Figure 8.8: The first few cantilever resonant mode shapes, calculated with the finite-element analysis software tool COMSOL.

The finite-element analysis (FEA) software tool COMSOL was used to numerically compute the mechanics of the cantilever. To do this, a geometrical model is built using the CAD interface within COMSOL, and then materials are assigned to each domain (all domains are silicon for this system). Then, certain physical constraints are imposed to set the appropriate boundary conditions for the numerical solver. In this case, it was sufficient to specify that the faces of the clamped section of the cantilever were to be fixed in space. Finally, a modal solver is executed in order to compute the eigenfrequencies and spatial shapes of the system's normal modes. The first four such modes are illustrated in Fig. 8.8.

The eventual, interferometric measurement will strictly measure the displacement of the cantilever system at the location of the beam spot on the mirror, which is of course a superposition of the displacements of all these normal modes. Therefore, we desire a method of adapting the analysis in Sec. 8.1.2 for a single-mode, discrete system into one that can describe multiple modes of a continuous system.

Generalizing the analysis to multiple modes is fairly straightforward; instead of a single admittance function, $Y(\omega)$, we instead have a series of admittance functions, $Y_n(\omega)$, one for each mode. The full system is described by the sum of these

admittances:

$$Y_{\text{tot}}(\omega) = \sum_{n=1}^{\infty} Y_n(\omega). \quad (8.32)$$

This would already be sufficient to describe multi-mode, discrete systems (e.g., multiple masses connected by massless springs, multi-stage pendulums with massless rods, etc.), but we are considering a continuous system whose generalized displacements are actually 3-D deformation fields. The displacement for any given mode can be described by the vector field

$$\mathbf{D}_n(t, x, y, z) = A_n(t)\mathbf{u}_n(x, y, z), \quad (8.33)$$

where $A_n(t)$ is the generalized amplitude of the mode n at time t and $\mathbf{u}_n(x, y, z)$ is its normalized displacement mode shape. Ultimately, however, we only need to consider the 1-dimensional motion of the aforementioned point on the mirror along the optical cavity axis, so we need a way to reduce this system. This is accomplished using a concept called the **modal mass**: by choosing a single point in the system, as well as a single dimension of motion, we can perform a volume integration of the normal mode shapes and, comparing it with the linear motion of the chosen point, we can convert this into an effective point mass for each mode. Therefore, this transforms our system into a multi-mode, effectively discrete, 1-dimensional problem. The modal mass at some point $\mathcal{P} = (x_{\mathcal{P}}, y_{\mathcal{P}}, z_{\mathcal{P}})$ and along the \hat{x} direction is equal to

$$M_n^{\mathcal{P}x} = \frac{\rho \int_V |\mathbf{u}_n|^2 dV}{[\mathbf{u}_n(\mathcal{P}) \cdot \hat{x}]^2}, \quad (8.34)$$

where ρ is the mass density and the integration is carried out over all moving domains.

Using this modal mass formalism, we can use (8.13) and (8.32) to compute an effective total admittance

$$Y_{\text{tot}}^{\mathcal{P}x}(\omega) = \sum_{n=1}^{\infty} Y_n(M_n^{\mathcal{P}x}, \omega). \quad (8.35)$$

Finally, the FDT can be applied to calculate the expected thermal displacement noise of the system at the point \mathcal{P} in the \hat{x} direction:

$$S_x^{\mathcal{P}x}(\omega) = \frac{4k_B T}{\omega^2} \text{Re}\{Y_{\text{tot}}^{\mathcal{P}x}(\omega)\}. \quad (8.36)$$

The results of this analysis for the lowest four modes of the cantilever are presented in Figs. 8.9 and 8.10. The transfer function plot is showing the mechanical response

from force to *displacement* (i.e., $\frac{Y_{tot}^x(\omega)}{i\omega}$), as this is often a more useful quantity to work with directly. As one can see, the response is almost completely dominated by the contribution from the fundamental mode (in blue, almost hidden), save for in very narrow bands near the resonances of the other modes. Note also that Mode 3 is all but completely absent. This makes sense if one examines the shape of that mode in Fig. 8.8: this is a torsional mode that has no first-order coupling to the chosen degree of freedom.

The estimated thermal noise displacement spectrum is plotted for two different values of ϕ , to demonstrate that the displacement noise far from resonance is proportional to $\sqrt{\phi}$. Note that there is no apparent contribution from Mode 3, as expected.




Figure 8.9: Cantilever mechanical transfer function calculated via FEA analysis, showing the individual contributions from each of the four lowest modes. In each case, the force is applied and the displacement is read out at the location of the beam spot on the center of the cantilever's mirror.




Figure 8.10: The cantilever thermal displacement noise calculated via FEA analysis, for $\phi = 10^{-8}$ and 10^{-6} .

*Chapter 9***CANTILEVER FABRICATION****9.1 Fabrication process****9.1.1 Overview**

The recipe used to create the cantilevers was adapted from a method used by the Chao group in Taiwan [162], which is based on the standard technique known as **hard mask etching**¹. The Chao group design was a more traditional cantilever resonator with a single thick end, used primarily to measure the loss of coatings applied onto them.

The initial step was to verify that these simpler cantilevers could be replicated using the shared process. The process was followed closely and simple cantilevers were generated and then evaluated as described in Chapter 10. Subsequently, work began on the development of a process that could create the more specialized resonators required for this experiment (as described in Sec. 8.3 above).

Many trials were conducted in the development of this process prior to settling on the ultimate recipe. This development was complicated in two primary ways: 1) the desire to have a hole through one end necessitates a multi-round process, each round requiring several steps, and 2) the need for the hole, as well as the size of the device, require special considerations to be taken throughout the processing to prevent contamination or damage to either side of the wafer, in order to maintain the devices' mechanical quality.

9.1.2 Detailed recipe

The final fabrication recipe is shown in the diagram in Fig. 9.1. The sections below describe the recipe in detail, step-by-step. All fabrication was performed at the Kavli Nanoscience Institute (KNI)² at Caltech.

¹This refers to the use of a “hard mask” material—such as Si_3N_4 or SiO_2 —to define the ultimate etch pattern, as opposed to a “soft” material, like an organic photoresist.

²The Kavli Nanoscience Institute, 1200 E California Blvd, MC 107-81, Steele Laboratory, Pasadena, CA 91125, USA.

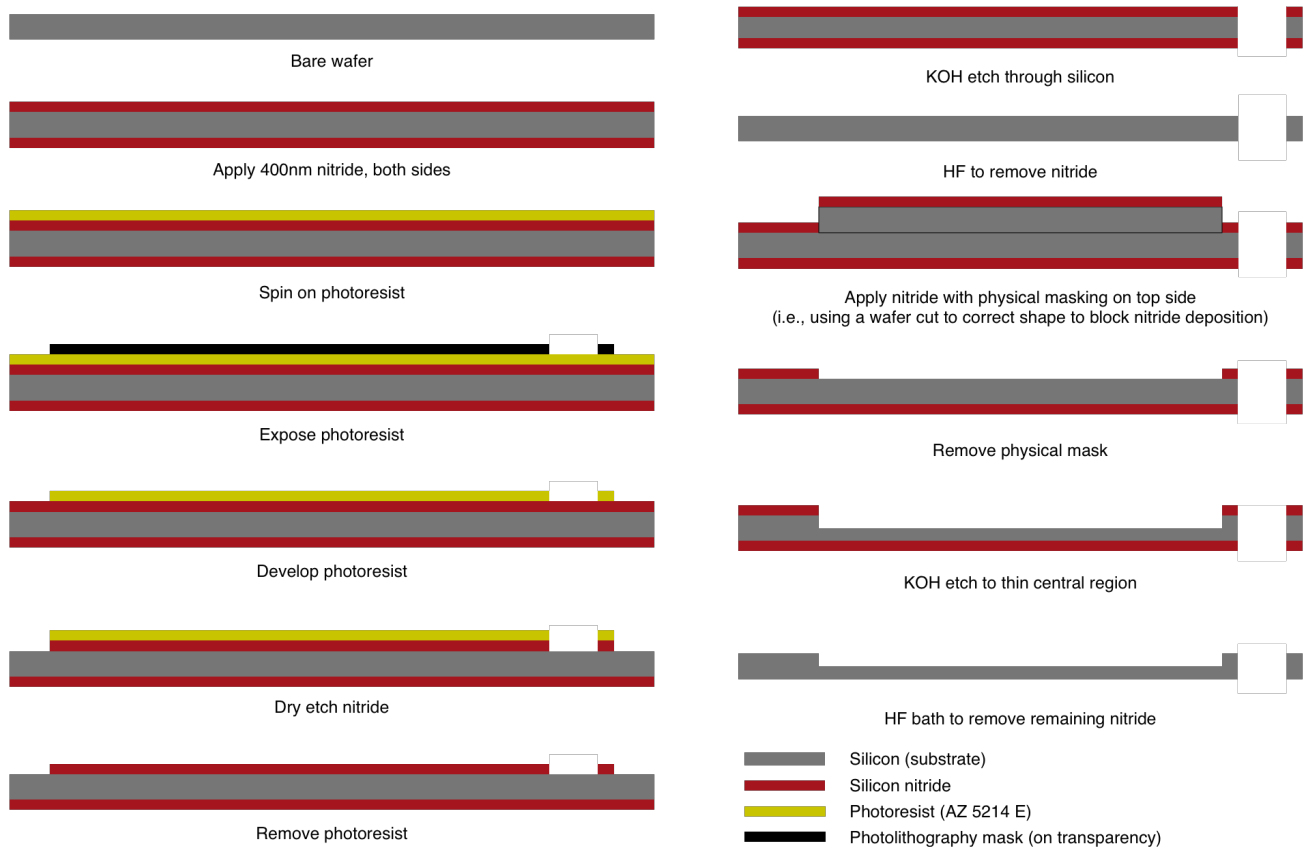


Figure 9.1: The cantilever fabrication process, starting with a bare silicon wafer at the top left and ending with a thinned cantilever with a hole in it at the bottom right. (The depiction is a cross section along the long dimension of the device, hence the appearance of the hole as a gap here.)

9.1.2.1 Wafer stock

All devices were fabricated from 100-mm diameter undoped silicon wafers with a thickness of $500\ \mu\text{m}$, in the $\langle 100 \rangle$ crystal orientation. At various times, both single-side-polished (SSP) and double-side-polished (DSP) wafers were used. The choice between SSP and DSP was not observed to affect the measured mechanical quality at the levels attained in this experiment.

9.1.2.2 Wafer cleaning

Prior to processing, all wafers are cleaned using a standard wafer cleaning process [163]. The first step is a solvent clean, where the wafers are soaked in warm (55°) acetone for ten minutes and then in room-temperature methanol for five minutes. The wafers are then rinsed in de-ionized (DI) water and blown dry.


Following the solvent clean, a so-called “RCA clean”³ is performed, using a combination of H₂O, NH₄OH and H₂O₂, in a roughly 5:1:1 ratio, at around 70°C for fifteen minutes. The wafers are then rinsed and dried again.

Finally, the wafers are soaked in a bath of dilute hydrofluoric (HF) acid for two minutes. Following a final DI rinse and dry, the wafers are ready for processing.

9.1.2.3 Silicon nitride hard mask deposition (I)

The first step in the process is to apply a roughly 400-nm-thick layer of silicon nitride (Si₃N₄) uniformly to both sides of the wafer. This is accomplished via **plasma-enhanced chemical vapor deposition (PECVD)** using an Oxford Instruments System 100 PECVD unit.

In this process, a radio-frequency-induced plasma is generated above the substrate with a particular mixture of gases, and the ensuing reaction causes a film layer to be grown on the substrate surface. The chemical equation for the deposition of the silicon nitride layer is:



This layer serves as the hard mask that will itself be patterned in order to determine the final etch pattern on the silicon wafer. A photo showing the wafer before and after the deposition is shown in Fig. 9.2.

9.1.2.4 Photolithography

The next step is to pattern the hard nitride mask. This is accomplished via **photolithography**.

First, a **photoresist** material (AZ 5214 E) is applied evenly over one surface of the wafer to a thickness of about 1.5 μm using a spinner⁴ running at around 4000 rpm for one minute. Then, the wafer is immediately placed on a hotplate at 110°C for 50 seconds to bake the photoresist.

³This process gets its name from the fact that it was developed by an employee of the Radio Corporation of America in the 1960s.

⁴A spinner is a device that retains a wafer and then spins it at a prescribed rotation rate for a certain period of time. During the spinning, a substance is pipetted onto the center of the wafer, and the centrifugal force causes the substance to be spread to a predetermined, mostly uniform thickness across the wafer. The wafer is retained by a small amount of suction provided by a vacuum pump applied to the center underside of the wafer by the spinner.




Figure 9.2: Wafer (a) before and (b) after deposition of a 400-nm-thick layer of silicon nitride using PECVD. The 100-mm wafer under process is resting on a 6-inch sacrificial carrier wafer, and the small rectangular pieces are there to assist in keeping the main wafer from slipping out of position during the chamber pump-down to vacuum.

The next step in the photolithography process is the exposure. In this step, the photoresist-clad wafer receives a controlled exposure of ultraviolet (UV) light, which is blocked in certain places as desired by a photolithography mask. These masks are commonly made of glass slides with a very-high-resolution pattern of chromium on one surface; in many applications not requiring extreme resolution—such as ours—this mask can be made of a transparency printed with high-opacity ink, taped to a glass slide. Both the wafer and the mask are loaded into a Suss Microtech MA6/BA6 contact aligner (see Fig. 9.3), which contains multi-axis actuators to ensure proper alignment of the desired pattern to the wafer orientation. When the desired alignment is achieved, the unit also performs the exposure, which takes on the order of ten seconds.

The mask pattern for this step is an array of rectangles in the shape of the eventual cantilever outlines, with circular holes at one end, as shown in Fig. 9.4.

Following exposure, the photoresist must be developed using an appropriate devel-




Figure 9.3: Demonstration of the operation of the contact aligner for photolithography: (a) The photomask, mounted on a retaining plate, being slid into the machine. (b) The photomask in position for exposure. (c) The wafer on the wafer drawer. For exposure, the wafer drawer is slid into the machine under the photomask plate. Then, an automated system puts the two in close contact, and the user is allowed to adjust their relative position and orientation before the illumination.

oper solution. In our case, the wafers are bathed in MF CD-26 developer for around one minute. The dissolving of the photoresist in the exposed areas can be seen by eye (as shown in Fig. 9.5(a)), so it is straightforward to know when the development is complete. The wafer is then rinsed in DI and blown dry.

9.1.2.5 Dry etching of silicon nitride

At this point, the silicon nitride layer is protected by photoresist in the regions that were not exposed and developed away. The next step is to transfer this pattern to the nitride mask via dry etching.

The dry etch is also a plasma process, performed on an Oxford Instruments System 100 ICP 380 etcher. Several plasma etches exist for the removal of silicon nitride; the one chosen based on expert recommendation was the so-called “pseudo Bosch etch” [164, 165], which uses sulfur hexafluoride (SF_6) and octafluorocyclobutane (C_4F_8) to produce a very vertical etch profile by passivating the vertical surfaces as the etch progresses. This process takes on the order of 5 minutes to remove the






Figure 9.4: Mask for the photolithography process (actual size). This pattern is printed with high resolution using high-opacity ink on a transparency, and the transparency is attached to a glass slide for insertion in to the contact aligner.



(a)



(b)

Figure 9.5: Demonstration (a) of the appearance of the photolithography pattern during photoresist development and (b) of its transfer to the silicon nitride hard mask layer during dry etching.

400-nm nitride layer in the desired regions. The appearance of the wafer at the completion of the etch is shown in Fig. 9.5(b).

Following completion of the dry etch, the photoresist is removed using remover solution, as it is no longer needed.

9.1.2.6 Wet etching of silicon (through-etch)

The next step is to perform the main removal of the desired sections of silicon. This is done using a wet etch: the sample is immersed in a temperature-controlled and circulated bath of 30% potassium hydroxide (KOH) solution, as shown in Fig. 9.6(a). At a temperature of 80°C, the etch proceeds at a rate of roughly 1 $\mu\text{m}/\text{min}$. Therefore, it takes on the order of 8 hours for the etch to proceed cleanly through the 500- μm -thick wafer.

Around 6 hours into the etch, the wafer is removed from the bath (see Fig. 9.6(b)), and slight mechanical pressure is applied to break the wafer into individual rectangular sections along the thinned channels between them (to prevent uncontrolled dissolution of the wafer into pieces). These rectangular sections are placed in a fixture and re-immersed in the bath until the etch is complete.

Once the etch has proceeded completely through the thickness of the pieces where desired, they are rinsed, dried, and then subjected to a brief bath in HF, which removes the nitride mask. The clean silicon pieces are then rinsed and blown dry once more. A photo of a sample at this stage is shown in Fig. 9.6(c).

9.1.2.7 Silicon nitride hard mask deposition (II)

At this stage, the only remaining task is to thin the central regions of the devices. Due to the difficulty of performing the full above patterning process on small rectangular pieces, and so as not to subject them to extra processing in any case, a very basic method of applying this second round of hard masking was devised.

First, the “bottom” side of the devices—the side that will not be etched in the thinning process—is coated once again with a film of silicon nitride via PECVD (typically only to 200-300 nm in this case, as the thinning etch is of shorter duration and therefore does not require as thick a mask). Then, the top side is also coated, however, a physical obstruction in the form of a sacrificial piece of silicon is aligned and placed on the central regions of the devices, as shown in Fig. 9.7(a). This obstruction acts to prevent any nitride from being deposited in this region. Therefore,




Figure 9.6: The KOH through-etch. (a) The wafer immersed in the KOH bath. (b) The wafer as it appears partially through the etch, prior to breaking into individual cantilever pieces. (c) The result of the through-etch: a rectangular piece of uniform thickness with a hole at one end.

at the end of the PECVD process, the result is a patterned hard mask without the need for photolithography and mask transfer via dry etching. The devices are now immediately ready to return to KOH for the thinning etch.




Figure 9.7: The second round of etching. (a) The samples resulting from the through-etch are physically masked during the second PECVD nitride deposition in order to the central region unprotected. (b) Two completed cantilevers, complete with a hole on one end and a thinned central region.

9.1.2.8 Wet etching of silicon (thinning etch)

The second round of wet etching in KOH proceeds very much like the first, however the etch is only allowed to proceed roughly halfway through the exposed material thickness. After an etch period of around 4 hours, the devices have a central-region thickness of about $250\ \mu\text{m}$, while the ends, which are protected by silicon nitride, retain their original thickness of $500\ \mu\text{m}$.

Following a final HF clean, DI rinse and drying, the cantilevers are complete and ready for use. Typically, they are packaged in sample carriers within a plastic bag purged with nitrogen gas, to prevent surface contamination and oxidation during

transport from the KNI lab to the LIGO lab. Two completed cantilevers are shown in Fig. 9.7(b).

CHARACTERIZING LOSS: RINGDOWN MEASUREMENTS

Prior to any interferometric measurements of the thermal displacement noise, the cantilevers are evaluated using the standard **ringdown** technique: an excitation¹ is used to ring the cantilever's motion up to some starting amplitude, and then the decay of that amplitude is recorded. Then, following the analysis in Sec. 8.1.1, the Q of the mode is determined from the decay time, τ .

The deflection of the resonator (i.e., the mode amplitude) is measured continuously using an **optical lever**: A HeNe laser beam is reflected off the surface of the cantilever and then directed onto a quadrant photodetector (QPD). The deflection angle is inferred from the local displacement of the reflected beam at the QPD.

10.1 The apparatus

This section describes the various pieces of equipment that are used to form the ringdown measurement setup.

10.1.1 The clamp

The most important component of the measurement setup is the clamp that holds the cantilever at one end. This component comprises two pieces of stainless steel screwed tightly together with the cantilever end sandwiched in between. The clamps were designed in SolidWorks and machined for this experiment.

An illustration of the design can be found in Fig. 10.1. Both pieces are made of steel, due to its high stiffness and density, to minimize the influence of clamping and recoil loss. The smaller piece is screwed into the larger piece—with the cantilever in between the two—using silver-plated 8-32 socket-head cap screws. Two steel dowel pins fit into holes in both steel clamp pieces to avoid skewing of the clamping. The other holes visible in the larger piece are for mounting to various other components of the experiment.

Proper alignment of the cantilever during clamping is paramount, especially during the interferometric measurements described in Chapter 11. To this end, an aluminum

¹This can be a more sophisticated, narrowband excitation from, e.g., an electro-static actuator, or it can be a very simple, broadband excitation from, e.g., banging one's fist on the table.




Figure 10.1: Stainless steel clamp used to hold the cantilever.

“alignment jig” was also designed. This jig attaches to the large part of the base and has a raised channel that serves as a reference position for the cantilever, as shown in Fig. 10.2. When the cantilever is in the correct position, the top piece of the clamp is tightened down onto the system. During the clamping process, the jig/clamp assembly is held down to an optical table with the three 1/4-20 screws visible in the bottom of the figure.

10.1.2 The cryostat

So that the loss can be evaluated as a function of temperature, the ringdown measurements are performed in a cryostat. The unit used for the ringdown measurements is an Infrared Laboratories (IRLabs) dewar-type cryostat. A tank above the experiment is filled with liquid nitrogen (LN_2), and this tank acts as a cold bath, thermally coupled to the experiment by a cold plate between the two. The experiment is conducted in vacuum, both to remove the damping effect of air molecules on the resonator, as well as to avoid condensation of water and ice on the system when cold.

A photo of the experimental volume, highlighting the important components, is shown in Fig. 10.3. All elements are attached to the cold plate, and a heat shield around the system prevents heating via radiative coupling with the room-temperature outer shell of the cryostat². The cantilever clamp assembly is attached to the cold plate via an adapter piece, which is made of PEEK for thermal stability: the relatively low thermal conductivity of the PEEK reduces the thermal link between the cold bath and the clamp, allowing the temperature of the clamp to be more readily elevated

²There is a small slit in this heat shield to allow the laser beam to enter and exit.




Figure 10.2: Clamp alignment jig.

by heating. The temperature of the clamp is monitored using a platinum resistive thermal device (RTD), and heat is applied to it using a resistive heater embedded in Kapton tape. An electro-static drive (ESD)—essentially two long electrodes spaced a few millimeters apart—is attached to the cold plate and held in close proximity to the cantilever. (Operation of the ESD will be described in the next section.) Finally, a 45° mirror directs the laser beam passing through the slit in the heat shield to and from the cantilever surface.

10.2 Operation

10.2.1 Thermal control

Control of the cantilever temperature is maintained by balancing cooling through thermal conduction of heat into the LN₂ cold bath with heating from the resistive heater. The height of the PEEK spacer was chosen based on the results of thermal modeling. In particular, it was chosen to be thick enough that the clamp/cantilever system could be kept at an elevated temperature without an unreasonable amount of heating power, but thin enough that the thermal response time of the system was




Figure 10.3: Photo of the ringdown measurement experimental volume.

manageable.

The temperature is measured using a 100- Ω RTD, which is a device that has a resistance that varies strongly with temperature. To read out the temperature, a “four-wire” sensing scheme is used, in which two wires are used to connect *each side* of the RTD (i.e., four wires total) to the electrical feedthrough at the cryostat wall. Through one pair of wires, an electric current is applied, setting up a voltage drop $V_{\text{out}} = IR_{\text{RTD}}$, where I is the applied current and R_{RTD} is the resistance of the RTD. The voltage V_{out} is measured using the other pair of wires. The reason for using two separate wire pairs is so that any resistance in the wires themselves do not influence the measured voltage, as would be the case if one simply measured the voltages of the current-carrying wires at the feedthrough.

The heater is a 25- Ω resistive heating element embedded in Kapton film. Two wires connect ports on the electrical feedthrough to the leads of this heater, and a voltage is applied across it to inject a heating power of $P_{\text{heat}} = \frac{V_{\text{heat}}}{R_{\text{heat}}} = \frac{V_{\text{heat}}}{25\Omega}$.

The thermal control system uses a digital loop running on a computer system not unlike those described in Sec. 2.1.2.1 for Advanced LIGO. The operation of the loop is as follows:

- The voltage across the RTD is acquired into the digital system through an analog-to-digital converter (ADC).
- The voltage is converted into a resistance using a (constant) current value provided by the user.
- The RTD resistance is converted to an absolute temperature using equations provided by the RTD manufacturer³.
- A desired temperature value is subtracted from the measured temperature, and this difference is used as an error signal for the digital control loop.
- A digital control servo filter conditions this error signal into a control signal to be fed (eventually) to the heater. The overall loop shape is typically a true integrator with a unity-gain frequency on the order of one 1 mHz, limited by the thermal response time of the system.
- The control signal is sent back out into the analog domain using a digital-to-analog converter (DAC).
- The DAC output is fed into a buffer or voltage-controlled voltage source, since the DAC can only drive a limited current.
- Finally, the output of the buffer or current source applies the voltage to the restive element.

This thermal control system allows for operation of the ringdown setup at nearly any temperature from just above 77 K to greater than 150 K.

10.2.2 Resonator readout and control

The signal from the QPD is fed into the same digital system described above for the thermal control. Once in the digital domain, this instantaneous displacement signal from the QPD can be measured using any of several independent readout and

³In narrow temperature regions, the function $T(R_{\text{RTD}})$ can be approximated linearly; since we require accurate temperature readings over a wide range of temperatures, a higher-order fit equation is used for this step.

control modules, called **mode ringers** for reasons that will become clear shortly. The structure of each of these modules is shown via the control screen depicted in Fig. 10.4.

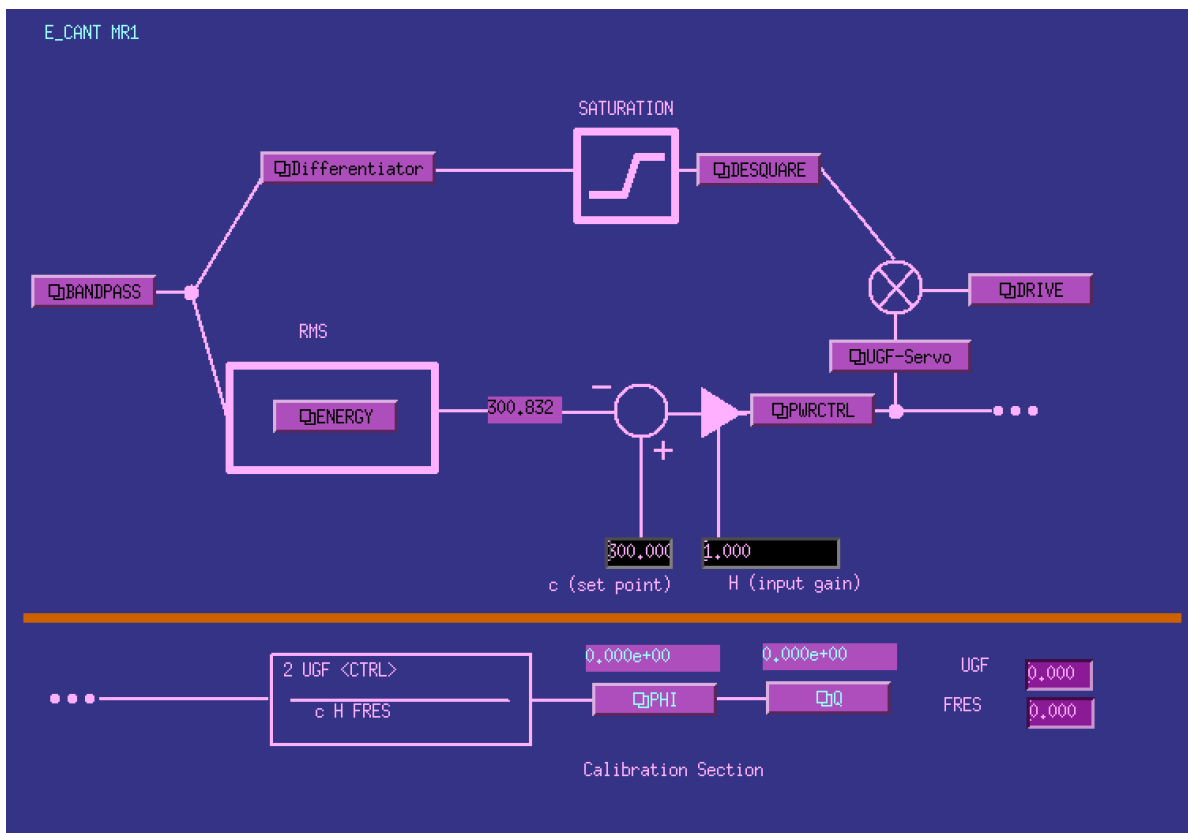


Figure 10.4: Control screen of the mode ringer module, showing the construction of the digital loop.

At the beginning of every mode ringer module is a bandpass filter. This is chosen to be a narrowband 4th-order butterworth bandpass filter centered on the frequency of the cantilever's resonant mode. The output of this filter therefore no longer has frequency components far from this mode of interest (in particular, there is no DC component, nor is there signal from any other resonant modes). From here, the signal is split into two paths.

On the lower path in the figure, the RMS amplitude of the bandpass filter output is taken⁴. In many cases, measuring this amplitude as a function of time is already

⁴This is accomplished via elementary operations in the model that are not depicted explicitly in the control screen: 1) the signal is squared, 2) the mean of the signal is taken by passing it through a low-pass filter of appropriately low frequency, and 3) the square of the output of the low-pass filter is computed.

enough to make a ringdown measurement, and, in practice, this “passive” method is the most commonly used.

For various reasons—in the presence of environmental noise that can affect the result of a ringdown measurement by exchanging energy with the system mechanically—it can be advantageous to perform Q measurements using the “active” technique devised by Smith [166]. The rest of the machinery in the depicted control screen are used to employ this technique, as described below.

Continuing along the lower path from the RMS amplitude, there is an amplitude set point that is subtracted⁵ from it, and then the difference is conditioned with a control filter (in the “PWRCTRL” block). The output of this control filter is thus a signed number proportional to the difference between the measured and desired amplitudes, with appropriate frequency shaping to maintain loop stability.

On the upper path, the signal from the bandpass filter passes through a differentiator, whose purpose it is to advance the phase of the signal by 90° . Within this block, the signal is also amplified by a large multiplication factor, so that its amplitude is $\gg 1$. This amplified, phase shifted signal is then *clipped* at ± 1 digital units, and then passed through a low-pass filter (“DESQUARE”) that rejects frequencies above the fundamental frequency of interest, with a maximum rejection at twice the mode frequency. The end result of these operations at the output of the desquaring filter is a constant-amplitude sine wave that is shifted by 90° from the input signal, but that is still always *phase-locked* with the signal. Crucially, this signal emerging from the upper path is ideally completely insensitive to amplitude fluctuations in the signal from the QPD.

To stabilize the cantilever ringing amplitude, a resonant feedback signal is generated by multiplying the output of the bottom path with that of the top path. Depending on the sign and magnitude of the control signal emerging from PWRCTRL, the applied signal is a sine wave either advanced or retarded from the cantilever’s motion by 90° , which is what is required to add or subtract energy from the cantilever, respectively.

This resonant control signal is sent out of the digital system via the DAC into a high-voltage amplifier, whose output is connected to the ESD. The amplifier scales the signal to a maximum level of ± 500 V for a maximum input, and also adds a DC

⁵On the control screen, this is indicated as an *addition*, but since the measured amplitude is *subtracted*, the resultant quantity is equal in magnitude to the measured value *minus* the desired value. The overall sign of the control filter is determined by the polarity of other components within and outside the digital domain.

offset of around 500 V. This leads to force actuation in the following way: 1) with no AC signal applied, the DC offset across the ESD electrodes induces an electric field, which in turn induces a charge on the surface of the cantilever, resulting in an attractive force; 2) when an AC signal is then applied, it leads to a first-order modulation of this attractive force, giving a linear coupling.

As described in [166], this technique is not only useful for stabilizing the resonator's amplitude, but also for making a *steady-state* measurement of the Q at this stable amplitude (as opposed to over a period where the amplitude rings down, as in the passive case). Since the control signal leaving PWRCTRL is proportional to the energy being resonantly added to (or subtracted from) the cantilever, then—even in the presence of random mechanical noise—the average of this signal should be proportional to the average amount of power we need to inject in order to keep the amplitude constant. This is clearly proportional to the loss angle, ϕ , via (8.8).

Conveniently, the calibration factor between this average control output and the loss angle is a simple function of the control loop shape and is independent of the individual transfer functions of the loop components (e.g., sensor, actuator, etc). The control loop typically has a UGF around 1-5 Hz.

10.3 Measurements

10.3.1 Expectations

The expected thermoelastic loss for the fundamental mode of the cantilevers is plotted as a function of temperature in Fig. 10.5. As can be seen, the contribution is on the order of 10^{-5} at room temperature, decreasing at lower temperature. Due to the dependence on CTE, the thermoelastic loss vanishes at 123 K. Therefore, we expect to measure a loss bounded only by the other loss processes at this temperature.

The loss from phonon-phonon damping is $< 10^{-8}$ and does not limit the measurement. However, the combination of clamping loss and surface loss has been observed to limit the Q of the lower modes in resonators of this type to the level of around 10^6 [150, 151, 162]. Since this commonly achievable loss is compatible with the acceptable limit calculated for LIGO Voyager in Sec. 8.2.4, the primary goal of this experiment was to fabricate devices at least as good as this, but little effort was allocated to improving fabrication methods over those established⁶; instead,

⁶Clearly, as described in Chapter 9, certain *modifications* to established fabrication methods were required due to the specifics of this experiment. However, the goal was simply to make cantilevers with these required modifications without sacrificing the quality achieved in standard cantilevers via established processes.




Figure 10.5: Calculated thermoelastic loss for the cantilevers as a function of temperature.

emphasis was placed on performing the direct, off-resonance displacement noise measurement on systems of this quality level.

10.3.2 Results

Typical ringdowns are shown in Fig. 10.6 for both the room temperature case and the cryogenic case. To ease in fitting, the natural log is taken of the vector containing the instantaneous amplitude as a function of time, and then this is fit to a line. This way, a simple linear fit can be used, with the linear fit parameter being equal to $-\frac{1}{\tau}$.

The results plotted show measured decay times of $\tau_{RT} = 430$ and $\tau_{cryo} = 1590$ for the room temperature and cryogenic cases, respectively. These ringdown measurements are performed *prior to the addition of the cylindrical mirrors*, which gives them a higher resonant frequency of about 76 Hz. Therefore, these correspond to a loss of around 10^{-5} at room temperature and of around 2.6×10^{-6} at 100 K.




Figure 10.6: Typical ringdown measurement results for the room temperature case and the cryogenic case.

The cryogenic loss is therefore not quite at the desired level, but not exceedingly high. On the other hand, the room temperature measurement is well explained by the sum of the expected thermoelastic loss and the anomalously high loss measured at low temperatures, which is assumed to contribute at room temperature, as well.

10.3.3 Analysis

The excess loss at low temperatures is not entirely understood—nor is it by other researchers who have observed qualitatively similar phenomena—but several potential culprits have been identified and investigated.

Ideal clamping losses from the finite stiffness of the steel clamp have been computed via FEA analysis. This is done using an integral like that in (8.22), but replacing the “surface” regions with the regions defined by the clamp. Assuming a Q of steel on the order of a few hundred to a thousand, this contribution does not come close to explaining the measured cantilever Q limitation at low temperatures.

The above calculation assumes a perfect interface between the various elements of the system under analysis (in particular, between the surfaces of the cantilever and those of the clamp), which can never be the case; in reality, surface effects known as “slip-stick” losses [150] can amplify the influence of mechanical loss in the clamp material. To investigate this effect, ringdowns were also performed using a different clamp design wherein the clamping surfaces were polished to a mirror finish, but the maximum attainable Q was not significantly improved. Therefore, it is not believed that clamp loss (including slip-stick effects) is a limiting source of loss in our system.

Another clamping-related effect that has been noted [162] is a non-repeatability of the cryogenic loss with successive re-clamping. That is, some trials exhibit a loss that is higher than the lowest achievable with a given sample, and, upon un-clamping and re-clamping, it is possible to achieve a better result. This effect is also not very well understood, and a variation in the measured loss from trial to trial (with reclamping in between each) was observed in our measurement, as well.

Finally, recoil losses are believed to have played a part in limiting the measured low-temperature Q of the cantilevers. In particular, the stiffness of the cryogenic system—i.e., the stiffness between the clamp and the rigid table to which the cryostat was mounted—was limited, both due to the construction of the cryostat and due to the semi-rigid wires attached to the clamping block for the RTD and heater. To investigate this hypothesis, room-temperature measurements were carried out in a simpler, more rigid vacuum chamber directly bolted to the table. Then, the clamp/cantilever assembly was moved to the cryostat and re-measured at room temperature. Additionally, the measurements in the cryostat were performed with and without the wires attached to the clamping block. The effect of moving from the rigid simple chamber to the cryostat was not conclusive above the variability from measurement to measurement; on the other hand, the measurement with the wires exhibited a noticeably higher loss than did either of the other two cases.

*Chapter 11***INTERFEROMETRIC INTERROGATION OF CANTILEVER
DISPLACEMENT NOISE**

This chapter describes the effort to directly interrogate the thermal displacement noise of the cantilevers using interferometry. Unfortunately, this ultimate goal was not realized by the time of this writing. However, significant progress was made in this effort, and a path forward has been identified.

The organization of the chapter is as follows. In Sec. 11.1, an overview of the experiment is given, including the optical layout and basic measurement scheme. Sec. 11.2 details the individual optical, electronic, and mechanical components of the experiment. Sec. 11.3 journals the progress made in this experiment over the course of multiple phases. Finally, in Sec. 11.4, a conclusion and path forward for this experiment are presented.

11.1 Overview**11.1.1 Design**

The optical layout for the experiment is shown in Fig. 11.1. From it, one can glean the basic scheme of the measurement. There are two parallel and largely independent optical systems (typically referred to as the “East” and “West” systems, due to their physical locations on the optical table). Each one interrogates one of two linear cavities on a common suspended platform within a large cryostat. Each cavity comprises a silicon ribbon cantilever (with a high-reflectivity mirror attached at the tip) on one side and a macroscopic commercial mirror on the other.

The laser from each system is locked to its respective cavity via feedback to the laser frequency using the Pound-Drever-Hall technique (see Appendix B). The cavities are nearly critically coupled, so a large fraction of the input laser power is transmitted through the end mirrors. Here, the transmitted beams from each cavity are spatially overlapped on a beamsplitter, and the resultant co-spatial beam is sensed on an RF photodetector (RFPD). If the absolute frequency difference between the two beams is within the bandwidth of this RFPD, then a strong beat signal is recorded electronically, and the frequency fluctuations of this beat signal encode the differential length fluctuations between the two cavities.

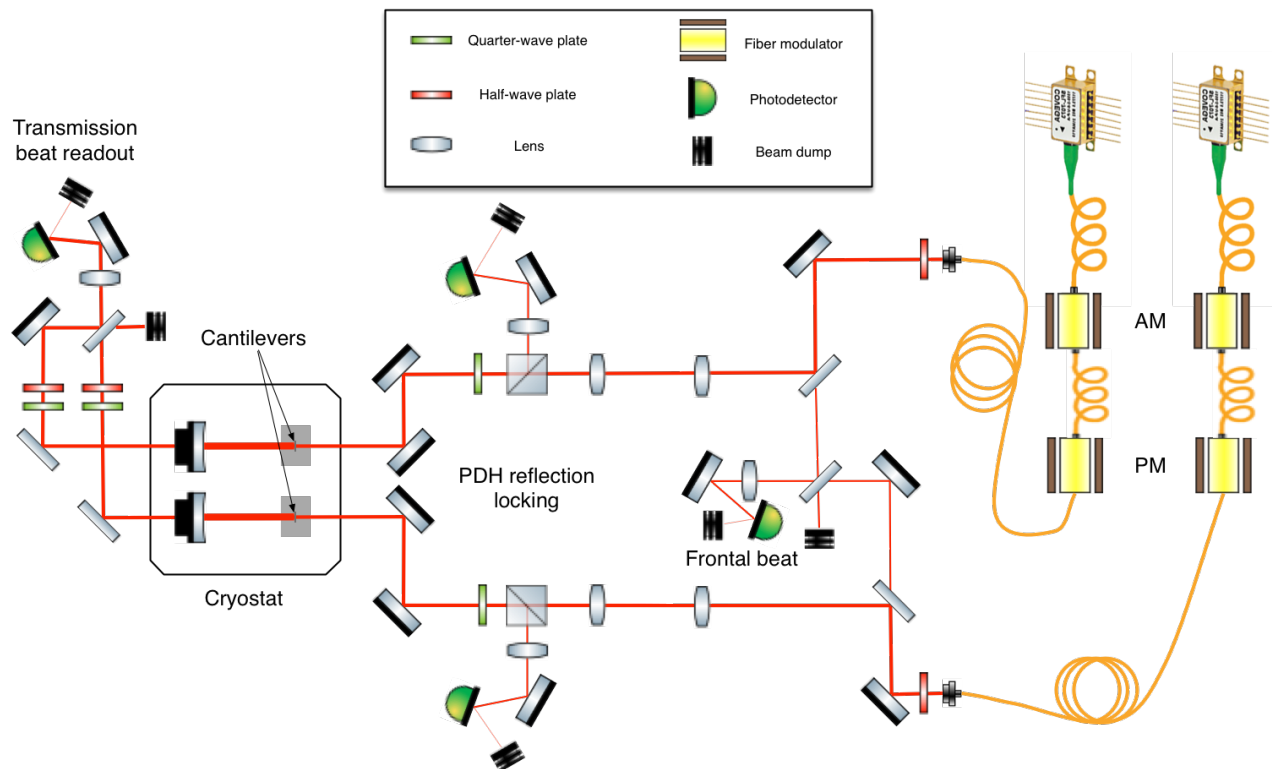


Figure 11.1: Optical layout for the interferometric measurement.

The phase modulators serve to apply the RF phase modulation sidebands required for PDH locking, while the amplitude modulators allow for independent control of the laser power, as the lasers have a strong coupling from drive current to frequency. The “frontal beat” RFPD is used to find an operating point at which the two lasers have the same frequency before locking the cavities.

The breadboard holding the optical cavities is suspended from the cryostat cold plate via a three-wire single-stage suspension that is a modified version of that used for the Advanced LIGO output mode cleaner (see Part III). As with that design, the upper ends of the suspension wires are attached to the tips of steel blade springs for added vertical isolation.

The goal of this design is to make a system that is strongly insensitive to environmental disturbances that affect both cavities in *common*. Additionally, the transmission beat measurement is strongly insensitive to the frequency noise of the lasers, since the PDH feedback loops suppress this noise in the transmitted light. Assuming that the two cantilevers exhibit the same level of thermal displacement noise, and that this noise is dominant over other *differential* sources of displacement noise, the

measured noise in the beat readout should be $\sqrt{2}$ times the noise calculated for a single cantilever.

11.1.2 Target sensitivity




Figure 11.2: Initial noise budget estimate for the interferometric experiment.

A preliminary noise budget constructed for the experiment is shown in Fig. 11.2. This model was built using code written by Wipf [167] for use with gravitational wave detectors and related experiments. In this model, the indicated noise sources are estimated as simple power-law functions based on applicable existing measurements. Likewise, control loop transfer functions are also estimated using simple pole-zero models. Then, a frequency domain model is constructed by the noise budgeting code that calculates the expected propagated contributions of each defined source to each measurement degree of freedom of interest.

This model was used to evaluate the feasibility of the measurement, given expectable performance of isolation systems and suppression loops. The plot above includes the effects of realistic laser frequency and intensity control loops, which were

determined not to be a major impediment. Based on these calculations, seismic noise was expected to be the limiting contributor of noise. To achieve the estimated performance in the plot, a certain amount of passive isolation was assumed, as well as some suppression of differential noise from symmetry in the physical plant. These are detailed further in the sections below.

11.2 Detailed description

This section gives a detailed description of the various individual components in the interferometric experiment.

11.2.1 Optical system

This section describes the various optical components in this experiment *apart from* the optical cavities themselves; these will be detailed in their various incarnations in Sec. 11.3.

11.2.1.1 Lasers, modulators and photodetectors

The light sources for the experiment are RIO PLANEX [168] external cavity lasers (ECLs) operating with a wavelength near 1550 nm and a maximum output of around 20 mW, with fiber-coupled outputs. The laser head temperatures are stabilized with commercial controllers, and the laser currents are provided by custom-made low-noise drivers [169]. The lasers have a wavelength tuning range of ± 30 pm; in order to ensure that the lasers would be within tuning range of each other—a requirement for the beat measurement—the two units were hand-selected to have close center wavelengths.

Both modulators are iXblue Photonics fiber-coupled units. The amplitude modulator is an MXAN-LN-10. Based on an internal Mach-Zehnder interferometer, this unit has a half-wave voltage of 5.5 V and an electro-optic bandwidth of around 10 GHz. The phase modulator is an MPX-LN-0.1, with a half-wave voltage of around 3.5 V and a bandwidth of 150 MHz¹. Both units have an insertion loss near 3 dB.

The specifications of the lasers and modulators were verified using an auxiliary test cavity with high mechanical stability (similar to the mode cleaning cavity described in Part II). Of particular importance was a verification of the laser frequency noise

¹The modulation frequencies for the east and west systems are 33.59 MHz and 32.70 MHz, respectively.

at the manufacturer-specified level of

$$f_n^{\text{laser}} \approx 100 \times \sqrt{\frac{1 \text{ kHz}}{f}} \frac{\text{Hz}}{\sqrt{\text{Hz}}} . \quad (11.1)$$

Sensing the reflected beams from the cavities for PDH locking are Thorlabs PDA255 broadband InGaAs photodetectors. These units have a bandwidth of 50 MHz and a transimpedance gain of 5000 V/A with 50- Ω termination. The inter-system beat measurements are made with New Focus 1611 broadband InGaAs photodetectors. These have a very large bandwidth of 1 GHz and an RF transimpedance of 700 V/A.

11.2.1.2 Focusing, isolation and polarization control

The beam profiles out of the collimators that couple the light from the fibers into free space were measured and used to calculate the required mode-matching telescope (MMT) lens focal lengths and positions. This MMT comprises two lenses between the steering mirrors directly after the fiber collimators and the polarizing beamsplitters.

The beams exiting the fibers through the collimators are linearly polarized, and they are rotated to P polarization (i.e., parallel to the table) by half-wave plates (HWPs) directly in front of the collimators. This polarization causes each input beam to pass through the polarizing beamsplitter (PBS) cube in its path and then through a quarter-wave plate (QWP) that circularizes its polarization, after which it continues on to its respective cavity. The reflected beam from each cavity retraces its path *back* through the QWP, which re-linearizes its polarization, only in the orthogonal orientation (i.e., S). As a result, the reflected beams are rejected by the PBSs and directed onto their respective RFPDs for PDH locking. In transmission of each cavity, a QWP and HWP in series act to re-linearize each beam's polarization and then rotate it to the P polarization, for which the recombining beamsplitter is designed.

The second mirror in each beam path is actually a beamsplitter with a transmission of 10%, allowing some light to pass to the frontal beat measurement setup.

A set of mirrors near the input side of the cryostat serve as compound periscopes, elevating the beam height from the standard level of 4 inches to the cavity height of around 6.2 inches. On the transmission side, due to the relatively few optical components, the height is maintained at 6.2 inches by increasing the length of the optics' pedestals.

All turning mirrors are commercial units with broadband NIR high-reflectivity coatings. All lenses have broadband NIR anti-reflectivity coatings.

11.2.2 Physical system

11.2.2.1 Cryostat

The cryostat used for the interferometric experiment is essentially a larger version of the one used in the ringdown setup in Chapter 10. Additionally, the cryostat has four large windows for proper insertion and extraction of the laser beams from the inner volume. A drawing of the cryostat, showing some important dimensions, is given in Fig. 11.3. Note that the cryostat sits atop an 0.5-inch base plate that is used to mount the system to the optical table (not pictured in the drawing), so the height of the window centers is shown to be shorter than the 6.2 inches mentioned in the previous section. The usable vertical distance from the cold plate down to the base of the volume is 10 inches.

Cryogenic measurements have not been performed with this system at the time of this writing; however, the system has a dewar for LN₂ (as with the smaller cryostat), as well as a heat shield with holes aligned with the windows. Likewise, no electrical connections into the vacuum envelope have been needed thus far, but a 10-pin feedthrough connector exists for this purpose.

The cryostat is pumped with a combination roughing/turbo pumping station and readily achieves a vacuum on the order of 10⁻⁶ Torr.

11.2.2.2 Suspension

As mentioned in Sec. 11.1.2, the influence of ground motion was of great concern when designing this experiment. To this end, a single-pendulum suspension stage was designed to hang the optical cavity platform from the downward-facing cold plate of the cryostat. This design incorporated many elements of the Advanced LIGO OMC suspension design (see Part III). In particular, the blade springs, blade spring clamps, wire clamps, and copper ferrules used to engage the lower end of the wires with the suspended breadboard were all taken directly from surplus aLIGO components or machined anew from the aLIGO drawings.

A diagram of the suspension is shown in Fig. 11.4. The major components of this suspension are:

- **Adapter plate:** A stainless steel plate that mates with the hole pattern on the




Figure 11.3: Drawing of the cryostat used to hold the ribbon cantilever cavities.

cold plate and provides holes for mounting of the blade spring clamps. This part was custom designed for this experiment.

- Blade springs:** Maraging steel blade springs pre-formed to be flat under a prescribed weight load. These blade springs, as well as the associated blade clamps attaching them to the adapter plate, were repurposed from the aLIGO suspension design. (The blade springs themselves were surplus parts, while the clamps were machined from existing drawings.)
- Steel wire:** 0.0079-inch steel music wire. The wire stock itself was taken from surplus aLIGO supply, as were the conical copper ferrules that are crimped onto the bottom ends of the wire and engage with conical sockets in the breadboard. The wire clamps that attach the wire to the tips of the blade springs were machined from existing drawings.




Figure 11.4: Diagram of the suspension from which the optical breadboard is hung within the cryostat.

- **Breadboard:** Aluminum breadboard onto which the optical cavity assembly is built. This was a custom design based on a modification of a commercially available part.

The overall suspension design was an optimization based on the space constraints of the system. Due to the area available on the cold plate, the maximum number of blade springs (and therefore suspension points) was determined to be three. The design load of the blade springs used was about 1.45 kg each, which specified a total payload mass of 4.35 kg. The length of the wires ($L_{\text{wire}} = 0.2194$ m) was chosen to be as long as possible, allowing for a small clearance between the bottom of the suspended breadboard and the base of the cryostat volume.

The predicted performance of this system was computed using the suspension modeling tool SUMCON [170], and the results of this analysis are shown in Fig. 11.5. As depicted, the system is stable, with all six mode frequencies lying between 0.8 Hz and 1.8 Hz². Above these resonant frequencies, coupling of motion from the optical table to the suspended breadboard falls off as f^{-2} .

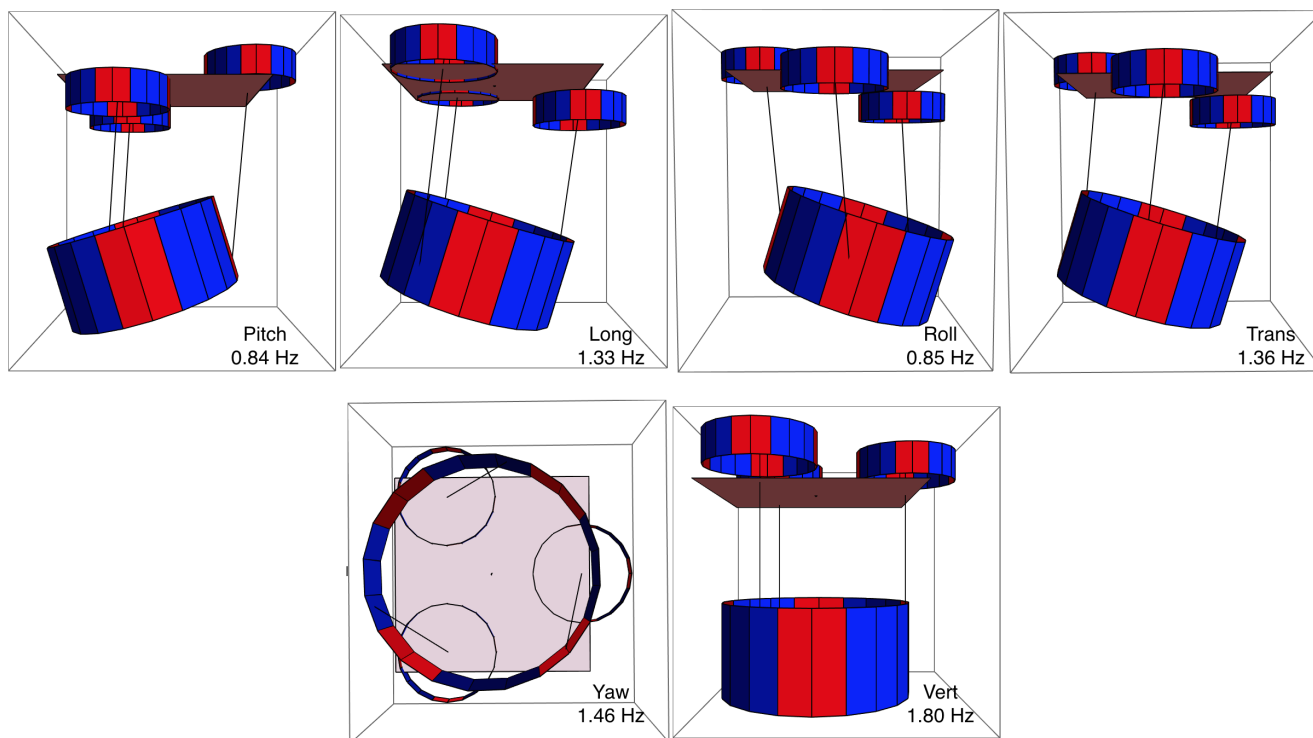


Figure 11.5: Modal analysis of the suspension.

11.2.3 Electrical system

The PDH locking loops are built primarily with commercially available components, with a few exceptions. First, the laser drivers are custom designed, as discussed above. The feedback servo modules are also custom designed, incorporating multiple switchable gain stages, a variable proportional gain, and input/output offset control. The RF sources used to provide the PDH phase modulation sidebands are a collection of commercially available Wenzel devices combined into a custom rack-mount chassis.

The beat readout system is also composed primarily of off-the-shelf components. Due to evolving requirements, this system has changed somewhat over the course

²This analysis is dependent on the moments of inertia of the suspended payload. Best guesses were used for these based on expected locations of optical components on the breadboard. The realized system exhibits mode frequencies not strongly divergent from those reported here.

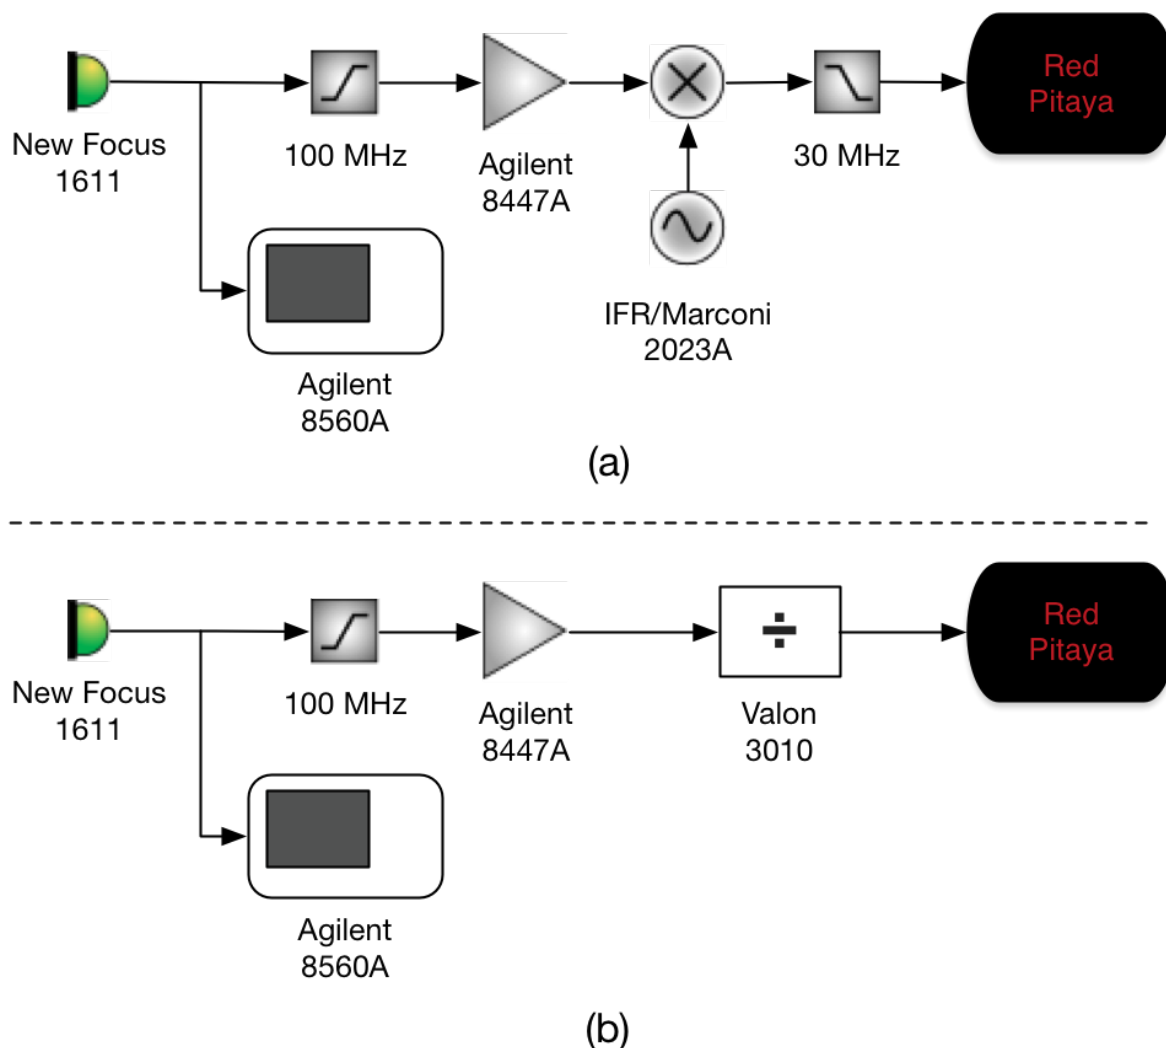


Figure 11.6: Two electronic chains used in the differential cavity beat note readout.

of the experiment. Two commonly used readout chains are shown in Fig. 11.6. Both systems begin with the beat PD, a New Focus 1611, and end with a Red Pitaya [171] FPGA-based digital lab instrument. This device can directly digitize signals with frequencies up to 64 MHz, and an onboard digital phase-locked loop enables frequency tracking with very high range and a bandwidth near 100 kHz. However, the Red Pitaya requires a specific input signal amplitude, and the beat signal frequency can in general be much greater than 64 MHz. Therefore, some additional conditioning is needed between the beat PD and the Red Pitaya.

For the majority of the experiment's lifetime, the scheme used was that found in Fig. 11.6(a). Due to the cavity design, as described later, the beat signal can be as

high as 1 GHz. To locate the beat frequency to begin with, a high-frequency signal analyzer is used directly at the PD. Once the base frequency is known, the signal can be demodulated by mixing with the output of a function generator tuned to a nearby frequency and then low-pass filtering.

This method above reduces the mean frequency of the signal to be acquired, but it does not reduce the amplitude of frequency deviations due to differential cavity motion. In some configurations, these deviations were observed to be too large themselves to be acquired properly by the Red Pitaya. After some searching, a solution to this problem was found in the form of a programmable frequency divider [172], as shown in Fig. 11.6(b). This method divides both the mean frequency and the deviations proportionally, so that the full, time-varying signal can be acquired by the Red Pitaya.

In both cases, variable amplification and passive inline attenuation were used to maintain appropriate Red Pitaya input levels of $\sim 1 V_{pk}$.

11.3 Experimental progress

This section details the progress made with this experiment towards an eventual measurement of off-resonance thermal displacement noise in the silicon ribbon cantilevers, divided into three phases: the test cavity phase, the dummy cantilever cavity phase, and the real cantilever cavity phase.

11.3.1 Phase I: Test cavities

11.3.1.1 Description

The first phase of interferometric testing was accomplished using test cavities comprising two identical macroscopic mirror assemblies (i.e., with no cantilevers). The purpose of this phase was to verify that the test setup was capable of making a measurement at the required sensitivity levels for the eventual cantilever measurements.

The test cavity payload is shown assembled in Fig. 11.7(a). Each cavity consists of a pair of 1-inch commercial mirrors, with each mirror held in a commercial kinematic mount (Thorlabs Polaris) atop a standard 1-inch pedestal base. The distance between the two mirrors' HR surfaces (i.e., the cavity length) is $L_{cav} = 7.25$ cm. Each mirror has a power transmission of $T = 165$ ppm and a radius of curvature of 0.5 m. This leads to a stable cavity with an FSR or $\nu_{FSR} = 2.14$ GHz and a finesse of $\mathcal{F} = 19000$.

The mirrors used for both sides of the cavity in this phase are also used as the end mirrors in the cantilever cavities. The cantilever mirrors are flat and have a slightly




Figure 11.7: Test cavities. (a) Assembled on the breadboard to be suspended. (b) Suspended and being lowered into the bottom section of the cryostat for testing.

higher transmission, however, the test cavities are a fairly close approximation of the eventual optical plant.

11.3.1.2 Measurements and analysis

The evolution of the beat noise over early part test cavity phase is shown in Fig. 11.8. At the top (in blue) is the displacement signal from a single cavity as measured by taking the control signal to the laser driver and calibrating it with the known actuation gain (in Hz per V). This signal contains the unsuppressed laser frequency noise, and is therefore very noisy.

Each successive trace in the plot was measured after a significant improvement in the experimental configuration (e.g., obtaining and measuring a beat signal instead of the aforementioned control signal, sealing the chamber, pumping it down to vacuum, etc.). Notably, a large improvement is had at frequencies below 100 Hz by “floating” the optical table on its pneumatic legs.




Figure 11.8: Evolution of the beat spectrum over part of the test cavity phase.

The remainder of the test cavity phase was spent identifying the source of and improving the noise limit in two distinct regions: 1) in the region from 100-600 Hz where large peaks were observed, and 2) in the region above 600 Hz where a broadband floor was observed.

The noise in the former region was presumed to be from mechanical resonances in the payload. To investigate this, a detailed FEA model of the payload was constructed, as shown in Fig. 11.9. The results of a modal analysis using this model were then compared with the observed forest of peaks in the region of interest, and several offensive modes were identified, as demonstrated in Fig. 11.10. It was predicted that this effect was exacerbated by the symmetry of the test cavity payload, which behaved somewhat like a tuning fork, so a modest improvement was expected when switching to cavities with cantilevers on one side.




Figure 11.9: Detailed FEA model of the test cavity payload.

The noise latter region was presumed to be from sensing noise or other electronic effects. Many avenues were pursued in trying to reduce this broadband noise floor, including a systematic swapping out of every electronic component in the locking loops from the photodetectors to the laser drivers, as well as in the beat measurement setup from the photodetector to the Red Pitaya. No conclusive culprit was found, however, anomalous dependence on the gain states of the two locking loops was observed. Ultimately, the noise was reduced by a factor of around 2 to 10^{-16} m/ $\sqrt{\text{Hz}}$ by iterative tuning of the two loops while monitoring the beat noise.

The best hypothesis formed was that there was a nonlinear coupling either up- or down-converting noise into this band. No specific mechanism was discovered, but similar effects could be created by intentionally, e.g., by injecting an RF line into the error point of the PDH locking servos.




Figure 11.10: Test cavity mode identification. The red lines indicate modes predicted by the FEA model shown in Fig. 11.9.

Finally, with the slight reduction in the high-frequency noise, it became clear that there were some apparent true displacement noise features hiding just beneath the broadband noise in the same band. To investigate these, a technique was devised by which the noise of each cavity could be compared independently to a quiet laser locked to a rigid reference cavity, typically used for a separate experiment. This technique is not ideal for the main measurement, since it does not take advantage of any mechanical common-mode suppression that exists between the two closely spaced cavities. Nevertheless, a broad feature around 3.2 kHz was revealed to in fact be a composite of four separate narrower features—two from each cavity—using this technique, as shown in Fig. 11.11.

11.3.2 Phase II: Dummy cantilever cavities

11.3.2.1 Description

The next phase incorporated fabricated cantilevers, but ones which did not have their central regions thinned. These “dummy” cantilevers were used to test the compatibility of all mechanical parts and verify operation of the entire system as





Figure 11.11: Inter-experiment beat measurement. The E and W cavities were individually compared against a third, much quieter cavity from another experiment. The broad feature near 3.2 kHz was revealed to be a composite of separate features from the E and W cavities.

designed without fear of breaking the devices.

Two modifications were necessary to build the dummy cantilever cavities: 1) Mirrors needed to be affixed to the dummy cantilevers, and 2) an adapter piece was needed between the suspended breadboard and the clamp/cantilever assembly.

The mirrors were attached to the cantilevers using **optical contacting**: when two extremely flat and clean surfaces are brought together, they can become bound due to the strong intermolecular forces present at these length scales. In a clean environment, the mirrors are cleaned using First Contact [173], then, carefully, the mirror is set over the hole in the cantilever and moderate pressure is applied for around one minute. Following this, the cantilever can be held from the base end and shaken to verify a successful bond. The process is depicted in Fig. 11.12.



(a)



(b)

Figure 11.12: The mirror is attached via optical contacting. (a) The mirrors and cantilevers are cleaned using First Contact. (b) A completed mirror/cantilever assembly.

The adapter piece was designed to mate with the hole pattern on the optical breadboard on the bottom and to provide holes for the cantilever clamps to mate with on the top. Care was taken so that these would be roughly centered in the cryostat windows both horizontally and vertically (though some maneuverability is afforded by the movable commercial end mirrors). Photos showing the attachment of the adapter piece, as well as demonstrating in-air locking of a dummy test cavity, are shown in Fig. 11.13. Due to payload mass constraints, these pieces were made out of aluminum.




Figure 11.13: Demonstration of the adapter piece between the cantilever clamps and the suspended breadboard. The construction is shown in (a), while (b) shows another angle and also demonstrates early in-air locking of the dummy cantilever cavity. (Note: For simplicity, the cavities were built and aligned rotated 180° from the usual orientation, as shown here.)

The cantilever mirrors are 1-cm-diameter silicon substrates with an HR coating of $T_{\text{cant}} = 300$ ppm on one side and an AR coating on the other side. They are flat with an 0.5° wedge, and the HR side is what is contacted to the hole on the cantilever. Due to the change in the cavity spatial mode, the MMT was slightly adjusted to ensure proper matching. Also, the increase in transmissivity of the input mirrors when transitioning from the test cavity case to the dummy cantilever cavity case reduced

the finesse to closer to $\mathcal{F} \approx 13000$, while the FSR stayed the same at $\nu_{\text{FSR}} \approx 2$ GHz. The cavities were aligned using an iterative process that ensured proper positioning of the beam spots on the centers of the mirrors³. The process was as follows:

- The cantilever/clamp assemblies were mounted to the breadboard using the adapter pieces.
- The breadboard was attached to the table rigidly at a height very close to where it sits when suspended, however, it was rotated 180° and slid one cavity length closer to the lasers. (The sliding was necessary to ensure that the new, rotated orientation matched the mode of the input beam.)
- Using the actuators on the input periscope mirrors, the beams were centered on the cantilever mirrors *and* retro-reflected along their input paths. This required iterating over the various input periscope mirror degrees of freedom.
- The macroscopic mirrors were installed such that the input beams were centered on them and still retro-reflected as before.

As a result of this process, the cavity modes themselves were aligned. The system was then suspended in its usual orientation and position, and alignment was found using the input periscope mirrors and searching for mode flashes in transmission.

11.3.2.2 Measurements and analysis

Immediately upon locking the cantilever cavities, it became clear that the resonant motion—even of the dummy cantilevers—could be large enough to destabilize the PDH locks by exceeding the laser actuation range. Additionally, the increased mechanical admittance of the cantilevers caused much stronger coupling of the resonant suspension motion. In particular, the yaw mode at ~ 1.5 Hz, which couples to differential motion between the cavities, induced large fluctuations to the control signals at this frequency.

Both effects were ameliorated significantly by increasing the actuation range of the laser drivers. Since the suppressed laser frequency noise was not a limiting source, this did not have any adverse effects.

³Due to the much smaller diameters, slight mis-positioning of the supported-mode beam spots from the centers of the cantilever mirrors could lead to strong diffractive losses.

Meanwhile, with respect to the resonant cantilever motion, a novel control scheme was developed using an adaptation of the technique detailed in Sec. 10.2.2, only using radiation pressure actuation (using the fiber amplitude modulators) in place of the ESD: since the admittance of the cantilevers is very high on resonance, small power fluctuations applied coherently can be used to ring their modes up and down. The input to the Moderinger bandpass filter in this case was the laser feedback control signal.

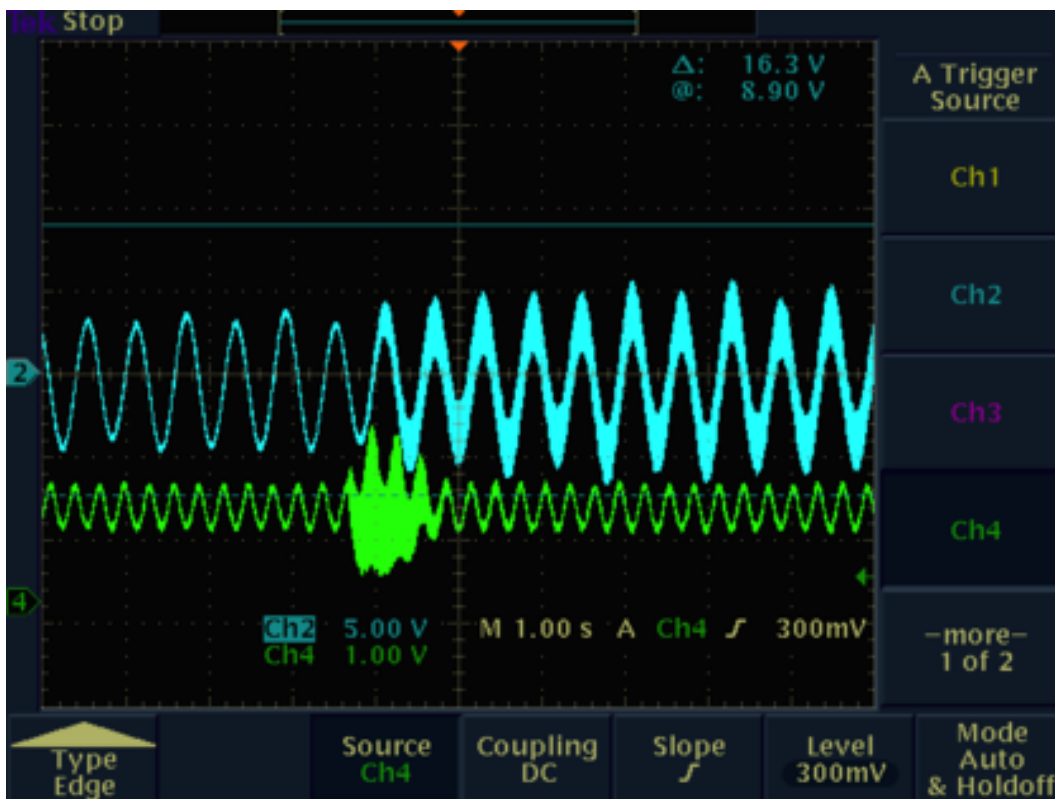


Figure 11.14: Demonstration of the radiation pressure cantilever actuation. The laser frequency feedback control signal is in cyan, while the transmitted power is in green. The large signal at 1.5 Hz is from the suspension yaw mode, and it persists throughout. At around $t = 3.75$ s, a radiation pressure excitation is engaged on the fundamental cantilever mode at around 110 Hz. Immediately, a 110-Hz fuzz becomes apparent in the control signal, indicating that the cantilever is moving at this frequency and the laser is being acted on to follow its motion. About 1 s after the excitation begins, the requested mode amplitude is reached, and the magnitude of the excitation is reduced, since it must only act when the mode amplitude deviates from the desired level. The amplitude control feedback remains engaged through the rest of the sampled interval.

A demonstration of this radiation pressure control is shown in Fig. 11.14. This technique worked well to keep the on-resonance motion of the cantilevers low

in the presence of environmental excitations, and also allowed for in-situ ringdown measurements of the cantilever Q s. A demonstration of such an in-situ measurement is illustrated in the control system screenshot in Fig. 11.15. Using this technique, the Q of the dummy cantilevers was estimated to be about 13000, likely dominated by clamping loss from the uniform-thickness geometry. This method was also used to measure the Q of the second bending mode at around 1 kHz, which was found to be considerably lower at around 1000.

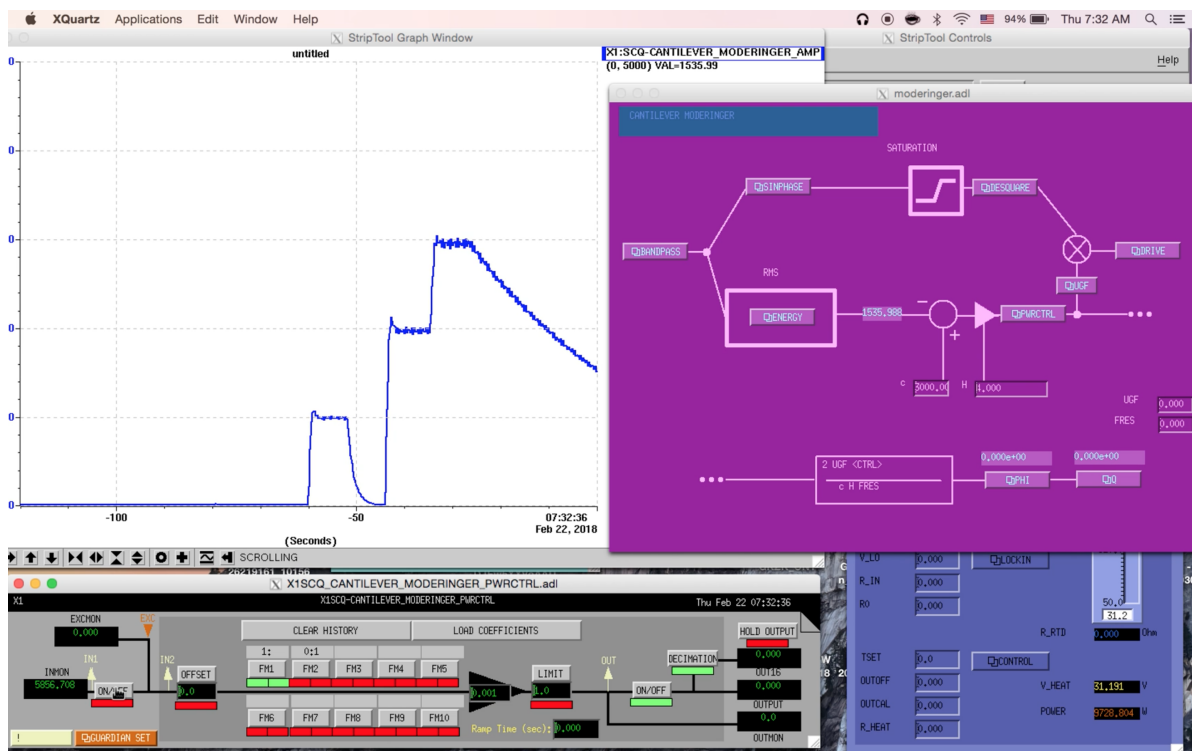


Figure 11.15: In-situ cantilever ringdown using radiation pressure excitation. The instantaneous cantilever fundamental mode amplitude is shown in the plot at the top left. At $t = -60$ s, RP control is engaged with a requested amplitude of 1000, and the cantilever is rung up accordingly. Then, the requested amplitude is switched to 0, then 2000, then 3000, before the system is disengaged and the cantilever is allowed to freely ring down.

The best beat measurement obtained with the dummy cantilever cavities is shown in Fig. 11.16, along with a thermal noise estimate for that system based on the in-situ measured Q s as well as a projection of the thermal noise from the real cantilevers assuming a uniform Q of 10^5 across all modes. As can be seen, the dummy cantilever beat was observed to be very nearly limited by the (elevated) thermal noise contribution in multiple bands.




Figure 11.16: Noise in the dummy cantilever cavities compared with expected thermal noise, along with a projection of thermal noise in the real cantilever cavities.

In order to compare the thermal noise estimates for the dummy cantilever case and the real cantilever case and determine whether the measurement of the latter was feasible, the change in the mechanical response—deflection of the cantilever per unit force on the cantilever—from the dummy to the real case also had to be considered. Such a consideration was made using a frequency-interpolated modal FEA analysis, as illustrated in Fig. 11.17.

Using the results in Figs. 11.16 and 11.17, it was possible to estimate what—if any—further environmental noise suppression would be necessary to make the thermal noise measurement with the real cantilevers. For example, from Fig. 11.16, it is predicted that the real cantilever system will exhibit the same displacement noise spectral density as the dummy cantilever at 60 Hz. Therefore, one might naively think that a modest improvement in the noise suppression here would be sufficient for the thermal noise measurement. However, judging from Fig. 11.17, it is clear that the real cantilevers will have a ~ 15 -dB increase in mechanical response to external




Figure 11.17: Comparison of the mechanical responses of the dummy and real cantilevers. The responses shown are the transfer function from force on the cantilever mirror to displacement of the mirror.

disturbances at 60 Hz, relative to the dummy cantilevers. Therefore, this additional noise reduction of 15 dB is needed to reach the goal.

11.3.3 Phase III: Real cantilever cavities

11.3.3.1 Description

Given the foregoing analysis, it was clear that significant suppression of environmental noise was necessary to reach the goal of measuring the thermal noise of the real cantilevers. Improvements at low frequencies are difficult to obtain without more involved modifications, as suggested in Sec. 11.4. Therefore, emphasis was placed on reducing the severity of the mechanical peaking in the middle-frequency band between 100-600 Hz.

During the transition from the test cavities to the dummy cantilever cavities, it was hoped that the removal of the symmetries in the suspended payload mass

distributions would lessen the effects of the payload body mode resonances. This did not happen, so, for the real cantilever measurement, a new, monolithic payload was designed.




Figure 11.18: Monolithic suspended payload design.

This new design is depicted in Fig. 11.18. The only newly machined piece is the large central component. Made of aluminum due to payload mass constraints, this part mates directly with the cantilever/clamp assemblies, as well as the 1-inch commercial mirrors in their Polaris mounts, significantly increasing the overall stiffness of the payload and shifting its body mode resonances up and out of the target band.

11.3.3.2 Measurements and analysis

The realized beat noise performance for the real cantilever cavities, compared with that for the dummy cantilever case, is shown in Fig. 11.19. As expected, the noise is improved in the middle frequencies from the stiffening of the payload as described above. Also as expected, the noise at lower frequencies has increased, due to the stronger coupling of (unchanged) environmental noise to the more responsive thinned cantilevers. Exactly as predicted, for example, the noise at 60 Hz lies about an order of magnitude above the estimated thermal noise contribution: this is the 15 dB increase in mechanical response on top of the excess noise measured at this frequency in the dummy cantilever case in Fig. 11.16.




Figure 11.19: Comparison of the measured noise in the real and dummy cantilever cases. Measurements are shown with the optical table floating (solid) and rigidly attached (dotted) to the floor. The thermal noise estimate for the realized room-temperature configuration is given in solid green, while that for an eventual cryogenic measurement is given in dotted green.

Therefore, the noise coupling model appears to be correct, and the main objective of the change to a monolithic payload has been largely achieved, as the noise between 100-600 Hz has been improved by up to two orders of magnitude. If the thermal noise estimate is accurate, however, this was not a sufficient improvement to intercept it. At the best frequency of just around 100 Hz, the beat noise is around a factor of 2 above the estimated thermal noise.

On the other hand, the scalloped shape of the measured noise in this middle-frequency band is undoubtedly suspicious. Given this, one can play devil's advocate and ask what kind of loss would be necessary to explain this feature using thermal noise. Such an attempt is made in Fig. 11.20. In the scaled thermal noise estimate plotted there, the loss of the second bending mode at around 500 Hz has been increased by a factor of 30 to 3×10^{-4} , while that for the mode around 3 kHz has been increased by 10000 to 0.1. Unfortunately, due to technical issues, it was not possible to measure the Q of these modes accurately in-situ.

11.4 Conclusion

In conclusion, significant progress has been made towards an eventual measurement of off-resonance thermal noise in the silicon ribbon cantilevers. Multiple phases have been undertaken, and improvements have been made with each transition from one to the next. Several future steps have been clearly identified, as described below.

Crucially, the environmental noise coupling from the ground to the differential cantilever cavity motion *must* be reduced in order to have a hope of measuring thermal noise in the band from 50 Hz to 1 kHz. This can be done in at least two ways:

1. **Extra isolation stage:** Most straightforwardly, an additional isolation stage with a resonant frequency at or below 10 Hz will vastly improve the visibility of thermal noise in the region of interest. This can take the form of a nested suspension stage mounted onto the suspended breadboard, from which another breadboard would hang just above, or it could be a spring-based isolation stage between the table and the cryostat.
2. **Symmetric cavities:** The performance could also be improved by increasing the mechanical common-mode rejection between the motion of both mirrors *within each cavity*. Since motion of the payload couples to cantilever mirror motion—relative to the macroscopic mirror with which it forms its




Figure 11.20: Attempt to explain the elevated noise level using anomalously low- Q higher-order bending modes. In this model, the loss of the second bending mode at around 500 Hz is scaled to be 30 times higher than the 10^{-5} of the fundamental mode, while that for the mode around 3 kHz has been scaled by 10000(!).

cavity—without suppression above the cantilever resonance, there is only common-mode rejection *between* the two cavities, and this is limited due to imperfections in the placement of the optical elements and the resting position of the suspension. On the other hand, if both mirrors in each cavity were attached to cantilevers, their motion relative to each other would be greatly reduced by comparison.

Option (1) is quite feasible with some modifications to the mechanical construction of the experiment. Option (2) is complicated by the difficulty in obtaining sufficiently small cavity mirrors with the required radii of curvature to make the cavity stable. Both of these areas are very worthy of further study.

Additionally, there is clearly much more interest in making this measurement when the source of the thermal displacement noise is not the rather large thermoelastic damping at room temperature. Therefore, this experiment should be conducted at low temperatures (at 125 K or below), where the off-resonance thermal displacement noise from bulk and surface effects can be studied directly, thereby corroborating models generated from existing ringdown measurements.

As the required displacement noise sensitivity only scales with the square root of the loss, it is likely that any of the improvements above would already lead to a test platform that is capable of measuring noise associated with interesting levels of loss on the order of $10^{-7} - 10^{-6}$.

BIBLIOGRAPHY

- [1] A. Einstein. Die grundlage der allgemeinen relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916. doi: 10.1002/andp.19163540702. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19163540702>.
- [2] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. Physics Series. W. H. Freeman, San Francisco, first edition edition, September 1973. ISBN 0716703440. URL <http://www.worldcat.org/isbn/0716703440>.
- [3] P. R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. WORLD SCIENTIFIC, 2017. doi: 10.1142/10116. URL <https://www.worldscientific.com/doi/abs/10.1142/10116>.
- [4] A. Einstein. Über gravitationswellen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, Seite 154-167., 1918.
- [5] P. R. Saulson. Josh Goldberg and the physical reality of gravitational waves. *General Relativity and Gravitation*, 43(12):3289–3299, Dec 2011. ISSN 1572-9532. doi: 10.1007/s10714-011-1237-z. URL <https://doi.org/10.1007/s10714-011-1237-z>.
- [6] H. Hertz. *Electric Waves*. Macmillan, 1893.
- [7] J. Weber. Detection and generation of gravitational waves. *Phys. Rev.*, 117: 306–313, Jan 1960. doi: 10.1103/PhysRev.117.306. URL <https://link.aps.org/doi/10.1103/PhysRev.117.306>.
- [8] J. A. Tyson and R. P. Giffard. Gravitational-wave astronomy. *Annual Review of Astronomy and Astrophysics*, 16(1):521–554, 1978. doi: 10.1146/annurev.aa.16.090178.002513. URL <https://doi.org/10.1146/annurev.aa.16.090178.002513>.
- [9] L. Ju, D. G. Blair, and C. Zhao. Detection of gravitational waves. *Reports on Progress in Physics*, 63(9):1317, 2000. URL <http://stacks.iop.org/0034-4885/63/i=9/a=201>.
- [10] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar psr1913 + 16. *Nature*, 277:437 EP –, 02 1979. URL <https://doi.org/10.1038/277437a0>.
- [11] M. Kamionkowski and E. D. Kovetz. The quest for b modes from inflationary gravitational waves. *Annual Review of Astronomy and Astrophysics*, 54(1): 227–269, 2018/11/05 2016. doi: 10.1146/annurev-astro-081915-023433. URL <https://doi.org/10.1146/annurev-astro-081915-023433>.

- [12] M. V. Sazhin. Opportunities for detecting ultralong gravitational waves. *Sov. Astron.*, 22:36–38, February 1978.
- [13] S. Detweiler. Pulsar timing measurements and the search for gravitational waves. *Astrophys. J.*, 234:1100–1104, December 1979. doi: 10.1086/157593.
- [14] F. B. Estabrook and H. D. Wahlquist. Response of doppler spacecraft tracking to gravitational radiation. *General Relativity and Gravitation*, 6:439–447, October 1975. doi: 10.1007/BF00762449.
- [15] R. X. Adhikari. Gravitational radiation detection with laser interferometry. *Rev. Mod. Phys.*, 86:121–151, Feb 2014. doi: 10.1103/RevModPhys.86.121. URL <https://link.aps.org/doi/10.1103/RevModPhys.86.121>.
- [16] R. Weiss. Electromagnetically coupled broadband gravitational antenna. Technical report, Massachusetts Institute of Technology, <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=38618>, 1972.
- [17] B. C. Barish and R. Weiss. LIGO and the Detection of Gravitational Waves. *Phys. Today*, 52(10):44, 1999. URL <https://doi.org/10.1063/1.882861>.
- [18] B. P. Abbott, R. Abbott, R. Adhikari, et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901, 2009. URL <http://stacks.iop.org/0034-4885/72/i=7/a=076901>.
- [19] R. Takahashi, K. Arai, D. Tatsumi, et al. Operational status of TAMA300 with the seismic attenuation system (SAS). *Classical and Quantum Gravity*, 25(11):114036, 2008. URL <http://stacks.iop.org/0264-9381/25/i=11/a=114036>.
- [20] T. Accadia, F. Acernese, M. Alshourbagy, et al. Virgo: a laser interferometer to detect gravitational waves. *Journal of Instrumentation*, 7(03):P03012, 2012. URL <http://stacks.iop.org/1748-0221/7/i=03/a=P03012>.
- [21] H. Grote and the LIGO Scientific Collaboration. The GEO 600 status. *Classical and Quantum Gravity*, 27(8):084003, 2010. URL <http://stacks.iop.org/0264-9381/27/i=8/a=084003>.
- [22] T. T. Fricke, N. D. Smith-Lefebvre, R. Abbott, et al. DC readout experiment in Enhanced LIGO. *Classical and Quantum Gravity*, 29(6):065005, 2012. URL <http://stacks.iop.org/0264-9381/29/i=6/a=065005>.
- [23] The LIGO Scientific Collaboration, J. Aasi, B. P. Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015. URL <http://stacks.iop.org/0264-9381/32/i=7/a=074001>.

- [24] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016. doi: 10.1103/PhysRevLett.116.061102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [25] The LIGO Scientific Collaboration and the VIRGO Collaboration. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. 2018. URL <https://arxiv.org/abs/1811.12907>.
- [26] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017. doi: 10.1103/PhysRevLett.119.161101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
- [27] D. V. Martynov, E. D. Hall, B. P. Abbott, et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93:112004, Jun 2016. doi: 10.1103/PhysRevD.93.112004. URL <https://link.aps.org/doi/10.1103/PhysRevD.93.112004>.
- [28] A. E. Siegman. *Lasers*. University Science Books, 1986.
- [29] J. Mizuno, K. Strain, P. Nelson, et al. Resonant sideband extraction: a new configuration for interferometric gravitational wave detectors. *Physics Letters A*, 175(5):273 – 276, 1993. ISSN 0375-9601. doi: [https://doi.org/10.1016/0375-9601\(93\)90620-F](https://doi.org/10.1016/0375-9601(93)90620-F). URL <http://www.sciencedirect.com/science/article/pii/037596019390620F>.
- [30] A. Buonanno and Y. Chen. Signal recycled laser-interferometer gravitational-wave detectors as optical springs. *Phys. Rev. D*, 65:042001, Jan 2002. doi: 10.1103/PhysRevD.65.042001. URL <https://link.aps.org/doi/10.1103/PhysRevD.65.042001>.
- [31] H. Miao. *Exploring Macroscopic Quantum Mechanics in Optomechanical Devices*. PhD thesis, University of Western Australia, 2010.
- [32] C. L. Mueller, M. A. Arain, G. Ciani, et al. The Advanced LIGO input optics. *Review of Scientific Instruments*, 87(1):014502, 2016. doi: 10.1063/1.4936974. URL <https://doi.org/10.1063/1.4936974>.
- [33] C. Zhao, J. Degallaix, L. Ju, et al. Compensation of Strong Thermal Lensing in High-Optical-Power Cavities. *Phys. Rev. Lett.*, 96:231101, Jun 2006. doi: 10.1103/PhysRevLett.96.231101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.96.231101>.
- [34] D. V. Martynov. *Lock Acquisition and Sensitivity Analysis of Advanced LIGO Interferometers*. PhD thesis, California Institute of Technology, 2015.

- [35] E. D. Hall. *Long-baseline laser interferometry for the detection of binary black-hole mergers*. PhD thesis, California Institute of Technology, 2017.
- [36] E. A. Quintero. *Improving the Performance and Sensitivity of Gravitational Wave Detectors*. PhD thesis, California Institute of Technology, 2018.
- [37] P. Kwee, C. Bogan, K. Danzmann, et al. Stabilized high-power laser system for the gravitational wave detector Advanced LIGO. *Opt. Express*, 20(10):10617–10634, 2012. URL <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-20-10-10617>.
- [38] R. W. P. Drever, J. L. Hall, F. V. Kowalski, et al. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(2):97–105, Jun 1983. ISSN 1432-0649. doi: 10.1007/BF00702605. URL <https://doi.org/10.1007/BF00702605>.
- [39] J. C. Driggers, J. Harms, and R. X. Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Phys. Rev. D*, 86:102001, Nov 2012. doi: 10.1103/PhysRevD.86.102001. URL <https://link.aps.org/doi/10.1103/PhysRevD.86.102001>.
- [40] J. C. Driggers. *Noise Cancellation for Gravitational Wave Detectors*. PhD thesis, California Institute of Technology, 2015.
- [41] D. Ugolini, C. Fitzgerald, I. Rothbarth, and J. Wang. Discharging fused silica optics occluded by an electrostatic drive. *Review of Scientific Instruments*, 85(3):034502, 2014. doi: 10.1063/1.4867248. URL <https://doi.org/10.1063/1.4867248>.
- [42] C. M. Caves. Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer. *Phys. Rev. Lett.*, 45:75–79, Jul 1980. doi: 10.1103/PhysRevLett.45.75. URL <https://link.aps.org/doi/10.1103/PhysRevLett.45.75>.
- [43] V. B. Braginskiĭ and Y. I. Vorontsov. Quantum-mechanical limitations in macroscopic experiments and modern experimental technique. *Soviet Physics Uspekhi*, 17(5):644, 1975. URL <http://stacks.iop.org/0038-5670/17/i=5/a=R02>.
- [44] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics. *Phys. Rev. D*, 65:022002, Dec 2001. doi: 10.1103/PhysRevD.65.022002. URL <https://link.aps.org/doi/10.1103/PhysRevD.65.022002>.
- [45] R. B. F. H. M. . R. A. V.P.L.S. Xxvii. a brief account of microscopical observations made in the months of june, july and august 1827, on the particles contained in the pollen of plants; and on the general existence

- of active molecules in organic and inorganic bodies. *The Philosophical Magazine*, 4(21):161–173, 1828. doi: 10.1080/14786442808674769. URL <https://doi.org/10.1080/14786442808674769>.
- [46] A. Einstein. Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der Physik*, 322(8):549–560, 1905. doi: 10.1002/andp.19053220806. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19053220806>.
- [47] Anonymous. Minutes of the philadelphia meeting december 28, 29, 30, 1926. *Phys. Rev.*, 29:350–373, Feb 1927. doi: 10.1103/PhysRev.29.350. URL <https://link.aps.org/doi/10.1103/PhysRev.29.350>.
- [48] H. Nyquist. Thermal Agitation of Electric Charge in Conductors. *Phys. Rev.*, 32:110–113, Jul 1928. doi: 10.1103/PhysRev.32.110. URL <https://link.aps.org/doi/10.1103/PhysRev.32.110>.
- [49] H. B. Callen and T. A. Welton. Irreversibility and Generalized Noise. *Phys. Rev.*, 83:34–40, Jul 1951. doi: 10.1103/PhysRev.83.34. URL <https://link.aps.org/doi/10.1103/PhysRev.83.34>.
- [50] B. Lantz, S. Danilishin, S. Hild, et al. Instrument Science White Paper 2018. Technical Report T1800133-V4, LIGO Laboratory, 2018. URL <https://dcc.ligo.org/LIGO-T1800133/public>.
- [51] R. X. Adhikari, N. Smith, A. Brooks, et al. LIGO Voyager Upgrade: Design Concept. Technical Report T1400226-v9, LIGO Laboratory, 2014. URL <https://dcc.ligo.org/LIGO-T1400226>.
- [52] C. M. Caves. Quantum-mechanical noise in an interferometer. *Phys. Rev. D*, 23:1693–1708, Apr 1981. doi: 10.1103/PhysRevD.23.1693. URL <https://link.aps.org/doi/10.1103/PhysRevD.23.1693>.
- [53] K. Goda, O. Miyakawa, E. E. Mikhailov, et al. A quantum-enhanced prototype gravitational-wave detector. *Nature Physics*, 4:472 EP –, 03 2008. URL <https://doi.org/10.1038/nphys920>.
- [54] H. Grote, K. Danzmann, K. L. Dooley, et al. First long-term application of squeezed states of light in a gravitational-wave observatory. *Phys. Rev. Lett.*, 110:181101, May 2013. doi: 10.1103/PhysRevLett.110.181101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.110.181101>.
- [55] T. L. S. Collaboration, J. Abadie, B. P. Abbott, et al. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nature Physics*, 7:962 EP –, 09 2011. URL <https://doi.org/10.1038/nphys2083>.

- [56] E. Oelker, T. Isogai, J. Miller, et al. Audio-band frequency-dependent squeezing for gravitational-wave detectors. *Phys. Rev. Lett.*, 116:041102, Jan 2016. doi: 10.1103/PhysRevLett.116.041102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.041102>.
- [57] J. Steinlechner, I. W. Martin, A. S. Bell, et al. Silicon-based optical mirror coatings for ultrahigh precision metrology and sensing. *Phys. Rev. Lett.*, 120:263602, Jun 2018. doi: 10.1103/PhysRevLett.120.263602. URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.263602>.
- [58] W. Z. Korth, A. Heptonstall, E. D. Hall, et al. Passive, free-space heterodyne laser gyroscope. *Classical and Quantum Gravity*, 33(3):035004, 2016. URL <http://stacks.iop.org/0264-9381/33/i=3/a=035004>.
- [59] B. Lantz, R. Schofield, B. O'Reilly, D. E. Clark, and D. DeBra. Review: Requirements for a ground rotation sensor to improve advanced ligo. *Bulletin of the Seismological Society of America*, 99(2B):980–989, 2009. doi: 10.1785/0120080199. URL <http://www.bssaonline.org/content/99/2B/980.abstract>.
- [60] K. Venkateswara, C. A. Hagedorn, M. D. Turner, T. Arp, and J. H. Gundlach. A high-precision mechanical absolute-rotation sensor. *Review of Scientific Instruments*, 85(1):015005, 2014. doi: 10.1063/1.4862816. URL <https://doi.org/10.1063/1.4862816>.
- [61] M. P. Ross, K. Venkateswara, C. A. Hagedorn, et al. Low-Frequency Tilt Seismology with a Precision Ground-Rotation Sensor. *Seismological Research Letters*, 89(1):67, 2017. doi: 10.1785/0220170148. URL <http://dx.doi.org/10.1785/0220170148>.
- [62] W. M. Macek and J. D. T. M. Davis. Rotation rate sensing with traveling-wave ring lasers. *Applied Physics Letters*, 2(3):67–68, 1963. doi: 10.1063/1.1753778. URL <http://link.aip.org/link/?APL/2/67/1>.
- [63] G. Sagnac. *Compt. Rend.*, 157(708), 1913.
- [64] W. W. Chow, J. Gea-Banacloche, L. M. Pedrotti, et al. The ring laser gyro. *Rev. Mod. Phys.*, 57:61–104, Jan 1985. doi: 10.1103/RevModPhys.57.61. URL <http://link.aps.org/doi/10.1103/RevModPhys.57.61>.
- [65] S. Ezekiel and S. R. Balsamo. Passive ring resonator laser gyroscope. *Applied Physics Letters*, 30(9):478–480, 1977. doi: 10.1063/1.89455. URL <http://link.aip.org/link/?APL/30/478/1>.
- [66] J. Belfi, N. Beverini, F. Bosi, et al. Rotational sensitivity of the G-Pisa gyrolaser. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 57(3):618–622, 2010. ISSN 0885-3010. doi: 10.1109/TUFFC.2010.1456.

- [67] F. Aronowitz. *Laser Applications*, pages 134–200. Academic, New York, 1971.
- [68] F. Zarinetchi and S. Ezekiel. Observation of lock-in behavior in a passive resonator gyroscope. *Opt. Lett.*, 11(6):401–403, Jun 1986. doi: 10.1364/OL.11.000401. URL <http://ol.osa.org/abstract.cfm?URI=ol-11-6-401>.
- [69] F. Zarinetchi. *Studies in optical resonator gyroscopes*. PhD thesis, Massachusetts Institute of Technology, May 1992. URL <http://dspace.mit.edu/handle/1721.1/13227>.
- [70] G. E. Stedman. Ring-laser tests of fundamental physics and geophysics. *Reports on Progress in Physics*, 60(6):615, 1997. URL <http://stacks.iop.org/0034-4885/60/i=6/a=001>.
- [71] K. U. Schreiber, T. Klügel, and G. E. Stedman. Earth tide and tilt detection by a ring laser gyroscope. *Journal of Geophysical Research: Solid Earth*, 108(B2), 2003. ISSN 2156-2202. doi: 10.1029/2001JB000569. URL <http://dx.doi.org/10.1029/2001JB000569>.
- [72] K. U. Schreiber, A. Velikoseltsev, M. Rothacher, et al. Direct measurement of diurnal polar motion by ring laser gyroscopes. *Journal of Geophysical Research: Solid Earth*, 109(B6), 2004. ISSN 2156-2202. doi: 10.1029/2003JB002803. URL <http://dx.doi.org/10.1029/2003JB002803>.
- [73] C. H. Rowe, U. K. Schreiber, S. J. Cooper, et al. Design and operation of a very large ring laser gyroscope. *Appl. Opt.*, 38(12):2516–2523, Apr 1999. doi: 10.1364/AO.38.002516. URL <http://ao.osa.org/abstract.cfm?URI=ao-38-12-2516>.
- [74] R. W. Dunn, D. E. Shabalin, R. J. Thirkettle, et al. Design and initial operation of a 367-m² rectangular ring laser. *Appl. Opt.*, 41(9):1685–1688, Mar 2002. doi: 10.1364/AO.41.001685. URL <http://ao.osa.org/abstract.cfm?URI=ao-41-9-1685>.
- [75] R. B. Hurst, G. E. Stedman, K. U. Schreiber, et al. Experiments with an 834 m² ring laser interferometer. *Journal of Applied Physics*, 105(11):113115, 2009. doi: 10.1063/1.3133245. URL <http://link.aip.org/link/?JAP/105/113115/1>.
- [76] *GG1320AN Digital Laser Gyro*. Honeywell Aerospace, Phoenix, AZ, USA.
- [77] *RGA-20 High Performance Gyro*. L-3 Communications, Budd Lake, New Jersey, USA.
- [78] J. Aasi, B. P. Abbott, R. Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015. URL <http://stacks.iop.org/0264-9381/32/i=7/a=074001>.

- [79] R. DeRosa, J. C. Driggers, D. Atkinson, et al. Global feed-forward vibration isolation in a km scale interferometer. *Classical and Quantum Gravity*, 29(21):215008, 2012. URL <http://stacks.iop.org/0264-9381/29/i=21/a=215008>.
- [80] F. Acernese, M. Agathos, K. Agatsuma, et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, 2015. URL <http://stacks.iop.org/0264-9381/32/i=2/a=024001>.
- [81] Advanced Virgo Baseline Design. Technical Report VIR-027A-09, The Virgo Collaboration, 2009. URL <https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf>.
- [82] J. Belfi, N. Beverini, F. Bosi, et al. Performance of G-Pisa ring laser gyro at the Virgo site. *Journal of Seismology*, 16(4):757–766, 2012. ISSN 1383-4649. doi: 10.1007/s10950-012-9277-8. URL <http://dx.doi.org/10.1007/s10950-012-9277-8>.
- [83] N. Robertson, R. Drever, I. Kerr, and J. Hough. Passive and active seismic isolation for gravitational radiation detectors and other instruments. *Journal of Physics E: Scientific Instruments*, 15(10):1101, 1982.
- [84] A. Giazotto, D. Passuello, and A. Stefanini. One-mile equivalent length interferometric pendulum for seismic noise reduction. *Review of scientific instruments*, 57(6):1145–1151, 1986.
- [85] K. Venkateswara, C. A. Hagedorn, M. D. Turner, T. Arp, and J. H. Gundlach. A high-precision mechanical absolute-rotation sensor. *Review of Scientific Instruments*, 85(1):015005, 2014. doi: <http://dx.doi.org/10.1063/1.4862816>. URL <http://scitation.aip.org/content/aip/journal/rsi/85/1/10.1063/1.4862816>.
- [86] V. Dergachev, R. DeSalvo, M. Asadoor, et al. A high precision, compact electromechanical ground rotation sensor. *Review of Scientific Instruments*, 85(5):054502, 2014. doi: <http://dx.doi.org/10.1063/1.4875375>. URL <http://scitation.aip.org/content/aip/journal/rsi/85/5/10.1063/1.4875375>.
- [87] F. Matichard and M. Evans. Review: Tilt-free low-noise seismometry. *Bulletin of the Seismological Society of America*, 2014.
- [88] S. Merlo, M. Norgia, and S. Donati. Fiber gyroscope principles. *Handbook of optical fibre sensing technology ed. JM Lopez-Higuera (John Wiley and Sons, NY, 2002)*, pages 331–347, 2000.
- [89] T. L. Gustavson, P. Bouyer, and M. A. Kasevich. Precision Rotation Measurements with an Atom Interferometer Gyroscope. *Physical Review Letters*, 78(1):2046–2049, March 1997.

- [90] T. L. Gustavson. *Precision rotation sensing using atom interferometry*. PhD thesis, stanford university, 2000.
- [91] D. Durfee, Y. Shaham, and M. Kasevich. Long-term stability of an area-reversible atom-interferometer sagnac gyroscope. *Physical Review Letters*, 97(240801), 2006.
- [92] K. Schwab, N. Bruckner, and R. Packard. Detection of the earth’s rotation using superfluid phase coherence. *Nature*, 386:585–587, 1997.
- [93] O. Avenel, P. Hakonen, and E. Varoquaux. Detection of the rotation of the earth with a superfluid gyrometer. *Phys. Rev. Lett.*, 78:3602–3605, May 1997. doi: 10.1103/PhysRevLett.78.3602. URL <http://link.aps.org/doi/10.1103/PhysRevLett.78.3602>.
- [94] G. A. Sanders, M. G. Prentiss, and S. Ezekiel. Passive ring resonator method for sensitive inertial rotation measurements in geophysics and relativity. *Optics Letters*, 6(11):569–571, 1981.
- [95] C. Hadziioannou, P. Gaebler, U. Schreiber, J. Wassermann, and H. Igel. Examining ambient noise using colocated measurements of rotational and translational motion. *Journal of seismology*, 16(4):787–796, 2012.
- [96] K. U. Schreiber and J.-P. R. Wells. Invited review article: Large ring lasers for rotation sensing. *Review of Scientific Instruments*, 84(4):041101, 2013. doi: <http://dx.doi.org/10.1063/1.4798216>. URL <http://scitation.aip.org/content/aip/journal/rsi/84/4/10.1063/1.4798216>.
- [97] S. W. Lloyd, S. Fan, and M. J. F. Digonnet. Experimental Observation of Low Noise and Low Drift in a Laser-Driven Fiber Optic Gyroscope. *Journal of Lightwave Technology*, 31(13):2079–2085, July 2013.
- [98] J. Chamoun, A. Evans, F. Mosca, and M. Digonnet. Low noise and low drift in a laser-driven fiber optic gyroscope with a 1-km coil. In *Proc. of SPIE Vol.*, volume 9157, pages 91570E–1, 2014.
- [99] J. Belfi, N. Beverini, F. Bosi, et al. A 1.82 m² ring laser gyroscope for nano-rotational motion sensing. *Applied Physics B*, 106(2):271–281, September 2011.
- [100] M. Grewal and A. Andrews. How good is your gyro [ask the experts]. *Control Systems, IEEE*, 30(1):12–86, Feb 2010. ISSN 1066-033X. doi: 10.1109/MCS.2009.935122.
- [101] R. W. P. Drever, J. L. Hall, F. V. Kowalski, et al. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31:97–105, 1983.

- [102] B. Willke, N. Uehara, E. K. Gustafson, et al. Spatial and temporal filtering of a 10-w nd:yag laser with a fabry–perot ring-cavity premode cleaner. *Opt. Lett.*, 23(21):1704–1706, Nov 1998. doi: 10.1364/OL.23.001704. URL <http://ol.osa.org/abstract.cfm?URI=ol-23-21-1704>.
- [103] K. Kokeyama, K. Izumi, W. Z. Korth, et al. Residual amplitude modulation in interferometric gravitational wave detectors. *J. Opt. Soc. Am. A*, 31(1): 81–88, Jan 2014. doi: 10.1364/JOSAA.31.000081. URL <http://josaa.osa.org/abstract.cfm?URI=josaa-31-1-81>.
- [104] N. Wong and J. L. Hall. Servo control of amplitude modulation in frequency-modulation spectroscopy: demonstration of shot-noise-limited detection. *JOSA B*, 2(9):1527–1533, 1985.
- [105] E. Jaatinen, D. J. Hopper, and J. Back. Residual amplitude modulation mechanisms in modulation transfer spectroscopy that use electro-optic modulators. *Measurement Science and Technology*, 20(2):025302, February 2009. doi: 10.1088/0957-0233/20/2/025302.
- [106] H. Grote. High power, low-noise, and multiply resonant photodetector for interferometric gravitational wave detectors. *Review of Scientific Instruments*, 78(5):054704, 2007. doi: 10.1063/1.2735559. URL <http://link.aip.org/link/?RSI/78/054704/1>.
- [107] *Aeroflex 2023 and 2024 AM/FM Signal Generators*. Aeroflex, 1998. URL http://www.aeroflex.com/ats/products/prodfiles/opsmanuals/2023_20240peratingManual.pdf.
- [108] S. R. Systems. Rubidium frequency standard, 2013. URL http://www.aeroflex.com/ats/products/prodfiles/opsmanuals/2023_20240peratingManual.pdf.
- [109] S. Schilt, N. Bucalovic, L. Tombez, et al. Frequency discriminators for the characterization of narrow-spectrum heterodyne beat signals: Application to the measurement of a sub-hertz carrier-envelope-offset beat in an optical frequency comb. *Review of Scientific Instruments*, 82(12):123116, 2011. doi: 10.1063/1.3670357. URL <http://link.aip.org/link/?RSI/82/123116/1>.
- [110] T. Fricke. *Homodyne detection for laser-interferometric gravitational wave detectors*. PhD thesis, Louisiana State University, 2011.
- [111] N. D. Smith-Lefebvre. *Techniques for Improving the Readout Sensitivity of Gravitational Wave Antennae*. PhD thesis, Massachusetts Institute of Technology, 2012.
- [112] K. Arai, S. Barnum, P. Fritschel, J. Lewis, and S. Waldman. Output Mode Cleaner Design. T1000276-v5, LIGO Laboratory, 2013. URL <https://dcc.ligo.org/LIGO-T1000276/public>.

- [113] K. Arai and W. Z. Korth. OMC optical design. Technical Report G1201111-v2, LIGO Laboratory, 2013. URL <https://dcc.ligo.org/LIGO-G1201111/public>.
- [114] W. Z. Korth and K. Arai. OMC DCPD characterization for aLIGO transition. Technical Report T1300552, LIGO Laboratory, 2013. URL <https://dcc.ligo.org/LIGO-T1300552/public>.
- [115] R. Abbott. In-vacuum DCPD for OMC. Technical Report D060572, LIGO Laboratory, 2006. URL <https://dcc.ligo.org/LIGO-D060572/public>.
- [116] L. Barsotti and C. Mueller. Mode matching investigations at LLO. Technical Report G1300909, LIGO Laboratory, 2013. URL <https://dcc.ligo.org/LIGO-G1300909/public>.
- [117] W. Z. Korth, H. Miao, T. Corbitt, et al. Suppression of quantum-radiation-pressure noise in an optical spring. *Phys. Rev. A*, 88:033805, Sep 2013. doi: 10.1103/PhysRevA.88.033805. URL <https://link.aps.org/doi/10.1103/PhysRevA.88.033805>.
- [118] T. J. Kippenberg and K. J. Vahala. Cavity optomechanics: Back-action at the mesoscale. *Science*, 321(5893):1172–1176, August 2008. URL <http://www.sciencemag.org/cgi/content/abstract/321/5893/1172>.
- [119] F. Marquardt and S. M. Girvin. Optomechanics. *Physics*, 2:40, May 2009. doi: 10.1103/Physics.2.40. URL <http://link.aps.org/doi/10.1103/Physics.2.40>.
- [120] M. Aspelmeyer, P. Meystre, and K. Schwab. Quantum optomechanics. *Physics Today*, 65:29, 2012.
- [121] V. B. Braginsky and F. Y. Khalili. *Quantum Measurement*. Cambridge University Press, 1992.
- [122] B. P. Abbott, R. Abbott, R. Adhikari, et al. LIGO: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, 72(7):076901, 2009. URL <http://stacks.iop.org/0034-4885/72/i=7/a=076901>.
- [123] J. D. Teufel, T. Donner, D. Li, et al. Sideband cooling of micromechanical motion to the quantum ground state. *Nature*, 475(7356):359–363, July 2011. ISSN 0028-0836. URL <http://dx.doi.org/10.1038/nature10261>.
- [124] A. D. O’Connell. Quantum ground state and single-phonon control of a mechanical resonator. *Nature*, 464:697–703, 2010. URL <http://dx.doi.org/10.1038/nature08967>.

- [125] J. Chan, T. P. M. Alegre, A. H. Safavi-Naeini, et al. Laser cooling of a nanomechanical oscillator into its quantum ground state. *Nature*, 478:89–92, Oct 2011.
- [126] S. Weis, R. Rivière, S. Deléglise, et al. Optomechanically induced transparency. *Science*, 330(6010):1520–1523, 2010. doi: 10.1126/science.1195596. URL <http://www.sciencemag.org/content/330/6010/1520.abstract>.
- [127] A. H. Safavi-Naeini, T. P. M. Alegre, J. Chan, et al. Electromagnetically induced transparency and slow light with optomechanics. *Nature*, 472(7341): 69–73, Apr 2011.
- [128] K.-J. Boller, A. Imamolu, and S. E. Harris. Observation of electromagnetically induced transparency. *Phys. Rev. Lett.*, 66:2593–2596, May 1991. doi: 10.1103/PhysRevLett.66.2593. URL <http://link.aps.org/doi/10.1103/PhysRevLett.66.2593>.
- [129] S. E. Harris. Electromagnetically induced transparency. *Physics Today*, 50 (7):36–42, 1997. doi: 10.1063/1.881806. URL <http://link.aip.org/link/?PTO/50/36/1>.
- [130] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics. *Phys. Rev. D*, 65:022002, Dec 2001. doi: 10.1103/PhysRevD.65.022002. URL <http://link.aps.org/doi/10.1103/PhysRevD.65.022002>.
- [131] Y. Ma, S. L. Danilishin, C. Zhao, et al. Narrowing the filter-cavity bandwidth in gravitational-wave detectors via optomechanical interaction. *Phys. Rev. Lett.*, 113:151102, Oct 2014. doi: 10.1103/PhysRevLett.113.151102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.113.151102>.
- [132] P. R. Saulson. Thermal noise in mechanical experiments. *Phys. Rev. D*, 42:2437–2445, Oct 1990. doi: 10.1103/PhysRevD.42.2437. URL <https://link.aps.org/doi/10.1103/PhysRevD.42.2437>.
- [133] D. F. McGuigan, C. C. Lam, R. Q. Gram, et al. Measurements of the mechanical q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30:621–629, 1978. ISSN 0022-2291. URL <http://dx.doi.org/10.1007/BF00116202>. 10.1007/BF00116202.
- [134] A. Ageev, B. C. Palmer, A. D. Felice, S. D. Penn, and P. R. Saulson. Very high quality factor measured in annealed fused silica. *Classical and Quantum Gravity*, 21(16):3887, 2004. URL <http://stacks.iop.org/0264-9381/21/i=16/a=004>.

- [135] R. Mihailovich and N. MacDonald. Dissipation measurements of vacuum-operated single-crystal silicon microresonators. *Sensors and Actuators A: Physical*, 50(3):199 – 207, 1995. ISSN 0924-4247. doi: 10.1016/0924-4247(95)01080-7. URL <http://www.sciencedirect.com/science/article/pii/0924424795010807>.
- [136] G. D. Cole, I. Wilson-Rae, K. Werbach, M. R. Vanner, and M. Aspelmeyer. Phonon-tunnelling dissipation in mechanical resonators. *Nat. Commun.*, 2: 231, March 2011.
- [137] C. Zener. Internal Friction in Solids II. General Theory of Thermoelastic Internal Friction. *Phys. Rev.*, 53:90–99, Jan 1938. doi: 10.1103/PhysRev.53.90. URL <https://link.aps.org/doi/10.1103/PhysRev.53.90>.
- [138] A. Akhiezer. *J. Phys. (Moscow)*, 1:277, 1939.
- [139] G. M. Harry, A. M. Gretarsson, P. R. Saulson, et al. Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings. *Classical and Quantum Gravity*, 19(5):897, 2002. URL <http://stacks.iop.org/0264-9381/19/i=5/a=305>.
- [140] V. Braginsky, M. Gorodetsky, and F. Khalili. Optical bars in gravitational wave antennas. *Physics Letters A*, 232(5):340–348, August 1997. ISSN 0375-9601. doi: 10.1016/S0375-9601(97)00413-1. URL <http://www.sciencedirect.com/science/article/pii/S0375960197004131>.
- [141] V. Braginsky and F. Khalili. Low noise rigidity in quantum measurements. *Physics Letters A*, 257(5-6):241 – 246, 1999. ISSN 0375-9601. doi: 10.1016/S0375-9601(99)00337-0. URL <http://www.sciencedirect.com/science/article/pii/S0375960199003370>.
- [142] A. Buonanno and Y. Chen. Signal recycled laser-interferometer gravitational-wave detectors as optical springs. *Phys. Rev. D*, 65(4):042001–, January 2002. URL <http://link.aps.org/doi/10.1103/PhysRevD.65.042001>.
- [143] T. Corbitt, Y. Chen, E. Innerhofer, et al. An all-optical trap for a gram-scale mirror. *Phys. Rev. Lett.*, 98:150802, Apr 2007. doi: 10.1103/PhysRevLett.98.150802. URL <http://link.aps.org/doi/10.1103/PhysRevLett.98.150802>.
- [144] T. Corbitt, C. Wipf, T. Bodiya, et al. Optical dilution and feedback cooling of a gram-scale oscillator to 6.9 mk. *Phys. Rev. Lett.*, 99:160801, Oct 2007. doi: 10.1103/PhysRevLett.99.160801. URL <http://link.aps.org/doi/10.1103/PhysRevLett.99.160801>.
- [145] A. Di Virgilio, L. Barsotti, S. Braccini, et al. Experimental evidence for an optical spring. *Phys. Rev. A*, 74:013813, Jul 2006. doi: 10.1103/PhysRevA.74.013813. URL <http://link.aps.org/doi/10.1103/PhysRevA.74.013813>.

- [146] D. E. Chang, K.-K. Ni, O. Painter, and H. J. Kimble. Ultrahigh- q mechanical oscillators through optical trapping. *New Journal of Physics*, 14(4):045002, 2012. URL <http://stacks.iop.org/1367-2630/14/i=4/a=045002>.
- [147] K.-K. Ni, R. Norte, D. J. Wilson, et al. Enhancement of mechanical q factors by optical trapping. *Phys. Rev. Lett.*, 108:214302, May 2012. doi: 10.1103/PhysRevLett.108.214302. URL <http://link.aps.org/doi/10.1103/PhysRevLett.108.214302>.
- [148] J. D. Thompson, B. M. Zwickl, A. M. Jayich, et al. Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane. *Nature*, 452(7183): 72–75, March 2008.
- [149] F. Marquardt, J. P. Chen, A. A. Clerk, and S. M. Girvin. Quantum theory of cavity-assisted sideband cooling of mechanical motion. *Phys. Rev. Lett.*, 99(9):093902–, August 2007. URL <http://link.aps.org/doi/10.1103/PhysRevLett.99.093902>.
- [150] S. Reid, G. Cagnoli, D. Crooks, et al. Mechanical dissipation in silicon flexures. *Physics Letters A*, 351(4):205 – 211, 2006. ISSN 0375-9601. doi: <https://doi.org/10.1016/j.physleta.2005.10.103>. URL <http://www.sciencedirect.com/science/article/pii/S0375960105016981>.
- [151] R. Nawrodt, C. Schwarz, S. Kroker, et al. Investigation of mechanical losses of thin silicon flexures at low temperatures. *Classical and Quantum Gravity*, 30(11):115008, 2013. URL <http://stacks.iop.org/0264-9381/30/i=11/a=115008>.
- [152] L. G. Prokhorov and V. P. Mitrofanov. Mechanical losses of oscillators fabricated in silicon wafers. *Classical and Quantum Gravity*, 32(19):195002, 2015. URL <http://stacks.iop.org/0264-9381/32/i=19/a=195002>.
- [153] C. Zener. Internal Friction in Solids. I. Theory of Internal Friction in Reeds. *Phys. Rev.*, 52:230–235, Aug 1937. doi: 10.1103/PhysRev.52.230. URL <https://link.aps.org/doi/10.1103/PhysRev.52.230>.
- [154] H. E. Bömmel and K. Dransfeld. Excitation and Attenuation of Hypersonic Waves in Quartz. *Phys. Rev.*, 117:1245–1252, Mar 1960. doi: 10.1103/PhysRev.117.1245. URL <https://link.aps.org/doi/10.1103/PhysRev.117.1245>.
- [155] D. F. McGuigan, C. C. Lam, R. Q. Gram, et al. Measurements of the mechanical q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30(5):621–629, Mar 1978. ISSN 1573-7357. doi: 10.1007/BF00116202. URL <https://doi.org/10.1007/BF00116202>.
- [156] J. Rodriguez, S. A. Chandorkar, C. A. Watson, et al. Direct detection of akhiezer damping in a silicon mems resonator. *Scientific Reports*, 9(1):2244,

2019. doi: 10.1038/s41598-019-38847-6. URL <https://doi.org/10.1038/s41598-019-38847-6>.
- [157] W. Z. Korth. What ϕ is needed for the Voyager suspension? Technical Report T1500505-v1, LIGO Laboratory, 2015. URL <https://dcc.ligo.org/LIGO-T1500505>.
- [158] B. Shapiro, D. Madden-Fong, and B. Lantz. LIGO III Quad Pendulum Conceptual Design Optimization. Technical Report T1300786-v8, LIGO Laboratory, 2014. URL <https://dcc.ligo.org/LIGO-T1300786>.
- [159] A. V. Cumming, L. Cunningham, G. D. Hammond, et al. Silicon mirror suspensions for gravitational wave detectors. *Classical and Quantum Gravity*, 31(2):025017, dec 2013. doi: 10.1088/0264-9381/31/2/025017. URL <https://doi.org/10.1088/0264-9381/31/2/025017>.
- [160] G. González. Suspensions thermal noise in the LIGO gravitational wave detector. *Classical and Quantum Gravity*, 17(21):4409–4435, oct 2000. doi: 10.1088/0264-9381/17/21/305. URL <https://doi.org/10.1088/0264-9381/17/21/305>.
- [161] D. F. McGuigan, C. C. Lam, R. Q. Gram, et al. Measurements of the mechanical q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30(5):621–629, Mar 1978. ISSN 1573-7357. doi: 10.1007/BF00116202. URL <https://doi.org/10.1007/BF00116202>.
- [162] S. Chao, J.-s. Ou, V. Huang, et al. Progress of coating development at NTHU. Technical Report G1200849-v1, National Tsing Hua University, 2012. URL <https://dcc.ligo.org/LIGO-G1200849/public>.
- [163] W. Kern. The evolution of silicon wafer cleaning technology. *Journal of The Electrochemical Society*, 137(6):1887–1892, 1990. doi: 10.1149/1.2086825. URL <http://jes.ecsdl.org/content/137/6/1887.abstract>.
- [164] J. S. Fakonas. *Quantum Interference and Entanglement of Surface Plasmons*. PhD thesis, California Institute of Technology, 2015.
- [165] M. D. Henry. *ICP Etching of Silicon for Micro and Nanoscale Devices*. PhD thesis, California Institute of Technology, 2010.
- [166] N. D. Smith. A technique for continuous measurement of the quality factor of mechanical oscillators. *Review of Scientific Instruments*, 86(5):053907, 2015. doi: 10.1063/1.4920922. URL <https://doi.org/10.1063/1.4920922>.
- [167] C. Wipf. Noise budgeting for advanced detectors. Technical Report G1400587, LIGO Laboratory, 2014.
- [168] RIO PLANEX. URL http://www.rio-lasers.com/_products/planex.html.

- [169] R. Abbott. High Stability Low Noise Laser Current Driver PCB. Technical Report D1200719-v6, LIGO Laboratory, 2012. URL <https://dcc.ligo.org/LIGO-D1200719/public>.
- [170] T. Sekiguchi. Sumcon. URL <https://granite.phys.s.u-tokyo.ac.jp/svn/LCGT/trunk/VIS/sumcon/>.
- [171] *Red Pitaya*. URL <https://www.redpitaya.com>.
- [172] *Valon 3010*. URL <https://www.valonrf.com/frequency-divider-2ghz.html>.
- [173] *First Contact*. URL <https://www.photoniccleaning.com>.
- [174] H. Miao, S. Danilishin, H. Müller-Ebhardt, and Y. Chen. Achieving ground state and enhancing optomechanical entanglement by recovering information. *New Journal of Physics*, 12(8):083032, 2010. URL <http://stacks.iop.org/1367-2630/12/i=8/a=083032>.
- [175] I. Wilson-Rae, N. Nooshi, W. Zwerger, and T. J. Kippenberg. Theory of ground state cooling of a mechanical oscillator using dynamical backaction. *Phys. Rev. Lett.*, 99(9):093901–, August 2007. URL <http://link.aps.org/doi/10.1103/PhysRevLett.99.093901>.
- [176] C. Genes, D. Vitali, P. Tombesi, S. Gigan, and M. Aspelmeyer. Ground-state cooling of a micromechanical oscillator: Comparing cold damping and cavity-assisted cooling schemes. *Phys. Rev. A*, 77(3):033804–, March 2008. URL <http://link.aps.org/doi/10.1103/PhysRevA.77.033804>.
- [177] G. J. Milburn and M. J. Woolley. An introduction to quantum optomechanics. *Acta Physica Slovaca*, 61:483–601, 2011.
- [178] D. F. Walls and G. Milburn. *Quantum Optics*. Springer-Verlag, 2008.

Appendix A

FEEDBACK CONTROL

A.1 Overview

Feedback is a powerful technique for the control of dynamical systems. At its most basic, feedback consists of taking the output of a system, conditioning it, and then re-injecting it into the system's input. Using an appropriate choice of conditioning, a **servomechanism** (or simply **servo**) can be built to maintain a system in a desired state. The basic structure of a generic feedback control loop is shown in Fig. A.1

The fundamental units of a feedback loop are:

- **Plant:** The physical system to be controlled. In our simple example, it is a single-input, single-output (SISO) system with input units of X and output units of Y . In general, the output Y can be a nonlinear function of X , in which case an important aspect of the feedback control is the **linearization** of the system near some chosen operating point.
- **Sensor:** A transducer used to convert the output of the physical system (in units of Y here) to a useful unit for conditioning. For an analog system, as in our case, this is typically volts; for a digital system, this would be digital unit counts.
- **Servo filter:** The element of the loop that performs the signal conditioning. The transfer function of the servo filter, H , is chosen based on the circumstances to achieve the desired loop performance and stability.
- **Actuator:** A transducer used to drive the plant (i.e., to convert the output of the servo filter in V into units of the plant input in X).

To understand the basic function of the loop, consider first the system without the control loop attached, as shown in the dashed box in Fig. A.1, and assume for now that $\delta Y = 0$. In this case, the output of the plant is

$$Y_p = P \delta X, \tag{A.1}$$

and it is freely evolving under the influence of the noise source δX . Now, if we connect the feedback loop, and again only consider the one noise term δX , the plant

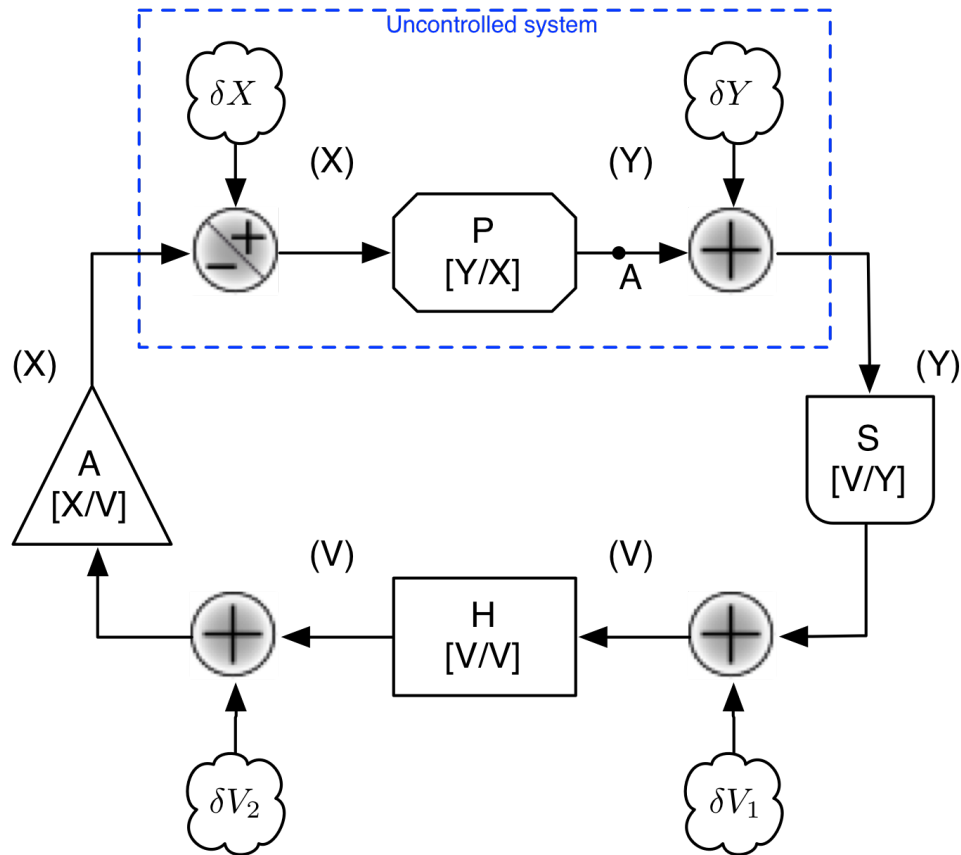


Figure A.1: Diagram of a generic feedback loop, showing the fundamental blocks: plant (P), sensor (S), servo filter (H), and actuator (A). The conversion gain units of each are shown in brackets, and the units of the signal at each point along the loop are shown in parentheses (“V” is meant to indicate volts, while “X” and “Y” represent two generic physical units). The elements in cloud outlines are noise sources at various points around the loop. The system as it exists without connecting the control loop is highlighted in the dashed box.

output becomes

$$Y_P = P \delta X - SHAP Y_P \quad (\text{A.2})$$

or, recursively,

$$Y_P = \frac{P \delta X}{1 + SHAP}. \quad (\text{A.3})$$

Defining the **open-loop gain**¹, $G \equiv SHAP$, this is

$$Y_P = \left(\frac{1}{1 + G} \right) P \delta X. \quad (\text{A.4})$$

Therefore, the loop has reduced the effect of δX on the plant output by the factor $(1 + G)^{-1}$.

Now, if we consider the effect of the rest of the noise sources, we arrive at

$$Y_P = \left(\frac{1}{1 + G} \right) (P \delta X - AP \delta V_2 - HAP \delta V_1 - SHAP \delta Y). \quad (\text{A.5})$$

Typically, S , A , and P are not easily tunable parameters (e.g., these are often determined by fixed elements of the experimental design, material properties, or by the available performance of transducer elements). On the other hand, the feedback servo filter H generally has a transfer function that is set by the user to determine the overall shape of G . For the servo to work, we must have $G \gg 1$. In this case,

$$Y_P \approx \frac{\delta X}{SHA} - \frac{\delta V_2}{SH} - \frac{\delta V_1}{S} - \delta Y. \quad (\text{A.6})$$

This illustrates one of the major limitations of feedback control. Note that the first two terms above, with H in the denominator, can be reduced by increasing the servo gain. In the limit as $H \rightarrow \infty$, these vanish entirely. The last two terms, however, are unaffected by H . Considering where these noise terms enter the loop, we can understand this effect:

- δY is a physical disturbance that is indistinguishable from a fluctuation in the output of the plant, Y_P . Therefore, the loop interprets it as such, and acts to correct it by pushing the plant accordingly.
- δV_1 is an electrical fluctuation between the sensor S and the servo filter H . To the loop, this is indistinguishable from a fluctuation $\frac{\delta V_1}{S}$ in the output of the plant, Y_P . Just as with δY , the loop acts to cancel this by actuating on the plant.

¹This quantity is called the “open-loop” gain because it is what you would calculate if you opened the loop in one place, then calculated the gain from one side of the opening to the other.

Taken together, noise sources like these are called **sensing noise**. As they cannot be mitigated by increasing the servo gain H , they must be addressed by other means².

A.2 Stability

The factor G is, in general, a complex function of frequency. Judging from the closed-loop gain factor $(1 + G)^{-1}$, we see that the loop blows up if $G = -1$. Therefore, at any frequency for which the open-loop gain magnitude is $|G| = 1$ —known as a **unity-gain frequency (UGF)**—the phase of the loop must not be at or near 180° ³.

An ideal loop has infinite gain and zero phase lag at all frequencies. Clearly, this is not possible in reality. Therefore, it is the job of the designer to gracefully roll off the gain of the loop with increasing frequency so that it satisfies the stability criterion everywhere. A common means of doing so is to give the loop a simple, shallow shape over the expected range of UGFs. For example, if G is a simple single-pole low-pass filter,

$$G(f) = \frac{G_0}{1 + \frac{if}{f_p}}, \quad (\text{A.7})$$

then, at high frequencies,

$$G(f \gg f_p) \approx \frac{-iG_0f_p}{f}. \quad (\text{A.8})$$

Therefore, the UGF can be placed any such frequency above f_p (by setting G_0 appropriately) and the constant phase of $-i$ (i.e., -90°) will prevent an oscillation.

Note that the shape of the loop below the UGF range does not matter; for example, extra pole/zero combinations may be used well below the UGF to increase the loop gain at low frequencies, so long as the loop shape eventually approaches the gentler rolloff as f increases.

The above is a simplified example wherein we have explicitly specified the overall shape of $G(f)$. In general, transfer functions of many elements in the loop—particularly physical system plants—can be quite complicated, and the business of shaping $H(f)$ to make a system stable transitions into an art form.

²In the case of physical disturbances like δY , they must typically be reduced at the source; electronic sensing noise like δV_1 can be improved by employing a better sensor (i.e., by making S bigger).

³This is sometimes known as the “Bode stability criterion”, and it is a special case of the more general **Nyquist stability criterion** in complex analysis.

POUND-DREVER-HALL LOCKING

Pound-Drever-Hall locking [38] is a technique developed for the purpose of laser frequency stabilization. Its basic function is to turn phase modulation into amplitude/intensity modulation in response to the displacement of a laser field from a cavity resonance. This resultant intensity modulation can be sensed using a photodetector, from which a linear error signal can be derived for feeding back to the laser frequency and/or the cavity length to maintain resonance. A typical PDH locking setup is shown in Fig. B.1.

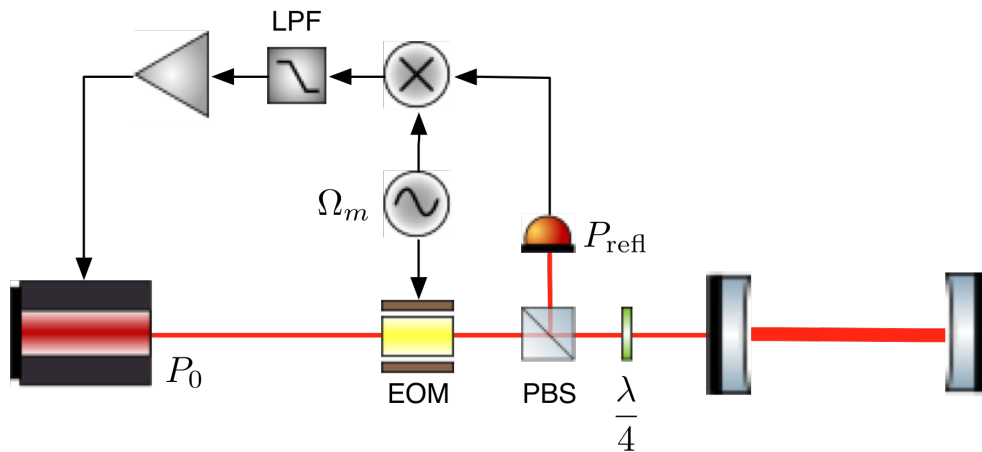


Figure B.1: A Pound-Drever-Hall locking setup. The input laser, with power P_0 , is passed through an EOM, which applies phase modulation sidebands at an RF angular frequency Ω_m before the beam is directed to a cavity. A polarizing beam-splitter (PBS) and quarter-wave plate (“ $\frac{\lambda}{4}$ ”) are used to isolate the reflected beam, with power P_{refl} onto an RF photodetector. The signal from the photodetector is mixed with the original oscillator signal, low-passed to remove the $2\Omega_m$ component, conditioned by a servo filter, and then fed back to the laser frequency to keep the laser locked. (Note: feedback can also be applied to the position of the mirrors to maintain lock. In practice, the more stable of the two quantities—laser frequency or cavity displacement—is used as a reference, while the other is actuated on. This directionality is often a function of Fourier frequency.)

In this scheme, RF phase modulation sidebands are applied to the laser beam using an EOM to create a field with multiple spectral components as in (2.34). Denoting

the modulation angular frequency as Ω_m , the field becomes

$$E = E_0 e^{i\omega_0 t} \left(J_0(\Gamma) + iJ_1(\Gamma)e^{i\Omega_m t} + iJ_1(\Gamma)e^{-i\Omega_m t} \right). \quad (\text{B.1})$$

If we take this field and reflect it from a Fabry-Pérot cavity on resonance, then— noting that the cavity reflectivity at $\Omega_m \gg 2\pi f_p$ is equal to unity, since these sideband fields are far from the cavity resonance—the reflected field is

$$E_{\text{refl}} = E_0 e^{i\omega_0 t} \left(\mathcal{R}J_0(\Gamma) + iJ_1(\Gamma)e^{i\Omega_m t} + iJ_1(\Gamma)e^{-i\Omega_m t} \right), \quad (\text{B.2})$$

where $\mathcal{R} = \mathcal{R}(\omega)$ is the frequency-dependent field reflectivity of the cavity. A photodetector placed in this reflected beam's path will detect a power

$$P_{\text{refl}} = E_{\text{refl}}^* E_{\text{refl}} \quad (\text{B.3})$$

$$\approx P_0 \left[J_0^2(\Gamma) \mathcal{R} \mathcal{R}^* - 2iJ_0(\Gamma)J_1(\Gamma) (\mathcal{R} - \mathcal{R}^*) \cos \Omega_m t \right], \quad (\text{B.4})$$

where we have neglected terms that are higher than first order in $J_1(\Gamma)$, and $P_0 \equiv E_0^2$. The first term above is the DC reflected power. The rest of the expression describes a power modulation at angular frequency Ω_m , proportional to the quantity $-i(\mathcal{R}(\omega) - \mathcal{R}^*(\omega)) = 2 \text{Im}\{\mathcal{R}(\omega)\}$. As we saw in (2.19), the cavity reflectivity on resonance, $\mathcal{R}(0)$, is purely real. Therefore, there is no signal at Ω_m in the reflected power on resonance. However, from (2.25), we expect a linear derivative of $\text{Im}\{\mathcal{R}(0)\}$ with respect to L :

$$\frac{1}{E_{\text{in}}} \frac{dE_{\text{refl}}}{dL} \Big|_{\text{res}} = \frac{d\mathcal{R}(0)}{dL} = 2ikG. \quad (\text{B.5})$$

Differentiating P_{refl} , we get:

$$\frac{dP_{\text{refl}}}{dL} = P_0 \left[J_0^2(\Gamma) \frac{d}{dL} (\mathcal{R} \mathcal{R}^*) - 2iJ_0(\Gamma)J_1(\Gamma) \left(\frac{d\mathcal{R}}{dL} - \frac{d\mathcal{R}^*}{dL} \right) \cos \Omega_m t \right]. \quad (\text{B.6})$$

For the first term,

$$\frac{d}{dL} (\mathcal{R} \mathcal{R}^*) = \mathcal{R} \frac{d\mathcal{R}^*}{dL} + \frac{d\mathcal{R}}{dL} \mathcal{R}^* = (\mathcal{R} - \mathcal{R}^*) \frac{d\mathcal{R}}{dL}, \quad (\text{B.7})$$

which is zero on resonance. Meanwhile, the part of remaining term in parentheses is

$$\frac{d\mathcal{R}}{dL} - \frac{d\mathcal{R}^*}{dL} = \frac{d\mathcal{R}}{dL} - \left(\frac{d\mathcal{R}}{dL} \right)^* = 4ikG \quad (\text{B.8})$$

on resonance. Therefore,

$$\frac{dP_{\text{refl}}}{dL} \Big|_{\text{res}} = 8kGP_0 J_0(\Gamma)J_1(\Gamma) \cos \Omega_m t. \quad (\text{B.9})$$

Once converted into the electrical domain by the photodetector, the resultant signal can be demodulated at Ω_m by mixing it with the same signal used to drive the phase modulator and then rejecting the sum-frequency term by applying a low-pass filter of $f_{LP} \ll \frac{\Omega_m}{2\pi}$. This results in an effective demodulated power signal of

$$\left. \frac{dP_{\text{refl}}^{\text{DM}}}{dL} \right|_{\text{res}} \equiv \mathcal{D}_x(0) \approx 4kGP_0J_0(\Gamma)J_1(\Gamma). \quad (\text{B.10})$$

The equivalent expression with respect to the laser frequency can be found as

$$\left. \frac{dP_{\text{refl}}^{\text{DM}}}{d\nu} \right|_{\text{res}} \equiv \mathcal{D}_\nu(0) = \frac{L}{\nu} \mathcal{D}_x(0) = 8\pi \frac{L}{c} GP_0J_0(\Gamma)J_1(\Gamma). \quad (\text{B.11})$$

Recalling that $\frac{2L}{c} = \nu_{\text{FSR}} = 2f_p\mathcal{F}$, and using (2.39), we can rewrite the above as

$$\mathcal{D}_\nu(0) = \begin{cases} \frac{2P_0J_0(\Gamma)J_1(\Gamma)}{f_p} & \text{critically coupled} \\ \frac{4P_0J_0(\Gamma)J_1(\Gamma)}{f_p} & \text{fully overcoupled} \end{cases}. \quad (\text{B.12})$$

Finally, by comparison with what was found in (2.42), we can determine the frequency dependence of $\mathcal{D}_{x/\nu}$ as

$$\mathcal{D}_{x/\nu}(f) = \frac{\mathcal{D}_{x/\nu}(0)}{1 + \frac{if}{f_p}}. \quad (\text{B.13})$$

Appendix C

SUPPLEMENTAL OPTICAL SPRING RADIATION PRESSURE
NOISE CALCULATIONS

In this appendix, we will show some additional details for the derivation of the formulas presented in the main text. We will first consider the ideal case without optical loss and show the leading-order terms in the large bandwidth and detuning limit. Then, we will show the effect of optical loss and next-order correction terms. Finally, we will consider the implementation of feedback and the closed-loop response of the system, which is relevant to actual experimental realization. Our notation here is nearly identical to that in Ref. [174].

C.0.1 Ideal situation—no optical loss and leading-order terms

In this section, we will consider the ideal situation for a typical optomechanical device, which has been extensively covered in the literature [149, 174–177]. We start with the standard Hamiltonian for the canonical optomechanical device, shown in the dashed box in Fig. 7.1:

$$\begin{aligned} \hat{\mathcal{H}} = & \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega_m^2\hat{x}^2 + \hbar\omega_c\hat{a}^\dagger\hat{a} + \hbar G_0\hat{x}\hat{a}^\dagger\hat{a} \\ & + i\hbar\sqrt{2\gamma}[\hat{a}_{\text{ext}}(t)e^{-i\omega_0 t}\hat{a}^\dagger - \hat{a}_{\text{ext}}^\dagger(t)e^{i\omega_0 t}\hat{a}]. \end{aligned} \quad (\text{C.1})$$

Here, the first two terms are the free Hamiltonian for the oscillator, with ω_m being the mechanical frequency; the third term is the free Hamiltonian for the cavity mode (ω_c is the cavity resonant frequency, and \hat{a} is its annihilation operator satisfying $[\hat{a}, \hat{a}^\dagger] = 1$); the fourth term describes the interaction between the oscillator and the cavity mode, with $G_0 = \omega_c/L$ being the coupling strength and L the cavity length; the remaining part is the coupling between the cavity mode with the external continuum $\hat{a}_{\text{ext}}(t)$, with coupling rate γ and $[\hat{a}_{\text{ext}}(t), \hat{a}_{\text{ext}}^\dagger(t')] = \delta(t-t')$, from which one can define the input operator \hat{a}_{in} (ingoing before interaction) and output operator \hat{a}_{out} (outgoing after interaction) through:

$$\hat{a}_{\text{in}} \equiv \hat{a}_{\text{ext}}(t_-), \quad \hat{a}_{\text{out}} \equiv \hat{a}_{\text{ext}}(t_+), \quad (\text{C.2})$$

according to the standard input-output formalism [178]. In the Hamiltonian, we have ignored those terms accounting for the dissipation mechanism of the mechanical

oscillator coupling to its thermal environment. We will later include their effects in the equation of motion for the oscillator.

Linearized Hamiltonian.—In the experiment, the cavity mode is driven coherently by a laser with a large amplitude at frequency ω_0 . We can therefore study the linearized dynamics by perturbing around the steady state. In the rotating frame of the laser frequency ω_0 , the corresponding linearized Hamiltonian for the system reads:

$$\begin{aligned} \hat{\mathcal{H}} = & \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega_m^2\hat{x}^2 + \hbar\Delta\hat{a}^\dagger\hat{a} + \hbar G_0\hat{x}(\bar{a}^*\hat{a} + \bar{a}\hat{a}^\dagger) \\ & + i\hbar\sqrt{2\gamma}[\hat{a}_{\text{ext}}(t)\hat{a}^\dagger - \hat{a}_{\text{ext}}^\dagger(t)\hat{a}]. \end{aligned} \quad (\text{C.3})$$

Here, the cavity detuning is the difference between the cavity resonant frequency and the laser frequency (i.e., $\Delta \equiv \omega_c - \omega_0$); \bar{a} is the steady-state amplitude of the cavity mode, and if we choose the phase reference such that the steady-state amplitude of the input field, \bar{a}_{in} , is real and positive, we have

$$\bar{a} = \frac{\sqrt{2\gamma}\bar{a}_{\text{in}}}{\gamma + i\Delta} = \frac{\sqrt{2\gamma}}{\gamma + i\Delta} \sqrt{\frac{P_{\text{in}}}{\hbar\omega_0}}, \quad (\text{C.4})$$

where P_{in} is the input laser power. These operators in the above Hamiltonian should be viewed as perturbed parts of the original ones and the quantum state they act on is also transformed correspondingly. For instance, the input state for \hat{a}_{in} is originally a coherent state (for an ideal laser), and now it is the vacuum state $|0\rangle$ with

$$\langle 0|\hat{a}_{\text{in}}(t)\hat{a}_{\text{in}}^\dagger(t')|0\rangle = \delta(t - t'). \quad (\text{C.5})$$

Equations of motion.—Given the above Hamiltonian, the cavity mode satisfies the Heisenberg equation of motion

$$\dot{\hat{a}}(t) + (\gamma + i\Delta)\hat{a}(t) = -iG_0\bar{a}\hat{x}(t) + \sqrt{2\gamma}\hat{a}_{\text{in}}(t), \quad (\text{C.6})$$

and it is related to the cavity output \hat{a}_{out} by the standard input-output relation:

$$\hat{a}_{\text{out}}(t) = -\hat{a}_{\text{in}}(t) + \sqrt{2\gamma}\hat{a}(t). \quad (\text{C.7})$$

Similarly, we can read off the equation of motion for the oscillator:

$$m[\ddot{\hat{x}}(t) + \gamma_m\dot{\hat{x}}(t) + \omega_m^2\hat{x}(t)] = \hat{F}_{\text{rad}}(t) + \hat{F}_{\text{th}}(t). \quad (\text{C.8})$$

Here, we have defined the radiation pressure

$$\hat{F}_{\text{rad}}(t) \equiv -\hbar G_0[\bar{a}^*\hat{a}(t) + \bar{a}\hat{a}^\dagger(t)]. \quad (\text{C.9})$$

In addition, we have added the damping term $m\gamma_m\dot{x}(t)$ and the associated thermal fluctuation force \hat{F}_{th} into the equation of motion, of which the correlation function is $\langle\hat{F}_{\text{th}}(t)\hat{F}_{\text{th}}(t')\rangle = 4m\gamma_mk_B T\delta(t-t')$ in the high-temperature limit $k_B T \gg \hbar\omega_m$.

Solution for the cavity mode.—The above linear equations of motion can be solved in the frequency domain. The solution for the cavity mode reads:

$$\hat{a}(\omega) = \frac{G_0\bar{a}\hat{x}(\omega) + i\sqrt{2\gamma}\hat{a}_{\text{in}}(\omega)}{\omega - \Delta + i\gamma}. \quad (\text{C.10})$$

From this, we can obtain the expression for the radiation pressure:

$$\hat{F}_{\text{rad}}(\omega) = -K_{\text{os}}(\omega)\hat{x}(\omega) + \hat{F}_{\text{noise}}(\omega). \quad (\text{C.11})$$

We introduce the optical spring coefficient K_{os} as:

$$K_{\text{os}}(\omega) \equiv \frac{2\hbar G_0^2|\bar{a}|^2\Delta}{(\omega - \Delta + i\gamma)(\omega + \Delta + i\gamma)}, \quad (\text{C.12})$$

and the quantum radiation pressure noise term as:

$$\hat{F}_{\text{noise}}(\omega) \equiv \frac{2\hbar G_0|\bar{a}|\sqrt{\gamma}}{\sqrt{\gamma^2 + \Delta^2}} \left[\frac{(\gamma^2 + \Delta^2 - i\gamma\omega)\hat{v}_1 + i\Delta\omega\hat{v}_2}{(\omega - \Delta + i\gamma)(\omega + \Delta + i\gamma)} \right] \quad (\text{C.13})$$

with $\hat{v}_1 \equiv (\hat{a}_{\text{in}} + \hat{a}_{\text{in}}^\dagger)/\sqrt{2}$ and $\hat{v}_2 \equiv (\hat{a}_{\text{in}} - \hat{a}_{\text{in}}^\dagger)/\sqrt{2}i$ being the vacuum fluctuation of the input amplitude and phase quadratures, respectively. The strength of the radiation pressure noise can be quantified by its power spectrum, which is defined through

$$\langle 0|\hat{F}_{\text{noise}}^\dagger(\omega)\hat{F}_{\text{noise}}(\omega')|0\rangle_{\text{sym}} \equiv \pi S_F(\omega)\delta(\omega - \omega'), \quad (\text{C.14})$$

where the subscript ‘sym’ denotes for symmetrization and the spectrum is a single-sided one. Notice that for vacuum input state $\langle 0|\hat{v}_k^\dagger(\omega)\hat{v}_l(\omega')|0\rangle_{\text{sym}} = \pi\delta_{kl}\delta(\omega - \omega')$, and therefore

$$S_F(\omega) = \frac{4\hbar^2 G_0^2|\bar{a}|^2\gamma(\gamma^2 + \omega^2 + \Delta^2)}{[(\omega - \Delta)^2 + \gamma^2][(\omega + \Delta)^2 + \gamma^2]}. \quad (\text{C.15})$$

For the case of large bandwidth and detuning in which we are interested, the above radiation pressure noise can be approximated as (up to zeroth order of ω):

$$\hat{F}_{\text{noise}}(\omega) \approx -\frac{2\hbar G_0|\bar{a}|\sqrt{\gamma}}{\sqrt{\gamma^2 + \Delta^2}}\hat{v}_1(\omega) \propto \hat{v}_1(\omega). \quad (\text{C.16})$$

This indicates that the quantum radiation pressure noise is mostly contributed by fluctuations in the amplitude quadrature of the input field. It can be directly measured at the cavity output using a photodetector, as we will see later—this is why we can

evade such noise by feeding back with an appropriate linear filter, which is the central idea of this work.

Solution for the mechanical oscillator.—Given the expression for the radiation pressure, we can write down the solution for the mechanical displacement \hat{x} as:

$$\hat{x}(\omega) = \frac{\hat{F}_{\text{noise}}(\omega) + \hat{F}_{\text{th}}(\omega)}{-m[\omega^2 - \omega_m^2 + i\gamma_m\omega] + K_{\text{os}}(\omega)}. \quad (\text{C.17})$$

As we can see, the mechanical susceptibility is modified into an effective one due to the optical spring effect. Since we are focusing on the case of large cavity bandwidth and detuning, the optical spring response K_{os} can be expanded as:

$$K_{\text{os}} \approx -\frac{2\hbar G_0^2 |\bar{a}|^2 \Delta}{\Delta^2 + \gamma^2} \left[1 + \frac{2i\gamma\omega}{\Delta^2 + \gamma^2} \right] \equiv m\omega_{\text{os}}^2 - i m\Gamma_{\text{os}}\omega, \quad (\text{C.18})$$

where ω_{os} is the optical spring frequency and Γ_{os} is the optical damping coefficient. We can then rewrite the mechanical displacement \hat{x} as

$$\hat{x}(\omega) = \chi_{\text{eff}}(\omega)[\hat{F}_{\text{noise}}(\omega) + \hat{F}_{\text{th}}(\omega)], \quad (\text{C.19})$$

where the effective mechanical susceptibility χ_{eff} is defined through:

$$\chi_{\text{eff}}^{-1}(\omega) \equiv -m[\omega^2 + i(\gamma_m + \Gamma_{\text{os}})\omega - (\omega_m^2 + \omega_{\text{os}}^2)]. \quad (\text{C.20})$$

In the negative-detuning case $\Delta < 0$, ω_{os} is positive and real, and the damping Γ_{os} is negative; in the positive-detuning case $\Delta > 0$, ω_{os} is purely imaginary and the damping Γ_{os} is positive. In both cases, the mechanical system is potentially unstable, especially when the intrinsic damping γ_m is small as in our proposed parameter regime. By introducing an additional laser with a different detuning frequency, we can combine two optical springs and achieve both positive frequency and damping—the so-called double optical spring. Such a scheme has been realized experimentally by Corbitt et al. [143]. We can therefore significantly upshift the mechanical resonant frequency while keeping the oscillator stable.

Solution for the cavity output.—From the input-output relation, the cavity output is given by

$$\hat{a}_{\text{out}}(\omega) = -\frac{\omega - \Delta - i\gamma}{\omega - \Delta + i\gamma} \hat{a}_{\text{in}}(\omega) + \frac{\sqrt{2\gamma} G_0 \bar{a}}{\omega - \Delta + i\gamma} \hat{x}(\omega). \quad (\text{C.21})$$

In the limit of high bandwidth and detuning, we can approximate this as

$$\hat{a}_{\text{out}}(\omega) = -\frac{\Delta + i\gamma}{\Delta - i\gamma} \hat{a}_{\text{in}}(\omega) - \frac{\sqrt{2\gamma} G_0 \bar{a}}{\Delta - i\gamma} \hat{x}(\omega). \quad (\text{C.22})$$

Similarly, for the classical amplitude at DC, we have the input-output relation:

$$\bar{a}_{\text{out}} = -\frac{\Delta + i\gamma}{\Delta - i\gamma} \bar{a}_{\text{in}}. \quad (\text{C.23})$$

The photodetector measures the power of the cavity output field:

$$\begin{aligned} \hat{P}_{\text{out}}(t) &= |(\bar{a}_{\text{out}}^* + \hat{a}_{\text{out}}^\dagger)(\bar{a}_{\text{out}} + \hat{a}_{\text{out}})| \\ &= |\bar{a}_{\text{out}}|^2 + \delta\hat{P}(t) + \hat{a}_{\text{out}}^\dagger \hat{a}_{\text{out}}. \end{aligned} \quad (\text{C.24})$$

It contains the classical DC part $|\bar{a}_{\text{out}}|^2$, and the leading-order time-varying component

$$\delta\hat{P}(t) \equiv \bar{a}_{\text{out}}^* \hat{a}_{\text{out}} + \bar{a}_{\text{out}} \hat{a}_{\text{out}}^\dagger \quad (\text{C.25})$$

that we are interested in, which, in the frequency domain, is given by

$$\delta\hat{P}(\omega) = \bar{a}_{\text{out}}^* \hat{a}_{\text{out}}(\omega) + \bar{a}_{\text{out}} \hat{a}_{\text{out}}^\dagger(\omega) \approx \sqrt{2} \bar{a}_{\text{in}} \hat{v}_1(\omega). \quad (\text{C.26})$$

This means that the photodetector mostly measures fluctuations in the amplitude quadrature of the input field, which is the main contributor to the quantum radiation pressure noise felt by the mechanical oscillator as shown by Eq. (C.16). Therefore, simply by feeding back the photodetector signal to the mechanical oscillator, we will be able to evade the quantum radiation pressure noise. The only limitation arises from the optical loss and the frequency dependence of the radiation pressure noise that we have ignored in Eq. (C.16) by assuming a large cavity bandwidth and detuning.

C.1 Realistic situation—optical loss and next-order corrections

In this section, we will analyze the effect of optical loss and also the next-order correction—frequency-dependence of the radiation pressure noise as well as non-zero response to the mechanical displacement in the photocurrent—due to finite cavity bandwidth. As mentioned in the main text, the optical loss will decrease the noise cancelation efficiency by introducing vacuum fluctuations—which we denote \hat{a}'_{in} —that are uncorrelated with \hat{a}_{in} . In terms of the equation of motion for the cavity mode, we have

$$\dot{\hat{a}} + (\gamma_{\text{tot}} + i\Delta)\hat{a} = -iG_0 \bar{a} \hat{x} + \sqrt{2\gamma} \hat{a}_{\text{in}} + \sqrt{2\gamma_\epsilon} \hat{a}'_{\text{in}}, \quad (\text{C.27})$$

where \bar{a} is modified into

$$\bar{a} = \frac{\sqrt{2\gamma} \bar{a}_{\text{in}}}{\gamma_{\text{tot}} + i\Delta} \quad (\text{C.28})$$

and we have introduced

$$\gamma_{\text{tot}} \equiv \gamma + \gamma_\epsilon = \gamma + c\epsilon/(4L) \quad (\text{C.29})$$

where ϵ is the roundtrip power loss factor in the cavity.

Modification of the radiation pressure.—Correspondingly, this will modify the radiation pressure [cf. Eq. (C.11)]:

$$\hat{F}_{\text{rad}}(\omega) = -K_{\text{os}}(\omega)\hat{x}(\omega) + \hat{F}_{\text{noise}}(\omega), \quad (\text{C.30})$$

where

$$K_{\text{os}} = \frac{2\hbar G_0^2 |\bar{a}|^2 \Delta}{(\omega - \Delta + i\gamma_{\text{tot}})(\omega + \Delta + i\gamma_{\text{tot}})}, \quad (\text{C.31})$$

and

$$\begin{aligned} \hat{F}_{\text{noise}} \equiv & \frac{2\hbar G_0 \sqrt{\gamma} |\bar{a}|}{\sqrt{\gamma_{\text{tot}}^2 + \Delta^2}} \left\{ \frac{(\gamma_{\text{tot}}^2 + \Delta^2 - i\gamma_{\text{tot}}\omega)\hat{v}_1 + i\Delta\omega\hat{v}_2}{(\omega - \Delta + i\gamma_{\text{tot}})(\omega + \Delta + i\gamma_{\text{tot}})} \right. \\ & \left. + \sqrt{\frac{\gamma_\epsilon}{\gamma}} \frac{(\gamma_{\text{tot}}^2 + \Delta^2 - i\gamma_{\text{tot}}\omega)\hat{v}'_1 + i\Delta\omega\hat{v}'_2}{(\omega - \Delta + i\gamma_{\text{tot}})(\omega + \Delta + i\gamma_{\text{tot}})} \right\}, \end{aligned} \quad (\text{C.32})$$

where $\hat{v}'_1 \equiv (\hat{a}'_{\text{in}} + \hat{a}'_{\text{in}}^\dagger)/\sqrt{2}$ and $\hat{v}'_2 \equiv (\hat{a}'_{\text{in}} - \hat{a}'_{\text{in}}^\dagger)/(\sqrt{2}i)$.

Again for large bandwidth and detuning, and keeping up to the next-order correction—leading order of ϵ and ω —we obtain:

$$K_{\text{os}} = -\frac{2\hbar G_0^2 |\bar{a}|^2 \Delta}{\Delta^2 + \gamma^2} \left[1 - \frac{4\gamma\gamma_\epsilon}{\gamma^2 + \Delta^2} + \frac{2i\gamma\omega}{\gamma^2 + \Delta^2} \right], \quad (\text{C.33})$$

and

$$\hat{F}_{\text{noise}} = -\frac{2\hbar G_0 \sqrt{\gamma} |\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}} \left[\hat{v}_1 + \frac{i\omega\Delta}{\gamma^2 + \Delta^2} \hat{v}_2 + \sqrt{\frac{\gamma_\epsilon}{\gamma}} \hat{v}'_1 \right]. \quad (\text{C.34})$$

Modification of the input-output relation.—Similarly, the input-output relation is also modified into

$$\begin{aligned} \hat{a}_{\text{out}} = & -\frac{\omega - \Delta - i(\gamma - \gamma_\epsilon)}{\omega - \Delta + i\gamma_{\text{tot}}} \hat{a}_{\text{in}} + \frac{2i\sqrt{\gamma\gamma_\epsilon}}{\omega - \Delta + i\gamma_{\text{tot}}} \hat{a}'_{\text{in}} \\ & + \frac{\sqrt{2\gamma} G_0 \bar{a}}{\omega - \Delta + i\gamma_{\text{tot}}} \hat{x}. \end{aligned} \quad (\text{C.35})$$

Modification of the photocurrent output.—The exact expression for the AC part of the photocurrent output $\delta\hat{P}(\omega)$ is quite complicated, however, in our stated limit, we have

$$\delta I(\omega) \equiv \sqrt{2}\bar{a}_{\text{in}}\hat{v}_1(\omega) + \delta\hat{P}_\epsilon(\omega) + \delta\hat{P}_\eta(\omega) + \delta\hat{P}_x(\omega), \quad (\text{C.36})$$

where the term $\delta\hat{P}_\epsilon$ contains the vacuum fluctuations $\hat{v}_{1,2}$ that are associated with optical loss:

$$\delta\hat{P}_\epsilon(\omega) = \frac{2\sqrt{2\gamma\gamma_\epsilon}\bar{a}_{\text{in}}}{\gamma^2 + \Delta^2}(\gamma\hat{v}'_1 - \Delta\hat{v}'_2), \quad (\text{C.37})$$

the additional noise term $\delta\hat{P}_\eta$, due to non-unity quantum inefficiency η of the photodetector (keeping to the first order of small $1 - \eta$), is

$$\delta\hat{P}_\eta(\omega) \approx \sqrt{2}\bar{a}_{\text{in}}\sqrt{1 - \eta}\hat{n}, \quad (\text{C.38})$$

and the term $\delta\hat{P}_x$ depends on the mechanical displacement:

$$\delta\hat{P}_x(\omega) = -\frac{2G_0|\bar{a}|^2\Delta(2\gamma_\epsilon - i\omega)}{\gamma^2 + \Delta^2}\hat{x}(\omega). \quad (\text{C.39})$$

Therefore, not only is there excess noise from the vacuum fluctuations introduced by the optical loss and non-unity quantum efficiency, but there is also a parasitic sensitivity to mechanical displacement, which is actually associated with the excess radiation pressure [cf. Eq. (C.34)], compared with the ideal case [cf. Eq. (C.16)].

C.2 Feedback and closed-loop response

The radiation pressure noise can be removed either by feedforward (i.e., the photocurrent output is fed forward to the mechanical oscillator as a force), or by feedback (i.e., the photocurrent output is fed back to the input field via an amplitude modulator). Here, we consider the implementation of the feedback scheme. Not only is it more robust against uncertainty in the model transfer functions, but also, as we will show, it can remove the negative damping in the optical spring and stabilize the mechanical oscillator, allowing in principle for a stable *single* optical spring.

According to the diagram shown in Fig. C.1, the photocurrent output is fed back to an amplitude modulator, which modulates the amplitude quadrature of the input field. The set of equations for relevant quantities describe such a feedback scheme go as follows, keeping up to the leading order of ϵ and ω :

$$\hat{x} = \chi_{\text{eff}}(\hat{F}_{\text{noise}} + \hat{F}_{\text{th}}), \quad (\text{C.40})$$

$$\hat{F}_{\text{noise}} = -\frac{2\hbar G_0\sqrt{\gamma}|\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}}\left[\hat{v}_1^{\text{loop}} + \frac{i\omega\Delta}{\gamma^2 + \Delta^2}\hat{v}_2 + \sqrt{\frac{\gamma_\epsilon}{\gamma}}\hat{v}'_1\right], \quad (\text{C.41})$$

$$\delta\hat{P} = \sqrt{2}\bar{a}_{\text{in}}\hat{v}_1^{\text{loop}} + \delta\hat{P}_\epsilon + \delta\hat{P}_\eta + \delta\hat{P}_x, \quad (\text{C.42})$$

$$\hat{v}_1^{\text{loop}} = \hat{v}_1 - K_c(\delta\hat{P}/\sqrt{2}\bar{a}_{\text{in}}). \quad (\text{C.43})$$

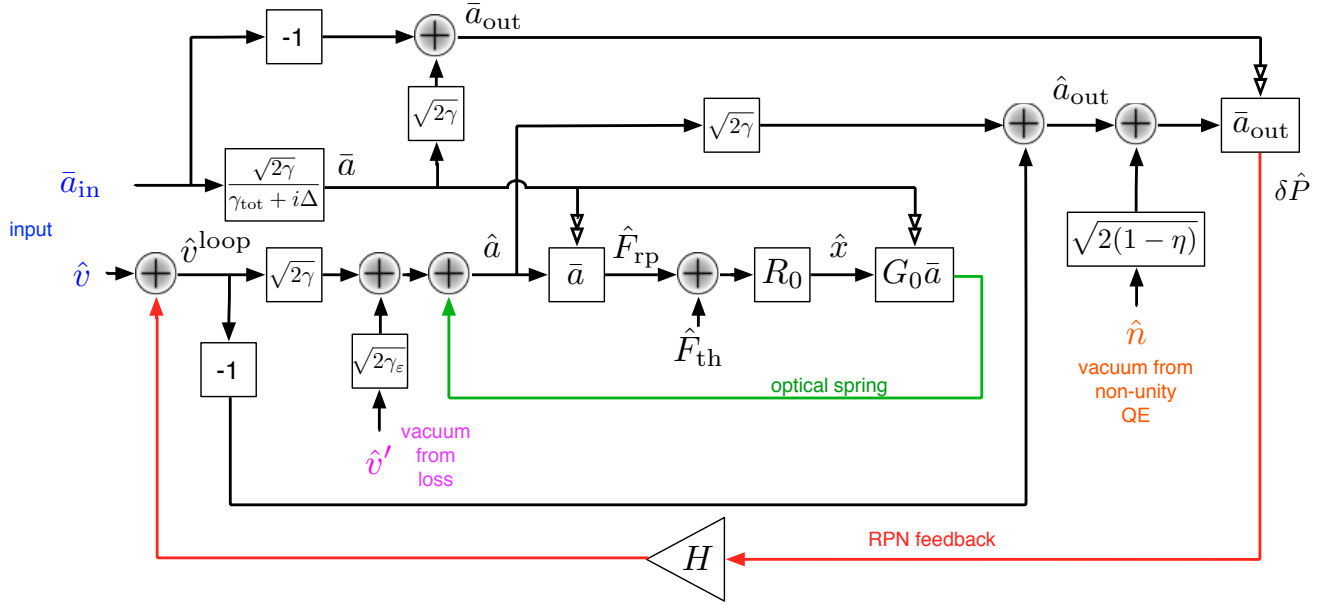


Figure C.1: A graphical representation of the feedback model described in the text. The effects of the optical spring and the active external feedback are shown explicitly, while the (static) resonance of the classical field \bar{a} is not. Using the input-output relations, the input fields at left are split into the prompt reflections and the portions that enter the cavity. Then, the leakage fields are summed with the prompt reflections to give the output. The feedback kernel is $H \equiv -K_c(\delta\hat{P}/\sqrt{2}\bar{a}_{\text{in}})$. Note that the double-headed arrows correspond to the setting of a parameter block, while single-headed arrows denote signal transmission as usual.

Here \hat{v}_1^{loop} is the in-loop amplitude quadrature after the amplitude modulator; K_c is the feedback kernel function and we intentionally leave out the factor $\sqrt{2}\bar{a}_{\text{in}}$ to simplify the equations.

We are interested in the motion of the mechanical oscillator when the feedback is turned on. Solving the above equations leads to

$$\hat{x} = \chi_{\text{eff}'}(\hat{F}_{\text{noise}'} + \hat{F}_{\text{th}}), \quad (\text{C.44})$$

where

$$\chi_{\text{eff}'}^{-1} = \chi_{\text{eff}}^{-1} - \frac{4\hbar G_0^2 |\bar{a}|^2 \gamma \Delta (2\gamma_\epsilon - i\omega)}{(\gamma^2 + \Delta^2)^2} \frac{K_c}{1 + K_c}, \quad (\text{C.45})$$

and

$$\begin{aligned}\hat{F}'_{\text{noise}} = & -\frac{2\hbar G_0\sqrt{\gamma}|\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}} \left[\frac{1}{1 + K_c} \hat{v}_1 \right. \\ & - \sqrt{\frac{\gamma_\epsilon}{\gamma}} \left(\frac{K_c - 1}{K_c + 1} \gamma^2 - \Delta^2 \right) \hat{v}'_1 + \frac{2\sqrt{\gamma_\epsilon\gamma}\Delta K_c}{K_c + 1} \hat{v}'_2 \\ & \left. - \frac{K_c}{1 + K_c} \sqrt{1 - \eta} \hat{n} + \frac{i\omega\Delta}{\gamma^2 + \Delta^2} \hat{v}_2 \right].\end{aligned}\quad (\text{C.46})$$

Ideal-feedback limit.—If we make $K_c \rightarrow \infty$, namely, in the ideal feedback limit, we have

$$\begin{aligned}\chi_{\text{eff}}^{-1}|_{K_c \rightarrow \infty} = & \chi_{\text{eff}}^{-1} - \frac{4\hbar G_0^2 |\bar{a}|^2 \gamma \Delta (2\gamma_\epsilon - i\omega)}{(\gamma^2 + \Delta^2)^2} \\ = & \chi_0^{-1} + K_{\text{os}} - \frac{4\hbar G_0^2 |\bar{a}|^2 \gamma \Delta (2\gamma_\epsilon - i\omega)}{(\gamma^2 + \Delta^2)^2} \\ = & -m \left[\omega^2 + i\gamma_m \omega - (\omega_m^2 + \omega_{\text{os}}^2) \right],\end{aligned}\quad (\text{C.47})$$

where we have plugged in the expression for χ_{eff} [cf. Eq.(C.20)] and K_{os} [cf. Eq.(C.33)]. Interestingly, the original negative damping Γ_{os} in K_{os} associated with the positive rigidity is canceled out, and the mechanical oscillator is stabilized. Therefore, using this feedback scheme, the resultant oscillator is stable with a shifted resonant frequency

$$\omega_m^{\text{new}} = \sqrt{\omega_m^2 + \omega_{\text{os}}^2}.\quad (\text{C.48})$$

Now, we quantify the residual radiation pressure noise on the mechanical oscillator. We have:

$$\begin{aligned}\hat{F}'_{\text{noise}}|_{K_c \rightarrow \infty} = & \frac{2\hbar G_0\sqrt{\gamma}|\bar{a}|}{\sqrt{\gamma^2 + \Delta^2}} \left\{ \sqrt{\frac{\gamma_\epsilon}{\gamma}} \frac{(\gamma^2 - \Delta^2) \hat{v}'_1 + 2\gamma\Delta \hat{v}'_2}{\gamma^2 + \Delta^2} \right. \\ & \left. - \sqrt{1 - \eta} \hat{n} - \frac{i\omega\Delta \hat{v}_1}{\gamma^2 + \Delta^2} \right\}.\end{aligned}\quad (\text{C.49})$$

The corresponding spectral density reads

$$S_F^{\text{res}} = \frac{4\hbar^2 G_0^2 \gamma |\bar{a}|^2}{\gamma^2 + \Delta^2} \left[\frac{\gamma_\epsilon}{\gamma} + 1 - \eta + \frac{\omega^2 \Delta^2}{(\gamma^2 + \Delta^2)^2} \right].\quad (\text{C.50})$$

The first term accounts for the effect of the optical loss; the second accounts for non-unity quantum efficiency of the photodetector; the third term accounts for a finite cavity bandwidth.

C.3 Proposed experimental setup

While the technique described in this paper is quite general, a possible experimental layout is shown in Fig. C.2. A laser's frequency is stabilized to the optical spring cavity length using the Pound-Drever-Hall (PDH)[101] locking technique. The laser is then detuned from the resonance by injecting an offset into the error point of the control loop. A second beam is picked off from the main laser and upshifted in frequency by an acousto-optic modulator (AOM). Once the detuning is set properly for both beams, the input power is ramped until the optical spring reaches the desired strength. At this point, the PDH frequency feedback to the laser can be disengaged, and—provided the mechanical and laser frequency stability are sufficiently high—the resonator is trapped in a passively stable potential by the optical spring forces.

Finally, the radiation pressure noise feedback described above is engaged, with the primary beam's signal fed back to the laser amplitude, and the secondary's to the AOM drive amplitude. The EOM, which imparts the phase modulation sidebands necessary for PDH lock, is also disengaged so as not to couple extra uncorrelated vacuum noise into the readout. In this operational configuration, the quantum radiation pressure noise is very strongly suppressed, limited only by the parasitic loss and finite-bandwidth effects detailed above. Classical laser amplitude noise—which is indistinguishable from its quantum counterpart here—is also suppressed by the loop.

The only remaining potential issues are laser frequency stability and drift of the mechanical system, which can drive the optical spring fields away from their optimal detunings. The former can be avoided using pre-stabilization (e.g., by locking the laser to an external frequency reference). The latter is not as simple to avoid, and will depend on the mechanical system in question. If necessary, a very weak PDH lock can be maintained using a low-frequency servo to ensure DC stability of the operating point. In this case, it may be possible to use weak enough control sideband fields that the RPN readout is still limited by the finite losses and bandwidth.

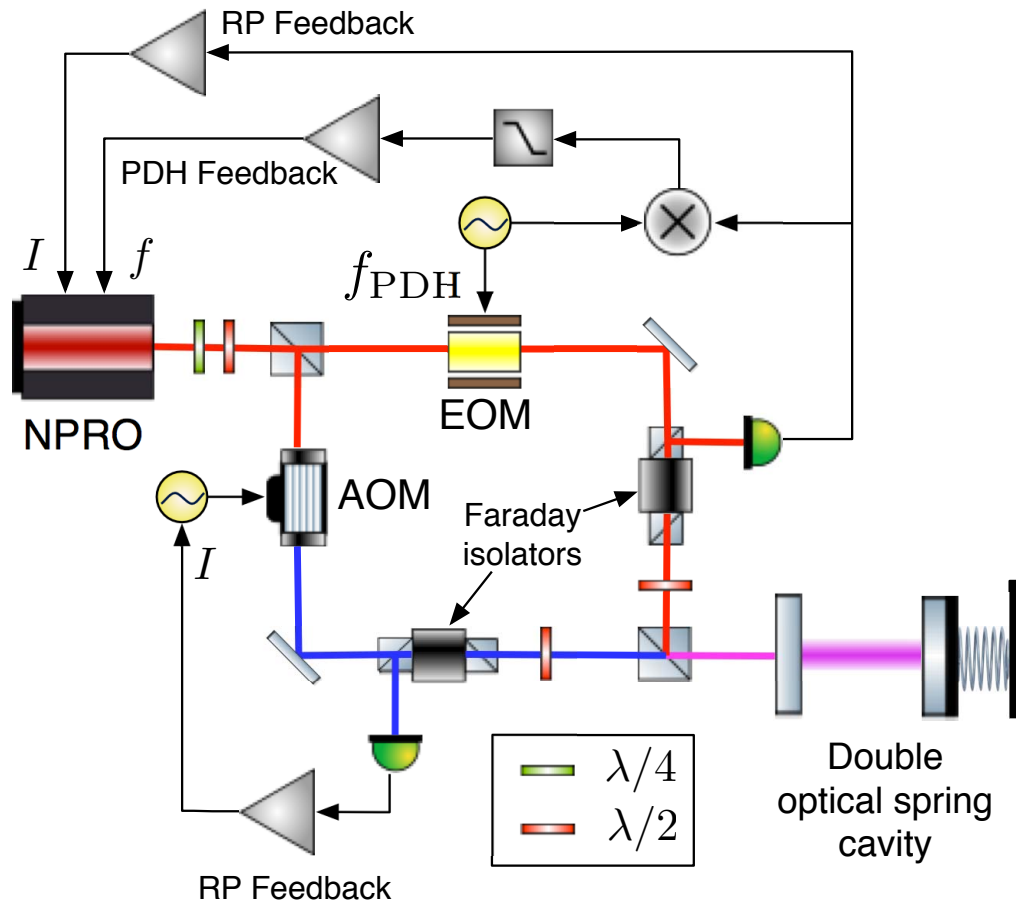


Figure C.2: The proposed experimental setup. A laser beam is split and one path is upshifted in frequency, allowing for independent control of the detuning of each field. Each beam's intensity can also be controlled by feeding back to the laser or the modulator, and these channels are used for the radiation pressure noise feedback. A Pound-Drever-Hall locking scheme is used to set the operating point before strengthening the stable optical spring.