

Advancement of cryogenic suspensions and optical cavities for gravitational wave detection

Thejas Seetharamu

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF
DOCTOR OF PHILOSOPHY

SCHOOL OF PHYSICS AND ASTRONOMY
COLLEGE OF SCIENCE AND ENGINEERING



University
of Glasgow

FEBRUARY 2026

To my grandmother, Pilli

Abstract

This thesis presents PhD work at the Institute for Gravitational Research and Caltech LIGO Lab that was done between October 2021 and February 2026 as part of the PhD programme. The overarching aim is to develop and validate key experimental technologies that enable precision displacement measurements at the levels required for gravitational-wave detectors and their future upgrades, with a particular focus on cryogenic instrumentation, low-noise sensing, and optical readout. The work spans two closely connected contexts: (i) the Glasgow Cryogenic Interferometer Facility (GCIF), a cryogenic GW detector prototype facility intended to directly measure coating thermal noise by monitoring the thermal noise in a short Fabry–Pérot cavity; and (ii) Advanced LIGO upgrade activities at the Caltech LIGO Laboratory, centred on the Output Mode Cleaner (OMC) and readout architectures relevant to squeezed-light operation and balanced homodyne detection.

A major component of this thesis addresses the suspension design required to realise a cryogenic coating-thermal-noise measurement in GCIF. Achieving the target sensitivity demands simultaneous control of multiple competing requirements: strong isolation from ground and platform motion at frequencies of interest, practical integration into a compact cryostat, and compatibility with cryogenic heat-load constraints. To meet these goals, a *common platform* suspension architecture is developed and analysed, in which the two suspension chains that form the cavity share a common top stage to enhance practical implementability and enable partial common-mode rejection. A combination of analytical modelling, numerical eigenmode studies, and state-space simulations is used to quantify longitudinal isolation, cross-coupling between degrees of freedom, and mode visibility across the chain. These studies inform design choices such as mass distribution,

wire parameters, and damping strategy. The outcome is a mechanically feasible suspension concept with a clear pathway to prototyping and commissioning, together with a modelling workflow that connects design parameters to measurable performance metrics and supports iterative refinement.

To control and read out the motion of cryogenic suspended stages, the thesis discusses development of a low-noise *cryogenic shadow sensor*, designed to provide robust local displacement sensing with minimal heat load. The sensor concept employs an infrared LED source, a flag attached to the moving stage, and an InGaAs quadrant photodiode (QPD) readout. The work presents a systematic experimental characterisation of LED and photodiode performance at cryogenic temperature, including optical efficiency, relative intensity noise, and electronic noise contributions. Particular emphasis is placed on identifying practical operating points that maximise displacement sensitivity and dynamic range while maintaining cryogenic compatibility. The sensor geometry is explored to understand how beam size, flag dimensions, and alignment tolerances shape linearity and noise performance, and how these trade-offs propagate into suspension damping capability. The resulting sensor design and measurement results establish a quantitative basis for sensor selection and integration in GCIF, and more broadly demonstrate an approach to engineering cryogenic-compatible displacement sensors where thermal budget and low-frequency stability are both critical. We achieved a shot-noise-limited displacement sensitivity of $\sim 5 \times 10^{-10} \text{ m}/\sqrt{\text{Hz}}$ for $f \gtrsim 100 \text{ Hz}$. We also discuss how this can be improved further.

The thesis also reports work at the Caltech LIGO Laboratory (January–October 2023, April–July 2024, and August 2025) related to optical filtering and readout in kilometre-scale detectors. A central component is the Output Mode Cleaner (OMC), which rejects higher-order spatial modes and radio-frequency sidebands at the interferometer output, thereby reducing technical noise couplings and improving the effectiveness of squeezed-light operation. This work contributes to the development and characterisation of enhanced OMCs that meet the high-throughput requirements of greater than 98% efficiency for the next upgrade of LIGO. These upgrades are closely linked to the adoption of bal-

anced homodyne detection (BHD) in the O5 era: an enhanced readout scheme that enables full use of squeezed light and reduces quantum noise by reading out the interferometer signal via interference with a local oscillator. In this context, the thesis highlights key practical complexities associated with deploying BHD on a kilometre-scale detector, and assesses the feasibility of a polarisation-based implementation as an alternative strategy. The final chapter develops the mathematical framework and modelling tools required to analyse polarisation BHD, and presents an experimental demonstration that validates the approach and informs requirements for robust, high-efficiency readout.

Contents

Abstract	iii
Acknowledgements	xxxi
Declaration	xxxv
1 Fundamentals of Gravitational Wave Detection	1
1.1 Gravitational-Wave Astronomy	1
1.2 Sources of gravitational waves	3
1.2.1 Compact Object Binaries	4
1.2.2 Continuous Sources	6
1.2.3 Burst Sources	7
1.2.4 Stochastic Sources	7
1.2.5 Notable Observations from LIGO	8
1.3 Gravitational-wave detectors	11
1.4 Noise Sources in GW Detectors	13
1.4.1 Thermal noise	14
1.4.2 Quantum noise	15
1.4.3 Seismic noise	16
1.5 Suspension System in LIGO	18
1.6 Shadow Sensors	20
1.7 Future upgrades to Advanced LIGO	21
1.7.1 The A+ upgrade	21
1.7.2 LIGO India	22
1.7.3 Post-O5 concepts: A#	23
1.8 Third-generation gravitational-wave detectors	24
1.9 Detector prototypes and technology demonstrators	27

1.9.1	Glasgow Cryogenic Interferometer Facility (GCIF)	28
1.10	Overview of thesis	30
2	Glasgow Cryogenic Interferometer Facility: Common Platform Suspension Design	33
2.1	Introduction	33
2.2	Pendulum Dynamics: State Space	34
2.3	Dynamics of a Single Stage Pendulum	37
2.3.1	Modes of a Single Stage Pendulum	40
2.4	Triple Stage Pendulum	49
2.4.1	Some Definitions	51
2.5	Split Triple Stage Pendulum: Common Platform Suspensions	54
2.5.1	Common-Platform vs. Conventional Triple Suspensions	54
2.5.2	Design Requirements and Considerations	56
2.5.2.1	Length Noise and Mode Placement	58
2.5.2.2	Local Control and Sensing Requirements	58
2.5.2.3	Mechanical and Cryogenic Feasibility	59
2.5.3	Resonance Mode Forest	60
2.6	Model A - Cage like Reaction Mass	61
2.7	Moving the COM to the centre of top mass	68
2.7.1	Takeaway	69
2.8	Making the final stage longer- More Crystal, less Metal	70
2.8.1	Takeaway	72
2.9	Massive Top Mass	73
2.9.1	Takeaway	76
2.10	Asymmetric chains	76
2.10.1	Takeaway	80
2.11	Conclusion and Future Development	80
3	Development of a Cryogenic Shadow Sensor	84
3.1	Introduction	84
3.1.1	Photodiodes	86
3.1.2	Light Emitting Diodes	88

3.2	GOSEM Electronics	89
3.2.1	Light Emitting Diode	89
3.2.2	Quadrant Photodiode	92
3.2.3	Noise Modelling and Measurements	95
3.2.4	Discussion	96
3.3	Python Modelling of Different Flag Geometries	96
3.3.1	Rectangular Flag	104
3.3.2	Circular Flag	110
3.3.3	Design implications	116
3.4	The Prototype	120
3.4.1	In-Air Characterisation using Circular Aperture Flag	121
3.4.2	Cryogenic Characterisation	126
3.4.2.1	LED Characterisation: I-V Curve	127
3.4.2.2	LED Characterisation: Optoelectronic efficiency	128
3.4.2.3	LED Characterisation: Relative Intensity Noise	132
3.4.2.4	LED Characterisation: Reliability Testing	132
3.4.2.5	QPD Characterisation: Dark noise	135
3.4.2.6	Discussion	139
3.4.2.7	Sensor Characterisation	140
3.5	Package Design	143
3.5.1	Cryogenic design drivers	146
3.5.2	Material selection and insulation strategy	147
3.5.3	Implementation in the cryogenic package	148
3.6	Notes and Future Development	149
4	A+ LIGO: Output Mode Cleaner Cavities	152
4.1	Introduction	153
4.1.1	Transition from DC Readout to BHD	154
4.1.2	Requirements and Impact	158
4.2	Cavity Design	159
4.3	Testing of Cavity Optical Components	160
4.3.1	Characterisation of Prism Mirrors	160

4.3.1.1	Wedge angle	161
4.3.1.2	Perpendicularity	162
4.3.2	Characterisation of the Curved Mirror Subassemblies	163
4.3.2.1	Tombstone Mounting Mirrors	164
4.3.2.2	Curved Mirrors	164
4.4	Sub-assembly matching and bonding	169
4.5	Cavity Assembly	176
4.6	Cavity Alignment Procedure	180
4.7	Cavity Characterisation	182
4.7.1	Cavity Transverse Mode Spacing and Free Spectral Range	182
4.8	Cavity Throughput Optimisation	190
4.9	PZT Characterisation	191
4.10	Cavity Bonding	195
4.11	Conclusion and Future Directions	196
5	Polarisation Balanced Homodyne Detection	198
5.1	Introduction	198
5.2	Birefringence in Optical Cavities	200
5.3	Revisiting BHD equations	202
5.4	Two-Polarisation BHD	203
5.4.1	BHD with HWP before the OMC	204
5.4.2	BHD with HWP after the OMC	206
5.4.3	Simultaneous Half-Wave Plate Mixing - Cancelling Birefringence Noise	208
5.4.4	Investigating Discrepancy	214
5.5	OMC6 Birefringence Measurement	219
5.6	Measuring the Optimal Mixing Angles for OMC6	222
5.7	Discussion	227
5.7.1	Future directions	228
6	Conclusion	229
6.1	Summary of contributions	230
6.1.1	Common-platform suspension modelling for GCIF	230

6.1.2	Cryogenic shadow sensor for local damping	231
6.1.3	A+ LIGO OMC cavities: design, assembly, and performance	232
6.1.4	Polarisation-based balanced homodyne detection with a single OMC	233
6.2	Outlook and recommended future work	234
6.2.1	GCIF suspension and controls integration	234
6.2.2	Cryogenic shadow sensor further work and miniaturisation	235
6.2.3	A+ OMC deployment and extension	235
6.2.4	Polarisation BHD refinement and scaling	235
6.3	Closing remarks	236

Appendices **237**

A	PCB Images	237
A.1	LED Driver	237
A.2	QPD Amplifier	238
B	Wave Plate Calibration Curves	239
B.1	Calibration of Waveplates In-Situ	241
C	Suspension Model Parameters for State Space Export	243
D	Suspension Modes for Common Platform Suspensions	245

List of Tables

2.1	Key mechanical parameters used in the suspension model of Sec. 2.6. The listed <i>d</i> -values are <i>effective</i> attachment offsets: the corresponding physical attachment distances are obtained by subtracting the relevant flexure length from these effective values.	62
2.2	Suspension mode list for the two-chain (I/E) common–platform suspension model with reaction cage. Refer to the diagram in Figure 2.8 for guidance on the DOF notations. The mode images can be found in Appendix D	64
2.3	Mode eigenvector components for the translational (<i>x,y,z</i>) and rotational (yaw, pitch, roll) DOFs for each stage for Mode 3 at 0.723078 Hz.	66
4.1	Measured wedge angles of prism mirrors. Mirrors with wedge angles exceeding 0.5° fall outside specification and are excluded from cavity construction. . . .	161
4.2	Prism mirror deviation from perpendicularity, as measured using a white light autocollimator. Perpendicularity better than 30 arcsec is required to minimise beam displacement due to the lack of internal pitch adjustment in the OMC cavity assembly. Deviations beyond this limit are marked non-conforming. . .	162
4.3	Tombstone mirror deviation from perpendicularity, as measured using a white-light autocollimator. Perpendicularity better than 30 arcsec is required to minimise beam displacement due to the lack of internal pitch adjustment in the OMC cavity assembly. Deviations beyond this limit are marked non-conforming.	164
4.4	Radius Of Curvature (ROC) of the 20 A+ OMC curved optics.	165
4.5	Curvature minimum and wedge angle data for the curved optics.	168
4.6	Wedge parameters and TIS for various CMs.	172
4.7	A+ OMC cavity performance parameters summary	183
4.8	Summary of measured and derived cavity parameters from power budget analysis of the bonded A+ OMC cavity 101.	191
4.9	Cavity Summary (after bonding)	196

4.10	Measured OMC parameters (OMC1–OMC3).	197
4.11	Measured OMC parameters (OMC4–OMC6).	197
6	Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: masses and inertias. Here, mass 1 denotes the upper/common platform, mass 2 the intermediate masses, and mass 3 the optics. Subscripts I and E refer to the input and end chains respectively. . . .	243
7	Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: attachment-point heights and wire separa- tions. Subscripts I and E refer to the input and end chains, respectively. . . .	244
8	Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: stage lengths and elastic properties. Sub- scripts I and E refer to the input and end chains, respectively.	245

List of Figures

1.1	Ring of freely falling particles in a GW that stretches and squeezes space-time. This figure illustrates two different (orthogonal) polarisation of GWs. These time series images depict how space-time curvature changes for $\omega t = 0, \pi/2, \pi, 3\pi/2$. in the plane of the paper for an incident GW that propagates out of the page.	2
1.2	Cumulative number of gravitational-wave detections as a function of month (UTC). The curve is constructed by converting event GPS times to UTC, binning detections by month, and taking the cumulative sum. Shaded vertical ribbons indicate the LIGO–Virgo–KAGRA observing runs O1–O4 (csv file).	5
1.3	GW151014 event as observed by LIGO Hanford (left) and LIGO Livingston (right). From top to bottom, are the strain signal, waveforms, residual noise, and a time-frequency representation of the strain data [10].	8
1.4	GW170817 chirp in time–frequency maps	10
1.5	A+ Design Curve with noise curves for all the fundamental noise sources in LIGO and Seismic noise. [44].	13
1.6	Drawing of Advanced LIGO quadruple pendulum with a quadruple reaction chain hanging behind it with its surrounding structure [39].	19
1.7	When the flag moves by a distance x , the amount of light falling on the photodiode changes. This is registered as V_{out}	20
1.8	Artist’s impression of the two thrid-generation GW detectors.	25
1.9	Strain sensitivity of current and future gravitational-wave detectors	26
1.10	GCIF Cryostat in the JIF lab (Kelvin Building, University of Glasgow). A section of the 10 m reference cavity beam tube is seen behind the cryostat in the photograph.	29
2.1	Single-stage pendulum degrees of freedom	38

2.2	Schematic of a rigid body suspended from a stationary platform by four wires. The parameters n_1 and n_2 denote the lateral distances of the wire attachment points from the z -axis along the centre-of-mass (COM) line of the body. The quantity d is the vertical offset of the attachment plane from the COM, l is the wire length, and s is the axial separation of the front and back attachment points.	39
2.3	Transfer functions for a representative single-stage pendulum. The orange trace shows the magnitude of the transfer function from longitudinal motion of the support to longitudinal motion of the suspended mass, $ H_x $. Below the pendulum resonance the mass follows the support motion with approximately unity gain, while near the resonance the response is amplified before rolling off with an asymptotic $1/f^2$ dependence at higher frequencies, providing seismic isolation in the observation band. The green trace shows the corresponding cross-coupled transfer function from support longitudinal motion to optic pitch, $ H_{x \rightarrow \text{pitch}} $. The blue trace displays the ratio $ H_{x \rightarrow \text{pitch}} / H_x $, which quantifies the longitudinal-to-pitch coupling: it peaks near the coupled resonance and tends to a constant at high frequency, reflecting the fixed geometrical ratio between translational and rotational response in the inertial regime. . . .	48
2.4	Longitudinal transfer functions from support motion to test-mass motion, $ H_x $, for single-, double-, and triple-stage pendulums. At low frequency all three systems follow the support motion with approximately unity gain, while around the pendulum resonances the response is amplified. Above the highest resonance the attenuation improves with each additional stage, approaching asymptotic slopes of $1/f^2$, $1/f^4$, and $1/f^6$ for the single-, double-, and triple-stage configurations respectively, illustrating the enhanced seismic isolation provided by multiple cascaded stages.	50
2.5	Split pendulum configuration with common top mass.	51
2.6	Triple-stage common platform suspension model	57
2.7	Mode visibility for the common-platform triple suspension. The model corresponds to a configuration in which the recoil masses are realised as a cage structure integrated with the top mass.	63

2.8	Split Pendulum diagram indicating the notation used for the degrees of freedom for each stage used in Table 2.2	64
2.9	Transfer function from a transverse force disturbance applied at the intermediate mass to the resulting yaw displacement of the optic, $H_{F_y \rightarrow yaw}(f)$, shown as magnitude and phase versus frequency.	66
2.10	Top view of the suspension model in Mathematica with common platform mass shown as semi-transparent for clarity. Resonance mode 3 at 0.723078 Hz. Here we see the dominant optic stage transverse motion coupling to top mass yaw motion that in turn couples to yaw motion of the bottom stages. . .	67
2.11	Cross-coupling transfer functions from longitudinal excitation at the suspension point to (left) optic vertical displacement and (right) optic pitch rotation. Each panel shows magnitude and phase as a function of frequency, evaluated for the common-platform suspension model.	67
2.12	Mode visibility for the common-platform triple suspension with COM of the common platform at the centre of the top mass.	68
2.13	Comparison of cross-coupling transfer functions for the common-platform suspension model with the top-mass centre-of-mass (COM) <i>centred</i> and with the COM <i>lowered</i> relative to the top-mass reference plane. In each panel, the magnitude and phase are overlaid for the two configurations (baseline COM-centred model and COM-lowered model), illustrating how the COM shift modifies the coupling pathways from the applied disturbance to the optic motion for pitch motion improving the performance.	69
2.14	Overlay of mode-visibility scatter plots for a set of lower stages length configurations. Opacity indicates the length setting for each configuration (from lowest length to highest length for the bottom stage), the total length of the suspension is maintained constant.	70

2.15	Variation of the roll and bounce mode frequencies in the observation band as a function of the final suspension-stage length, with the upper-stage lengths held fixed. The observation band contains two degenerate roll modes and two degenerate bounce modes. As the final-stage length increases, the out-of-phase bounce mode becomes increasingly cross-coupled to pitch motion of the upper mass, and the annotation in the plot indicated the mode-visibility value for this mode.	71
2.16	Comparison of cross-coupling transfer functions for the common-platform model as the lower-stage lengths are varied. Each panel overlays magnitude and phase for five length sets ($l_1 : l_2 : l_3$): 0.290 : 0.290 : 0.290, 0.290 : 0.270 : 0.310, 0.290 : 0.230 : 0.350, 0.290 : 0.190 : 0.390, and 0.290 : 0.150 : 0.430 m. The transfer functions shown are (a) Structure longitudinal excitation to optic vertical displacement, (b) Structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation.	72
2.17	Overlay of mode-visibility scatter plots for several common-platform mass configurations. Each point marks the visibility of a normal mode at its eigen-frequency; colour encodes DOF and marker encodes where the mode is sensed (platform, cross-coupled, or intermediate mass). Opacity monotonically maps the mass setting from lightest (most transparent) to heaviest (least transparent), enabling direct visual comparison of frequency shifts and visibility changes across configurations.	73
2.18	Comparison of cross-coupling transfer functions for the top-heavy common-platform configurations, parametrised by the common-platform mass m_1 (8.483, 12.725, 16.966, and 21.208 kg). Each panel overlays the magnitude and phase for the corresponding coupling path: (a) Structure longitudinal excitation to optic vertical displacement, (b) structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation.	75
2.19	Common Platform Model highlighting asymmetry in d_4 between the chains.	77

2.20	Comparison of cross-coupling transfer functions for the common-platform suspension model under d_4 -value asymmetry between the two suspension chains. The symmetric baseline ($\Delta d_4 = 0$) is compared against three asymmetric cases: $\Delta d_4 = -1$ mm, $\Delta d_4 = +1$ mm, and $\Delta d_4 = +5$ mm. Each panel overlays the magnitude and phase for (a) structure longitudinal excitation to optic vertical displacement, (b) structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation. The transfer functions are taken to the end chain in (a) and (b) and input chain in (c). . . .	78
2.22	Example of mode-shape localisation for $\Delta d_4 = +5$ mm, where pitch/roll motion becomes more strongly associated with one optic chain, reducing the degree of common-mode rejection.	78
2.21	Differential (common-mode rejected) cryostat seismic length-noise projection for the common-platform model under Δd_4 asymmetry, overlaid with the Cryo-Triple large-upper-blade (LUB) reference model [81].	79
2.23	Preliminary CAD rendering of the cryostat showing the current work-in-progress integration of the common-platform suspension assembly. The model is primarily used for envelope, clearance, and access checks, and does not represent a final mechanical design.	82
3.1	Closed-loop LED driver circuit using the ADA4891 operational amplifier and an N-channel MOSFET for low-side current regulation. The op-amp senses the voltage across the 0.2Ω current-sense resistor and adjusts the MOSFET gate drive to regulate LED current based on the <i>Drive_Input</i> control voltage. The feedback loop ensures precise LED current control independent of LED forward voltage or supply variations, making it suitable for optical sensing applications such as shadow sensors.	91
3.2	Comparison of measured and simulated noise spectral densities of the QPD amplifier output: The traces show noise measurements for two output channels (AB-CD and BD-AC) obtained using a Moku spectrum analyser, the modelled electronic noise from LTSpice simulations, and the baseline noise floor of the Moku. The simulation aligns well with the measured spectra above 10 Hz, validating the noise performance of the QPD circuit and we see many ground loops on top of the baseline curve.	95

3.3	Rectangular flag moving across a QPD and intercepting the incident beam—the green highlighted edge shows the leading edge.	102
3.4	Concept and response of shadow sensing with a rectangular flag and QPD. (a) Schematic illustration of the flag partially blocking the incident Gaussian beam on a QPD. (b) Corresponding simulated normalised and raw differential signal measured from the QPD as the flag translates across the beam.	105
3.5	Heat map characterisation of a QPD shadow sensor with a rectangular flag. Axes are normalised to the QPD radius R and the total beam power is held fixed.	107
3.6	Concept and response of shadow sensing with a rectangular flag and QPD. (a) Schematic illustration of the flag partially blocking the incident Gaussian beam incident on the QPD. (b) Corresponding simulated normalised and raw differential signal measured from the QPD as the flag translates across the beam.	112
3.7	Heat map characterisation of a QPD shadow sensor with a circular aperture flag. Axes are normalised to the QPD radius R and the total beam power is held fixed.	114
3.8	Circular aperture flag: simulated QPD response for varying beam radius w and aperture diameter d in the region of interest of Figure 3.7 (max transmission, low shot-noise floor). Normalised (solid) and raw (dashed) signals are plotted versus the flag-centre position, normalised to the QPD active radius.	117
3.9	Rectangular aperture flag: simulated quadrant-photodiode response for variations in flag width and beam radius, for operating points taken from the thermal maps (Figure 3.5). Normalised (solid) and raw (dashed) signals are plotted against the flag-centre position, normalised to the QPD active radius R	119
3.10	Characterisation of the GOSEM in air. (a) Optical/mechanical setup with LED illumination, flag aperture, and QPD detection. (b) Measured raw and normalised sensing curves with local-slope responsivities.	122

3.11	Relative intensity noise (RIN) spectra of the GOSEM output shown in a single plot. Traces display the normalised RIN with the flag displaced (green), the raw RIN with the flag displaced (orange), the RIN with the flag centred (blue), the calculated shot-noise limits for the displaced and centred configurations (dashed lines), and the dark-noise floor of the readout electronics (black). The comparison illustrates that normalisation suppresses excess LED intensity fluctuations when the flag is displaced, lowering the apparent noise floor toward the shot-noise level. The right-hand axis converts voltage spectral density to equivalent displacement noise using the measured responsivity.	123
3.12	Cryogenic setup for component-level testing of the shadow sensor. The image shows the integration of the FCI InGaAs Q3000 photodiode (QPD1 and QPD2 systems) and the LED1200L sources inside the cryostat. Dedicated temperature monitors are installed near both QPDs and the LED system to track thermal conditions during cooldown to 5 K. QPD2 is additionally configured for tests with an LED beam injected from outside the cryostat.	127
3.13	Comparison of forward I-V characteristics for two near-infrared LEDs measured during cool-down in the cryostat. The Thorlabs LED1200L exhibits abnormal behaviour and poor efficiency at low temperatures, whereas the Hamamatsu L13072-0120L performs more reliably and is better suited for cryogenic operation.	129
3.14	Comparison of LED optoelectronic efficiency across temperature. Subfigures (a) and (b) show η for the Thorlabs LED1200L and Hamamatsu L13072-0120L, respectively.	131
3.15	Thorlabs LED relative intensity noise. (a) Broadband spectra across temperatures; (b) temperature trend at a representative line frequency of 117 Hz. Both panels use the same normalisation to total QPD optical power.	133

3.16	Comparison of relative intensity noise (RIN) of the Hamamatsu and Thorlabs LEDs measured under cryogenic conditions. Each curve corresponds to the unbiased (grounded) configuration of the shadow sensor at different baseplate temperatures. The solid lines represent the Hamamatsu LED and the dashed lines represent the Thorlabs LED, with colour indicating the corresponding baseplate temperature. The plot illustrates that the Thorlabs LED 1200L exhibits higher RIN levels across the measured frequency range.	134
3.17	Amplitude spectral density (ASD) of the QPD differential dark noise (AB – CD) versus frequency for temperatures from 5 to 287 K (colour-coded; plasma map, see bar). Solid lines denote biased operation ($V_{\text{bias}} = 2.3 \text{ V}$); dashed lines denote unbiased. Spectra are estimated with a Welch periodogram (Hanning window) and then downsampled onto logarithmic frequency bins using the per-bin median to suppress narrow spectral lines, revealing the broadband floor. The grey trace shows a room-temperature electronics-noise reference with the input shorted. The plot highlights the temperature- and bias-dependence of the low-frequency noise, while the broadband floor above $\sim 100 \text{ Hz}$ varies only moderately across temperatures.	136
3.18	Temperature dependence of the QPD responsivity. The quantity plotted is $R = V_{\text{QPD}} / (G P_{\text{LED}})$ with $G = 330 \text{ k}\Omega$, i.e. the QPD output voltage normalised by the LED optical power and the transimpedance gain. Each curve is then normalised to its room-temperature point so that $R(295\text{--}300 \text{ K}) = 1$. Two independent measurement runs (different LED powers) are shown; markers denote the measured points and lines guide the eye. The responsivity increases with mount temperature and approaches its room-temperature value near 300 K.	137
3.19	Transfer-function measurements of the QPD mounted next to the cryostat window. (a) Representative 5 K sweep highlighting the flat passband used for characterisation. (continued on next page)	138
3.21	Image showing photograph of the flag used in the setup in Figure 3.20	141

3.20	Photograph of the cryostat used for helium testing of the shadow sensor. The setup includes the QPD holder, the flag mounted on a translation stage moved via a motion feedthrough, a heater for controlled thermal loading, cold anchors for thermalisation, and temperature sensors placed near the LED and on the cold plate. The flag is a circular aperture that has an aperture diameter of 2.5 mm.	141
3.22	Cryogenic shadow-sensing response of the QPD (flag diameter 2.5 mm). The black curve shows the raw differential voltage and the red curve is normalised by the total sum signal to compensate LED intensity changes. Dashed lines indicate local linear fits (tangents) to the normalised curve: at the centre the responsivity is -4.085 V/mm , and at -1.00 mm it is -15.179 V/mm	142
3.23	Relative intensity noise (RIN) spectra measured at 47 K for the shadow sensor. The orange and green traces correspond to the raw and normalised responses when the flag is displaced from the centre, respectively. The normalisation effectively suppresses LED intensity fluctuations, reducing the low-frequency excess noise towards the shot-noise limit (dashed lines). The blue curve shows the RIN with the flag centred, while the black trace represents the dark noise floor of the readout electronics. The equivalent displacement noise can be read from the right-hand axis where the voltage spectral density has been scaled using the measured responsivity of the sensor when the flag is centred. This gives a worse case upper limit for the displacement noise floor as the responsivity is lowest when the flag is centred.	143
3.24	IRLED carrier assembly. Left: cross-sectional view of the IRLED carrier. Right: exploded view showing the constituent parts.	144
3.25	Photodiode carrier assembly. Left: cross-sectional view of the photodiode carrier. Right: exploded view showing the constituent parts.	145
3.26	Exploded CAD views of the (left) LED carrier assembly and (right) QPD carrier assembly.	146
3.27	GOSEM assembly.	147

4.1	Simplified schematic of the Advanced LIGO interferometer operated with DC readout. A pre-stabilised laser is injected into the power-recycling cavity, formed by the power-recycling mirror (PRM) together with the input test masses (ITMs) and the recycling optics (PR2, PR3), to build up carrier power at the beamsplitter. The arms are Fabry–Pérot cavities formed by the ITMs and end test masses (ETMs). The differential arm-length signal exits at the antisymmetric port, where a small intentional offset provides carrier light as a local oscillator for DC readout. The output field is shaped by the signal-recycling cavity (SRM with SR2, SR3 and the ITMs), spatially filtered by the output mode cleaner (OMC), and detected at the photodiode (PD) [116, 104].	155
4.2	Standard balanced homodyne detection (BHD) configuration in the A+ interferometric gravitational-wave detector: This schematic we see how the local oscillator (LO) and interferometer output fields are combined at the BHD beam splitter. No vacuum ports are open, ensuring minimal quantum noise contribution. The output beams pass through orthogonal Output Mode Cleaners (OMC A and OMC B) and are detected at photodiodes PD1 and PD2. The phase shifter, which is basically a mirror, allows control over the quadrature being measured.	157
4.3	CAD of A+ Output Mode Cleaner Cavity: This figure highlights the bowtie geometry of the cavity. The beam enters through flat mirror 1 and exits through flat mirror 2. The grey surrounding the cavity is the bonding template used to guide the optics bonding process. The cavity consists of two flat mirrors and two curved mirrors. The curved mirror sub-assemblies are held against a couple micrometers to aid cavity alignment [94].	159
4.4	Components of the curved mirror subassembly: curved optic, piezoelectric transducer (PZT), and tombstone mounting prism.	163
4.5	Shift in the position of the curvature minimum of the curved optic manifests as wedging of the curved optic.	166
4.6	ZYGO curvature minimum measurement setup in the Downs Lauriston Labs, Caltech.	167

4.7	In the left graph we see the position vectors of curvature minimum for various clockings of the fiducial. In the graph on the right the curvature minimum positions after de-clocking is shown i.e. projecting the position vectors of various clockings to 12 o' clock. Red - 12 o' clock, green - 3 o' clock, blue - 6 o' clock, black - 9 o' clock.	168
4.8	Diagrams indicating PZT wedge direction and the length of the arrows indicate the amount of wedging. The number in the bracket is the wedge angle	170
4.8	Diagrams indicating PZT wedge direction and the length of the arrows indicate the amount of wedging. The number in the bracket is the wedge angle	171
4.9	Subassembly preparation and bonding process for the A+ LIGO Output Mode Cleaner curved mirrors using a bonding fixture. Left: Subassembly bonding step showing EP30 epoxy applied to the PZT surface. Middle: Close-up of a fully stacked subassembly consisting of a curved mirror, a PZT actuator, and a mounting prism held in a precision bonding fixture. Right: Four completed subassemblies placed inside an oven for curing. This subassembly is later mounted on the optical breadboard along with the flat mirrors to form the full cavity.	171
4.10	Scatter map for three of the A+ OMC curved mirrors. Measurements and plots by Dr. Liyuan Zhang, LIGO Lab, Caltech [183].	174
4.11	Example curved mirror subassembly glueing instruction sheet for A+ OMC curved mirror subassembly.	175
4.12	Inspection and cleaning of the OMC glass breadboard using high-purity IPA wipes and halogen lighting to remove and detect particulate and fibrous contaminants prior to template bonding.	177
4.13	Cavity bonding template mounted and aligned on the OMC glass breadboard. The template is centred such that the cavity waist coincides with the beam waist, and mirrors are placed in their designated cutouts.	178
4.14	Bonding template setup used for positioning the curved cavity mirrors (CM1 and CM2). The micrometer adapter plates are adjusted to set the nominal cavity length, and the mirrors are placed such that they simultaneously contact the micrometers and template pads. Visual and tactile checks are performed to ensure proper seating and alignment before final securing.	179

- 4.15 Fitting of transfer function between excitation of the EOM and the OMC transmission PD: As the EOM sideband is scanned, the various HOM sidebands beat with the carrier. When the carrier sideband gets closer to FSR, the transfer function sees a dip in frequency which is seen in the plot above between 250 MHz and 300 MHz. This scan is done with a pitch misalignment thus the HOMs in the plot are TEM_{x0} modes for $x = 1,2,3,..$. The OMC mirrors are bonded here (*Courtesy: M. Nakano, A+ OMC team*). 185
- 4.16 Higher Order Mode (HOM) distribution in the optical cavity, showing the Gouy phase (expressed in units of FSR) as a function of mode order. Each colour represents a different RF sideband frequency used for controlling the interferometer and auxiliary optical cavities in LIGO. The data corresponds to the configuration where PZT1 is held at 0 V and PZT2 at 50 V, resulting in measured mode spacings of $TMS_y/FSR = 0.21983$ for yaw degree of freedom and $TMS_p/FSR = 0.21950$ for pitch degree of freedom. The vertical black line at the centre marks the laser fundamental mode carrier resonance frequency. This measurement was taken immediately after bonding the optics on the cavity. 186
- 4.17 Transverse mode distribution as a function of PZT drive voltage, showing the frequency locations of higher-order modes (HOMs) and their sidebands relative to the fundamental mode. Each coloured band corresponds to a specific HOM and its frequency shift as the cavity length is varied through PZT actuation. The vertical spread of each band arises from cavity astigmatism, which causes the transverse mode spacing (TMS) to vary between the sagittal and tangential Gouy phase limits. The black dashed lines indicate the minimum required frequency separation from the TEM_{00} carrier (two linewidths) for reliable OMC operation. The green dashed region marks the optimal PZT range where modal degeneracies are avoided while satisfying this separation criterion—labelled as the "useful PZT range" on each plot. The top panel corresponds to length tuning via PZT1, and the bottom panel to tuning via PZT2. 187

4.18	Higher Order Mode (HOM) distribution in OMC cavity with roundtrip length of 111.7 cm: In this plot, we see the Gouy phase (in units of FSR) as a function of mode order. Each colour corresponds to a different RF sideband frequency used in the LIGO interferometer control scheme. In this configuration, PZT1 is held at 0 V and PZT2 at 50 V, resulting in measured mode spacings of $TMS_y/FSR = 0.2191$ and $TMS_p/FSR = 0.2186$. This HOM configuration is not suitable for OMC operation due to a coincidental resonance of the 9th-order mode sideband with the fundamental mode, which violates the requirement of avoiding higher-order mode overlap within at least two linewidths of the fundamental resonance. Such modal degeneracy can lead to excess power leakage through the OMC and reduced filtering performance.	189
4.19	DC Response of PZT Actuators. (a) and (b) correspond to measurements for PZT 1 and PZT 2, respectively. The plots on the left show cavity transmission as a function of PZT drive voltage, scanning across several free spectral ranges (FSRs). Prominent peaks correspond to the fundamental TEM_{00} mode and higher-order modes (HOMs). Separate traces are shown for upward and downward voltage sweeps. Using the known cavity geometry to determine the FSR and transverse mode spacing (TMS), the plots on the right convert this data into PZT displacement versus applied voltage, characterising the DC response of each actuator (<i>Courtesy: M. Nakano, A+ OMC team</i>).	193
4.20	Cavity beam spot position verification by imaging beam spots on the cavity mirrors. The cavity axis appears to be well centred by eye, with beam spots located within 1 mm of the mirror centres.	195
4.21	OMC 101 mirrors bonded.	196
5.1	OMC Circulating Power: This plot shows exaggerated birefringence effects in an optical cavity. $\delta\mathcal{F}/\mathcal{F}$ is 0.1 and $\delta f_0/\gamma$ is 0.5. The measured cavity birefringence for the A+ OMC cavity number 6 are 0.025 (0.0005) and 0.0038 (0.0012) respectively.	201

5.2	Schematic of the A+ LIGO interferometer output with polarisation homodyne readout chain: A local oscillator (LO) field is picked off from the input beam path and reflected off of a phase shifter mirror and a HWP to control its phase angle and power relative to the . This LO field is then recombined with the signal field E_{IFO} at a PBS. The resulting beams are filtered through a single OMC to remove higher-order spatial modes and control sidebands. Balanced homodyne detection (BHD) is performed using a PBS. The red and yellow paths represent the LO and signal fields, respectively, with polarisation states shown schematically at various points.	204
5.3	After the OMC, a half-wave plate (HWP) is inserted to rotate the polarisation angle by 45° prior to balanced homodyne detection (BHD), which is performed using a PBS and a pair of photodiodes (P1 and P2). Red and yellow lines represent the LO and signal paths, respectively, with polarisation states indicated schematically.	206
5.4	Simplified schematic of the A+ LIGO interferometer output with polarisation homodyne readout chain: The main interferometer consists of two orthogonal arm cavities formed by input (ITM) and end test masses (ETM), with the beamsplitter (BS) at the center. The power recycling cavity, formed by PRM, PR2, and PR3, enhances circulating laser power, while the signal recycling cavity, formed by SRM, SR2, and SR3, shapes the frequency response of the interferometer. A local oscillator (LO) field is picked off from the input beam and directed through a phase shifter mirror and a power-adjusting HWP-PBS combination before being combined with the signal field E_{IFO} at a polarizing beamsplitter (PBS). The combined beam is filtered by a single output mode cleaner (OMC), which suppresses higher-order spatial modes and unwanted sidebands. Before and after the OMC, half-wave plates (HWP) are inserted to rotate the polarisation angle by 45° , combined, prior to balanced homodyne detection (BHD), which is performed using a PBS and a pair of photodiodes (P1 and P2). Red and yellow lines represent the LO and signal paths, respectively, with polarisation states indicated schematically.	214

5.5	<p>Experimental setup to measure cavity birefringence: Schematic of the optical setup used for characterising the output mode cleaner (OMC) on the transport fixture. A 1064 nm Innolight laser (output power ~ 60 mW) is passed through a Faraday isolator to prevent back-reflections, and then phase modulated at 32 MHz using a resonant electro-optic modulator (EOM) for Pound-Drever-Hall (PDH) locking. The beam is mode-matched into the OMC using a telescope consisting of two plano-convex lenses ($f = 35$ mm and $f = 125$ mm). A PBS shown downstream the EOM cleans the polarisation entering the cavity. A BS directs a small fraction of light to a power reference photodiode for monitoring input power. The main beam is passes through a PBS to clean the polarisation. The input polarisation is adjusted at 45 degree angle using a combination of a half-wave plate (HWP_0) and a quarter-wave plate (QWP). After interacting with the OMC cavity, the reflected light is directed to a reflection photodetector and a CCD camera for PDH locking and alignment diagnostics, respectively. The transmitted beam is analyzed using a balanced homodyne detection (BHD) setup comprising a polarizing beam splitter (PBS) and two photodiodes (PD_1 and PD_2). The OMC cavity optics are mounted on a glass breadboard which is held in a dedicated transport fixture to maintain mechanical stability and optical alignment.</p>	219
5.6	<p>Normalised p-polarisation cavity resonances. Each trace shows the photodiode signal for a single resonance window after converting the fitted drive voltage to a normalised detuning $f = (V - V_0)/\gamma$. γ here was extracted by calculating the difference in voltage at half maximum of the resonance peak for every peak. The data used in the fitted using Lorentzians $L(f) = A/(1 + f^2)$. The dashed line at $f = 0$ marks the fitted resonance centre. The per-peak voltage-domain normalisation collapses distinct scans onto a common lineshape suitable for averaging and comparison.</p>	221

- 5.7 **Differential signal fit $L_p - L_s$.** The measured differential photodiode signal versus the normalised detuning f . The scatter is fit with offset $+ \frac{A}{1+f^2} - \frac{A + \delta A}{1 + [(1 + \beta)(f - \delta f_0)]^2}$, where δf_0 is the relative p-s resonance offset, β encodes the linewidth (finesse) asymmetry, and δA captures contrast (amplitude) mismatch. This parameterisation separates resonance spacing, finesse difference, and contrast into orthogonal fit knobs for subsequent statistical analysis. The data is averaged over 54 peaks. 221
- 5.8 **Distributions of birefringence and lineshape parameters from OMC cavity scans.** For each resonance window the p-peak is fit in voltage (see Figure 5.6) to obtain (V_0, γ, A) , and the differential model offset $+ \frac{A}{1+f^2} - \frac{A + \delta A}{1 + [(1 + \beta)(f - \delta f_0)]^2}$ (see Figure 5.7) yields (clockwise from top left) the DC *offset* of the differential channel, the polarisation resonance separation δf_0 (reported in MHz via the cavity calibration), the finesse difference $\Delta \mathcal{F} = \beta \mathcal{F}_0$ with $\mathcal{F}_0 = 378$, and the contrast mismatch δA . Each panel annotates the sample mean μ and standard deviation σ of the inlier set, as printed on the plots. 222
- 5.9 **Tabletop setup for polarisation BHD** The input laser setup is same as that used for birefringence measurement in Figure 5.5. Waveplate stack HWP₀-QWP provides control of the LO-IFO relative phase (homodyne angle) and also helps maintain a similar power ratio between the two orthogonal polarisations across measurements. Downstream, HWP₁ and HWP₂ before and after the OMC balance the rotation before and after the OMC enabling cancellation of birefringence-induced noise while preserving net 45 degrees rotation. 223
- 5.10 Characterisation of the HWP₀-QWP angle pairs that realize the input power ratio $P_p : P_s = 10:1$. Plotted are the measured QWP angles θ_2 versus HWP₀ angles θ_1 (two measurement series). Angles are reported relative to calibration offsets $\theta_1 = \text{HWP}_0 - \theta_{1,0}$ and $\theta_2 = \text{QWP} - \theta_{2,0}$ 224
- 5.11 Phase difference δ between the *s*- and *p*-polarized outputs evaluated at the waveplate settings from Figure 5.10. The phase is inferred from the balanced-homodyne relation $P = 2\sqrt{P_p P_s} \cos \delta$ with $P \equiv E_y - E_x$ and $(P_p, P_s) \equiv (P_{\text{LO}}, P_{\text{IFO}})$. Two measurement series are shown. 224

5.12	Output powers P_1 and P_2 versus HWP0 angle θ_1 after the HWP–QWP–HWP stage set to 22.5° . Curves show the modeled intensities $ E_x ^2$ (p) and $ E_y ^2$ (s) for two series; markers show the corresponding measured channel fractions (data).	225
5.13	Optimal polarisation angle as a function of homodyne angle. The solid curve shows the model prediction for Θ_{opt} from Eq. 5.27, evaluated using the measured birefringence parameters. The markers denote the measured HWP ₁ angles that minimise the OMC length-to-homodyne noise transfer at the indicated homodyne angles.	226
5.14	Determination of the optimum HWP ₁ mixing angle for a range of homodyne angles. Each panel shows the measured peak level of the driven line (in dBm) versus the HWP ₁ FA angle, while HWP ₂ is counter-rotated so that the homodyne angle remains fixed at the value indicated in the panel (θ_{HD}). The minimum in each curve identifies the HWP ₁ setting that minimises the OMC length-to-homodyne signal noise transfer for that θ_{HD}	227
6	Photograph of the LED driver printed circuit board. The circuit diagram for this is found in Figure 3.1	237
7	Photograph of the QPD amplifier printed circuit board.	238
8	Calibration of HWP0, the power ratio half wave plate. The normalised photodiode signals in the s - and p -polarisation channels are shown as a function of the HWP0 rotation angle. The data are fitted with the expected sinusoidal responses, $f_s(\theta) = A_s \sin^2(2(\theta - \theta_{0,s})) + \text{off}_s$ and $f_p(\theta) = A_p \cos^2(2(\theta - \theta_{0,p})) + \text{off}_p$, confirming the orthogonal mixing of the two polarisations. Vertical dashed and dotted lines indicate the fitted extrema (peaks/dips) within the scan range, which define the effective fast-axis reference angles used for subsequent polarisation-BHD alignment and balancing.	239

9	Calibration of the quarter-wave plate (QWP), the power ratio quarter wave plate. The photodiode signal after the PBS is shown as a function of QWP rotation angle. The data are fitted with the expected periodic dependence of the transmitted power for an initially linearly polarised field passing through a QWP and analysed by a PBS, $f(\theta) = 2A \sin^2(\theta - \theta_0) \cos^2(\theta - \theta_0) + \text{off}$. The vertical dashed and dotted lines indicate the fitted extrema (dip and peak) within the scanned range, which define the effective QWP fast-axis reference angle used for setting the desired polarisation state in the polarisation-BHD experiment.	240
10	Calibration of the first mixing half-wave plate (HWP1)	240
11	Calibration of the second mixing half-wave plate (HWP2)	241
12	Calibration of the first mixing half-wave plate (HWP1) in the presence of all the other waveplates upstream i.e. HWP0 and QWP, and other optical components (see Figure 5.9)	241
13	Calibration of the first mixing half-wave plate (HWP1) in the presence of all the other waveplates upstream i.e. HWP0, QWP and HWP1, and other optical components (see Figure 5.9)	242
14	Common platform suspension modes (1–6).	246
15	Common platform suspension modes (7–12).	247
16	Common platform suspension modes (13–18).	248
17	Common platform suspension modes (19–24).	249
18	Common platform suspension modes (25–30).	250

Acknowledgements

My journey in gravitational waves began in the summer of 2018, during my second year as an undergraduate, when I walked into Prof. Umakant Rapol's lab with an enthusiasm for quantum optics. He was sceptical that my excitement would endure, and gave me one week to find out for myself — I went back. Witnessing matter wave interferometry in action, and experiencing firsthand how rarely things behave the way textbooks promise, was genuinely exhilarating; there was something deeply compelling about coaxing a small, broken corner of the universe back into working order. Prof. Rapol then told me he was building a laboratory for LIGO India training, that I would be its first student, and asked whether I would like to work on constructing a Fabry-Perot cavity to stabilise a 1064 nm laser — I said yes without hesitation, even though I did not fully appreciate what gravitational waves were until I took courses in Cosmology and General Relativity a year later. I am deeply grateful to Prof. Rapol for placing those first stepping stones, for introducing me to Prof. Giles Hammond, and for jointly supporting my visit to the University of Glasgow in the summer of 2019. I would also like to sincerely thank Dr. Suresh Doravari of IUCAA for his steady guidance and for ensuring I was always moving in the right direction. Without these three individuals, I would not be where I am today.



I would like to express my deep gratitude to Prof. Giles Hammond, my primary supervisor, for his behind-the-scenes support and for teaching me to approach every project by returning to first principles. My most profound thanks, however, go to Dr. Bryan Barr and Andrew Spencer, who guided me at every step of my PhD. Bryan was instrumental in ensuring smooth transitions between my visits to Caltech and my work on Glasgow

projects, and his practical support in the laboratory — from tracking down suitable lenses for the shadow sensor to noise measurements — was invaluable. I find myself reminiscing about our thought-provoking conversations every time I encounter a whiteboard, whether in rooms 253 or 252b, or anywhere else. I would also like to thank Andrew for his support and discussions, going above and beyond despite not being my official supervisor. Without these two people, I think I would have felt utterly lost across all those projects. I hold your support during my PhD in the highest regard.

I would like to thank the Caltech A+ LIGO team for being such an exceptional group to work with. I am grateful to Dr. Calum Torrie and Dr. Gabriele Vajente for recognising contributions with humility and generously supporting multiple visits to Caltech. I would like to thank Dr. Koji Arai for his thought-provoking discussions and steadfast support across those visits — I will never forget how graciously he accommodated my last-minute requests, both inside and outside the laboratory, without a moment's hesitation. I would like to extend my deepest thanks to Dr. Masayuki Nakano for being such an exceptional colleague, mentor, and friend. Masayuki helped me see things differently in the lab, through beautiful graphs, elegant equations, and — on at least one occasion outside of work hours — sake. I look forward to future collaborations, and to finally writing that paper on polarisation balanced homodyne detection together.

I would like to thank Stephen Appert for being an outstanding manager of the A+ OMC team during my visits, as well as a wonderful friend outside of work. I am grateful for all the support he offered, both in and out of the workplace. He taught me to pay close attention to fine details and to approach even the most seemingly obvious of things with a sceptical eye. His emphasis on the importance of thorough documentation and careful note-taking is a lesson I carry with me. I would also like to thank Dr. Camille for being an amazing colleague.



I would like to thank the wonderful experimental group at Glasgow for their help with laboratory work and for truly fruitful chats — even the small ones left a great impression on me: Giles Hammond, Ian Martin, Peter Murray, Mark Barton, Bryan Barr, Andrew Spencer, Karl Toland, Graeme Eddols, Jack Callaghan, Gregoire Lacaille, Jennifer Docherty, and Ardiana Nela. I would also like to thank Jamie Scott, Colin Craig, and Stevie Craig for their help with IT and in the mechanical workshop respectively. I would like to thank Joe Bayley for very useful discussions on Markov Chain Monte-Carlo.

I would like to thank Dr. Abhinav Prasad for his help with electronic circuits, valuable discussions about sensors, and assistance with troubleshooting electronics. I would like to thank Dr. Mark Barton for his unparalleled insights on suspension dynamics and his mentorship in suspension modelling. I would like to thank Russell Jones for his support with engineering aspects, CAD modelling, and helping realise a suitable package for the shadow sensor, among a long list of other contributions. Thanks to Karl Toland and Abhinav Prasad for those invaluable conversations that helped me navigate difficult decisions during trying times.

Moving halfway across the world to a place very different from anything I had known came with its share of challenges. I am deeply grateful to everyone below for making that transition smooth and for making me feel at home so far from it. Special thanks to the first generation of 253 colleagues: Ross Johnston, Federico Stachurski, Eungwang Seo (Harry), Jess Irwin, Leigh Smith, Naren Nagarajan, Storm Colloms, Christian Chapman-Bird, and Victoria Graham. Thanks to all the generations of members of room 253 (August, Natalia, Josh, Elizabeth, Varun) for putting up with my relentless coffee grinding. Thanks to Jack Callaghan for workout breaks and for being a true foodie buddy, and to Phoebe for being a great friend and for her help in securing my new position at Strathclyde. To the book club gang — Ben, Maddie, Ryan, Sarah, Rob, Hector, Gav, Kyrie, and Lewis — thanks for everything, you guys are just amazing. Thanks also to the badminton crew — Ryan, Tyler, Wrishik, Dwayne, Dev, and Ardiana — for making my PhD life spicy with fiercely competitive baddie. You are all the core reason I am happy I chose Glasgow and continue to do so. aRDiana, thank you for being part of this journey. I said I would



finish, and I did. Soon it will be your turn — and you will do great. I would also like to thank all my friends in California who made my time there so memorable: Radhika Bhatt, Madeline Lesovsky, Alexey Yermakov, Simon Tait, Ian Macmillan, Stephen Appert, Masayuki Nakano, Srinath, and James Gardener. To long lasting friendships.

Finally my family, to my parents, grandparents, Yashu and Usha — there are no words that could adequately describe how important you are to me. Thank you for supporting me throughout, for making possible everything I have ever tried to achieve, and for being there for me always, unconditionally.



Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution

Thejas Seetharamu

Fundamentals of Gravitational Wave Detection

1.1 Gravitational-Wave Astronomy

The notion that gravity allows disturbances that propagate through space has roots that predate general relativity. In the 1870s, Clifford advanced the geometric idea that what we perceive as matter and gravitation could be understood as variations in the curvature of space, implicitly inviting the possibility of dynamical (time-dependent) curvature rather than instantaneous action-at-a-distance [60]. A more explicit wave picture appears in Poincaré’s relativistic discussions of gravity, where he argued that a Lorentz-consistent theory suggests retarded gravitational influence and “ondes gravifiques” propagating at the speed of light [129]. Gravitational waves entered their modern form with Einstein’s development of the theory of general relativity; in 1916 he showed that the field equations admit wave-like solutions—small perturbations of spacetime that propagate outward from accelerating masses—establishing gravitational radiation as a generic prediction of the theory [66].

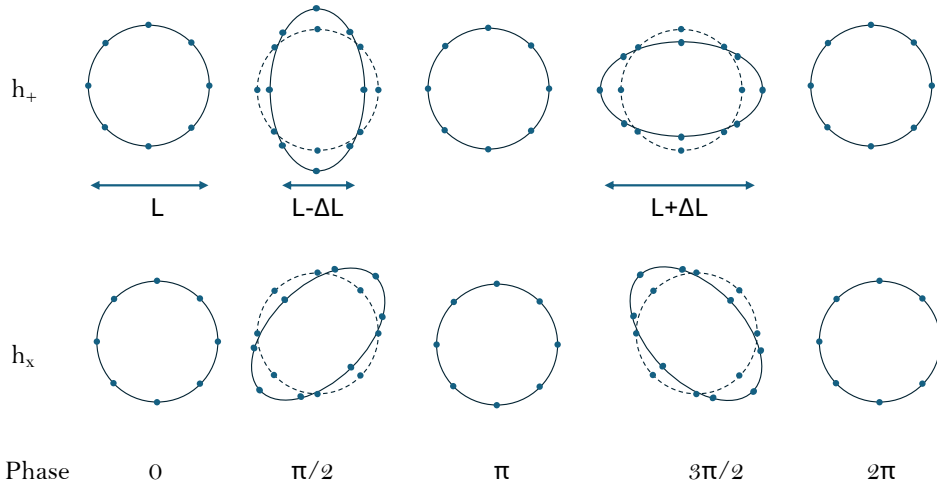


Figure 1.1: Ring of freely falling particles in a GW that stretches and squeezes space-time. This figure illustrates two different (orthogonal) polarisation of GWs. These time series images depict how space-time curvature changes for $\omega t = 0, \pi/2, \pi, 3\pi/2$. in the plane of the paper for an incident GW that propagates out of the page.

Within general relativity, gravity is not treated as a force but as the curvature of space-time caused by energy and momentum [66, 147] of an object. Gravitational waves are produced by accelerating mass distributions with a time-varying quadrupole moment and induce a characteristic tidal stretching and squeezing of freely falling test particles, corresponding to two independent transverse polarisations [147, 65].

Figure 1.1 illustrates the effect of a passing gravitational wave on a ring of freely falling test particles. As the wave propagates perpendicular to the plane of the ring, it induces a time-dependent tidal distortion, alternately stretching and compressing the separation L between particles along orthogonal directions. The dimensionless gravitational-wave strain is defined as the fractional change in separation,

$$h(t) = \frac{\Delta L(t)}{L}, \quad (1.1)$$

where L is the unperturbed separation and $\Delta L(t)$ is the wave-induced differential displacement. The measurement of ΔL enables detection of gravitational waves.

1.2 Sources of gravitational waves

Gravitational waves can, in principle, be emitted by any accelerating mass distribution with a time-varying quadrupole moment. Everyday objects therefore generate gravitational radiation, but the resulting strains are extraordinarily small and far beyond the reach of current detectors. Detectable signals require compact, relativistic sources with large masses and strong accelerations, typically of astrophysical or cosmological origin.

In the wave zone far from an isolated source [48], the leading-order metric perturbation in general relativity can be written in terms of time derivatives of the source multipole moments. Unlike electromagnetism, where dipole radiation is generic, gravitational radiation has no monopole or dipole contribution for isolated systems. The mass monopole corresponds to the total mass–energy M , which is conserved, so $\dot{M} = 0$ and it cannot radiate. The mass dipole moment is proportional to the centre-of-mass position; its first time derivative is the total linear momentum, which is also conserved for an isolated system. As a result, the leading radiative term is the mass quadrupole moment Q_{ij} , giving the familiar quadrupole formula

$$h_{ij}^{\text{TT}}(t, r) \simeq \frac{2G}{c^4 r} \ddot{Q}_{ij}^{\text{TT}}\left(t - \frac{r}{c}\right), \quad (1.2)$$

where r is the source distance and TT denotes the transverse–traceless projection [149]. This immediately shows why strong GW sources require large masses undergoing rapid, non-spherically-symmetric acceleration.

The first direct detection of gravitational waves was made on 14 September 2015 and originated from the merger of two stellar-mass black holes [7]. Gravitational-wave sources are commonly grouped into three broad families, described in the following subsections.

1.2.1 Compact Object Binaries

Coalescing binaries of compact objects are among the most important sources for direct gravitational-wave (GW) detection. They produce strong, repeatable signals whose main features are well understood from general relativity, and they populate the observing band of ground-based interferometers for a wide range of masses and distances [145, 14]. The observed waveform depends on intrinsic source properties (component masses and spins, and—when neutron stars are involved—finite-size matter effects) and extrinsic parameters (luminosity distance, sky location, and the inclination of the orbital plane relative to the line of sight). These factors set both the signal amplitude and how long the signal remains in-band.

Across all compact-binary classes, the emission mechanism is the same: as the binary orbits, it loses energy and angular momentum through GW radiation, causing the orbital separation to shrink and the orbital frequency to increase. This produces the characteristic *chirp* in which both the GW frequency and amplitude rise with time. For stellar-mass sources, the inspiral can last from seconds to minutes in the sensitive band, followed by a rapid merger and the ringdown of the remnant compact object. The relative prominence of these stages depends mainly on the total mass: lower-mass binaries spend more time in inspiral, while higher-mass binaries merge at lower frequencies and may be observed primarily through their late inspiral and merger [145]. There have been over 200 detections/candidate events as shown in Figure 1.2 [6, 13, 6, 15, 6, 2].

The compact-binary population is typically grouped into three source classes:

- **Binary Neutron Star (BNS):** two neutron stars in orbit. BNS signals can remain in-band for relatively long durations and, at the highest frequencies, exhibit measurable finite-size effects due to tidal deformation. The first observed BNS event was GW170817 [9].

- **Binary Black Hole (BBH):** two black holes in orbit. BBH mergers often produce high signal-to-noise transients dominated by the late inspiral, merger, and ringdown, and they currently make up the largest class of detections in the LIGO–Virgo–KAGRA catalogues [4, 14]. The first direct detection of gravitational waves was from a BBH merger event GW150914 from a luminosity distance of about 410 Mpc [7].
- **Neutron Star–Black Hole (NSBH):** a neutron star and a black hole in orbit. NSBH mergers have been observed in GWs (e.g. GW200105 and GW200115). Whether an electromagnetic counterpart is produced depends on whether the neutron star is tidally disrupted outside the black hole horizon, which in turn depends on the mass ratio and black-hole spin [16].

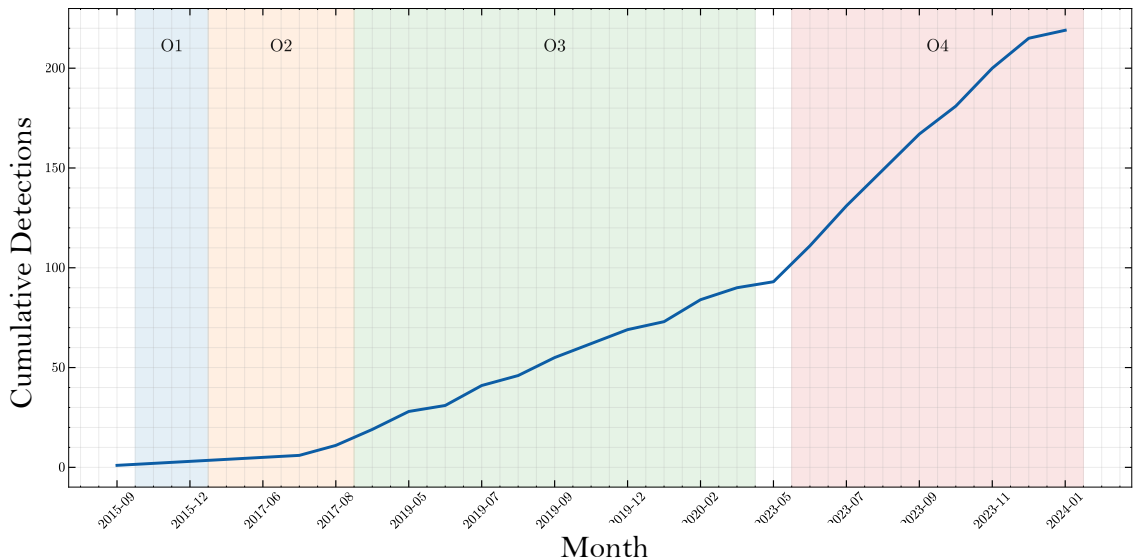


Figure 1.2: Cumulative number of gravitational-wave detections as a function of month (UTC). The curve is constructed by converting event GPS times to UTC, binning detections by month, and taking the cumulative sum. Shaded vertical ribbons indicate the LIGO–Virgo–KAGRA observing runs O1–O4 (csv file).

In the quadrupole approximation, the characteristic strain amplitude from a compact binary inspiral scales as [109, 145]

$$h(f) \sim 10^{-21} \left(\frac{\mathcal{M}_c}{1.2 M_\odot} \right)^{5/3} \left(\frac{f}{100 \text{ Hz}} \right)^{2/3} \left(\frac{100 \text{ Mpc}}{d} \right). \quad (1.3)$$

where \mathcal{M}_c is the chirp mass and d the luminosity distance to the source. For a binary neutron star system at $d \sim 100\text{Mpc}$ emitting near $f \sim 100\text{Hz}$, this corresponds to a characteristic strain $h \sim 10^{-21}$.

1.2.2 Continuous Sources

Continuous gravitational-wave (CW) sources produce long-duration, nearly monochromatic signals, most commonly expected from rapidly rotating neutron stars with a small non-axisymmetric deformation (e.g. “mountains” supported by the crust or magnetic stresses) or from long-lived oscillation modes [136]. In contrast to compact-binary chirps, CW signals are extremely weak and are therefore searched for by integrating coherently (or semi-coherently) over weeks to months while correcting for Doppler modulation from the Earth’s motion and for intrinsic spin evolution. To date, no CW signal has been confidently detected, and LIGO–Virgo–KAGRA searches have instead placed stringent upper limits on CW strain from both wide-parameter-space all-sky surveys (e.g. using O3 data) and targeted searches for known pulsars (including O4a analyses) [105, 3].

For a steadily rotating, non-axisymmetric neutron star, the intrinsic continuous-wave strain can be estimated as [136]

$$h_0 \simeq \frac{4\pi^2 G I \epsilon f^2}{c^4 d}, \quad (1.4)$$

which gives $h_0 \sim 10^{-26}$ – 10^{-25} for optimistic Galactic sources ($d \sim \text{kpc}$, $f \sim 10^2$ – 10^3Hz , $\epsilon \lesssim 10^{-6}$).

1.2.3 Burst Sources

Burst sources are short-duration gravitational-wave transients (from milliseconds to minutes) whose waveform is not known accurately or source is not known. Instead, “burst” searches typically look for statistically significant excess power and coherence across a detector network with minimal assumptions about signal morphology, making them sensitive to a broad range of unexpected or poorly modelled phenomena [11]. Astrophysical scenarios that may produce burst-like emission include core-collapse supernovae, magnetar giant flares/short bursts, and other violent, rapidly changing processes; dedicated targeted analyses use external electromagnetic triggers (e.g. optically observed supernovae or gamma-ray activity) to reduce background and improve sensitivity [164, 18]. To date, no confident burst detection from a non-compact-binary source class has been reported, and current results are expressed as upper limits on burst strain (and corresponding constraints on plausible emission models) from both all-sky and targeted searches [11, 164, 18].

1.2.4 Stochastic Sources

A stochastic gravitational-wave background (SGWB) is a persistent, noise-like GW signal formed by the superposition of many individually unresolvable sources (astrophysical foregrounds) and/or by processes in the early Universe. Experimentally, SGWB searches are performed by cross-correlating the strain data from separated detectors, since a genuine background produces a correlated contribution that accumulates with observing time while most instrumental noise remains uncorrelated between sites. Using Advanced LIGO and Advanced Virgo data through O3, no SGWB has been detected in the ground-based band and analyses instead place upper limits on the isotropic and anisotropic background amplitude [19][17]. At much lower frequencies (nanohertz), pulsar timing arrays have reported evidence consistent with a gravitational-wave background, commonly interpreted as arising from a population of supermassive black-hole binaries [23].

1.2.5 Notable Observations from LIGO

The landmark event GW150914 marked the first direct detection of gravitational waves and established compact-binary coalescences as observable astrophysical sources. The signal was consistent with the inspiral and merger of two stellar-mass black holes, confirming that such binaries exist in nature and demonstrating that the radiated gravitational-wave energy can be measured directly from the waveform. Parameter estimation indicated a binary black-hole origin with source-frame component masses of $\sim 36M_{\odot}$ and $\sim 29M_{\odot}$, producing a remnant black hole of $\sim 62M_{\odot}$ and dimensionless spin ~ 0.67 . This detection also enabled the first precision tests of the strong-field, dynamical regime of general relativity using gravitational-wave data alone [8].

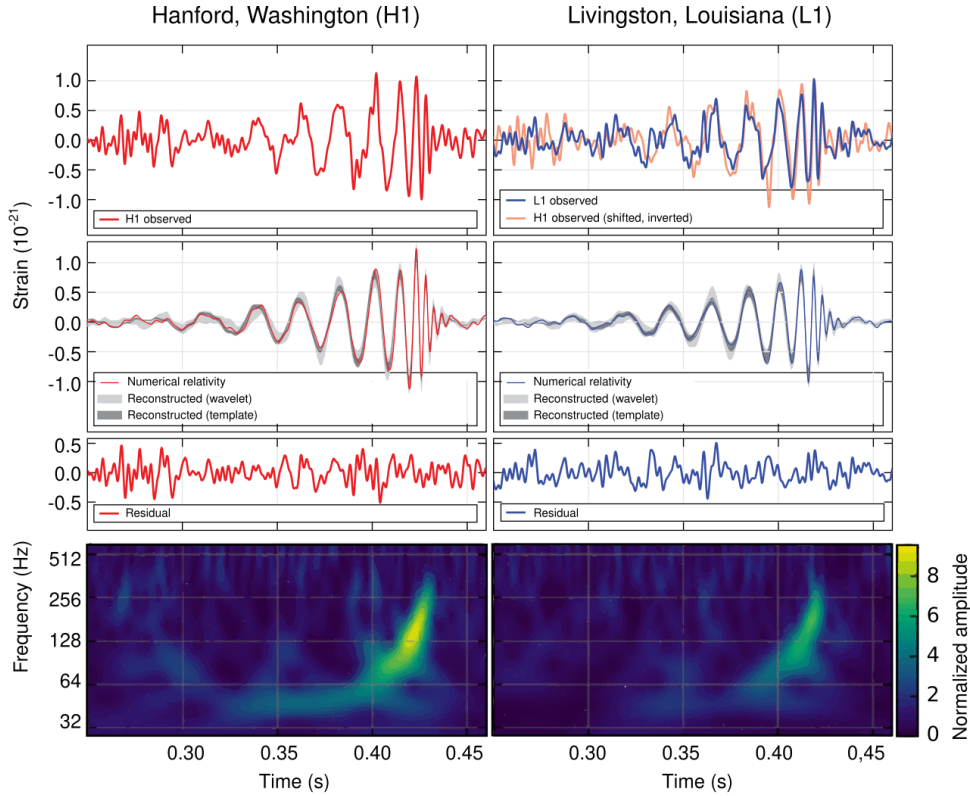


Figure 1.3: GW150914 event as observed by LIGO Hanford (left) and LIGO Livingston (right). From top to bottom, are the strain signal, waveforms, residual noise, and a time-frequency representation of the strain data [10].

The strain plot from the event is shown in Figure 1.3. It was recorded during Advanced LIGO’s first observing run with a consistent waveform observed in both instruments and an inter site arrival-time delay of about 7 ms. In the detectors’ most sensitive band, the strain data show the characteristic “chirp”: the instantaneous frequency rises from ~ 35 to $\sim 250\text{Hz}$ over $\approx 0.2\text{s}$ as the amplitude grows toward merger, reaching a peak strain of order 10^{-21} . The event was recovered with high network signal-to-noise ratio (SNR ≈ 24) and an extremely low false-alarm probability, making it a statistically robust detection despite the limited two-detector network and correspondingly coarse sky localization (hundreds of square degrees) [8].

GW170817 provided the first observation of a binary neutron star inspiral, and it had an especially broad impact because it inaugurated gravitational-wave multi-messenger astronomy. In addition to the gravitational-wave detection, the event was associated with electromagnetic emission, enabling joint constraints that are not accessible with gravitational waves alone. The long inspiral in band also allowed informative constraints on matter effects through tidal signatures, linking the observed waveform to the neutron-star equation of state [5].




Figure 1.4: Time–frequency representations of data containing the binary-neutron-star inspiral GW170817, observed by the LIGO–Virgo network: LIGO Hanford (top), LIGO Livingston (middle), and Virgo (bottom). Time is shown relative to the coalescence time (17 August 2017, 12:41:04 UTC). The colour scale shows the amplitude in each detector, normalised by that detector’s noise amplitude spectral density, revealing the characteristic “chirp” track that sweeps upward in frequency toward merger. In the LIGO data, independently observable noise sources and the short instrumental transient in Livingston were subtracted as part of the analysis. [5]

GW190814 revealed a compact binary with an extreme mass ratio and a secondary component in the $\sim\text{few-}M_{\odot}$ range, a regime that is particularly interesting because it overlaps the putative neutron-star/black-hole “mass gap”. The high signal-to-noise ratio and strong parameter estimation from the detector network made this system an important case study for population inference and for understanding the maximum neutron-star mass and the astrophysical formation channels of highly asymmetric compact binaries [12].

Finally, GW200105 and GW200115 provided the first confident detections of neutron-star–black-hole coalescences. These events completed the observational triad of CBC source classes (BBH, BNS, and NSBH), and they are key for constraining compact-object populations and merger rates, as well as for understanding when (and whether) NSBH mergers produce observable electromagnetic counterparts [16].

1.3 Gravitational-wave detectors

Early experimental efforts to detect gravitational waves focused on resonant-mass antennas called Weber bars, in which a passing wave would excite a high- Q vibrational mode of a massive cylinder. While these early searches did not yield a confirmed detection, they established the basic experimental challenge: measuring strains far below those accessible to conventional metrology. The field shifted toward laser interferometry in the 1970s, motivated by the prospect of broadband sensitivity and scalable baselines. This development ultimately led to kilometre-scale Michelson interferometers: the LIGO facilities in the United States, the 3 km Virgo detector in Europe, and technology testbeds such as GEO 600. The upgrade to second-generation instruments (Advanced LIGO and Advanced Virgo) delivered the sensitivity required for the first direct detection of gravitational waves (GW150914, observed on 14 September 2015), marking the beginning of gravitational-wave astronomy. The global network has since expanded to include

underground, cryogenic interferometers such as KAGRA, and is expected to evolve further with planned third-generation ground-based observatories and future space-based detectors targeting lower-frequency bands [127, 180, 20, 181] [177, 7, 104, 95] [110, 84, 67].

Current ground-based detectors use precision laser interferometry that infer the gravitational-wave strain from a differential change in the cavity lengths of the two perpendicular arms of a Michelson Interferometer. A 1064 nm wavelength laser is split at a beamsplitter and sent down the arms, where highly reflective suspended test masses act as mirrors; in kilometre-scale observatories the arms incorporate resonant optical cavities to increase the phase response to small displacements. The interferometer is operated near a dark fringe so that a gravitational wave produces a measurable change in the output light power (or phase) at the antisymmetric port. The fundamental observable is the dimensionless strain,

$$h(t) = 2 \frac{\Delta L}{L}, \quad (1.5)$$

where L is the arm length and ΔL is the gravitational-wave-induced differential arm-length change. In practice, multiple coupled control loops maintain the optical resonances and mirror alignment, while auxiliary subsystems (vacuum, seismic isolation and suspensions, and the readout chain) suppress environmental disturbances and convert the optical signal into a calibrated $h(t)$ time series [104, 177] [95, 181].

The existing ground-based gravitational-wave detectors are presently in a commissioning-and-upgrade phase between observing campaigns, during which major subsystems are being installed and validated to improve broadband sensitivity. In the LIGO observatories, this transition is captured by the upgrade path from Advanced LIGO to the A+ configuration, with further details deferred to a dedicated section. The A+ upgrade targets an approximately factor-of-two improvement in strain sensitivity relative to the Advanced

LIGO baseline, achieved primarily by suppressing the dominant mid-band thermal noise (via improved low-loss mirror coatings) and reducing quantum noise across a wider band using frequency-dependent squeezing implemented with a dedicated filter cavity [107, 44, 101, 64].

1.4 Noise Sources in GW Detectors

The design sensitivity of existing ground based gravitational wave is limited by shot noise at high frequencies and mirror coating thermal noise at mid frequencies and quantum radiation pressure noise and thermal noise at low frequencies. In figure we see the A+ LIGO design noise budget.

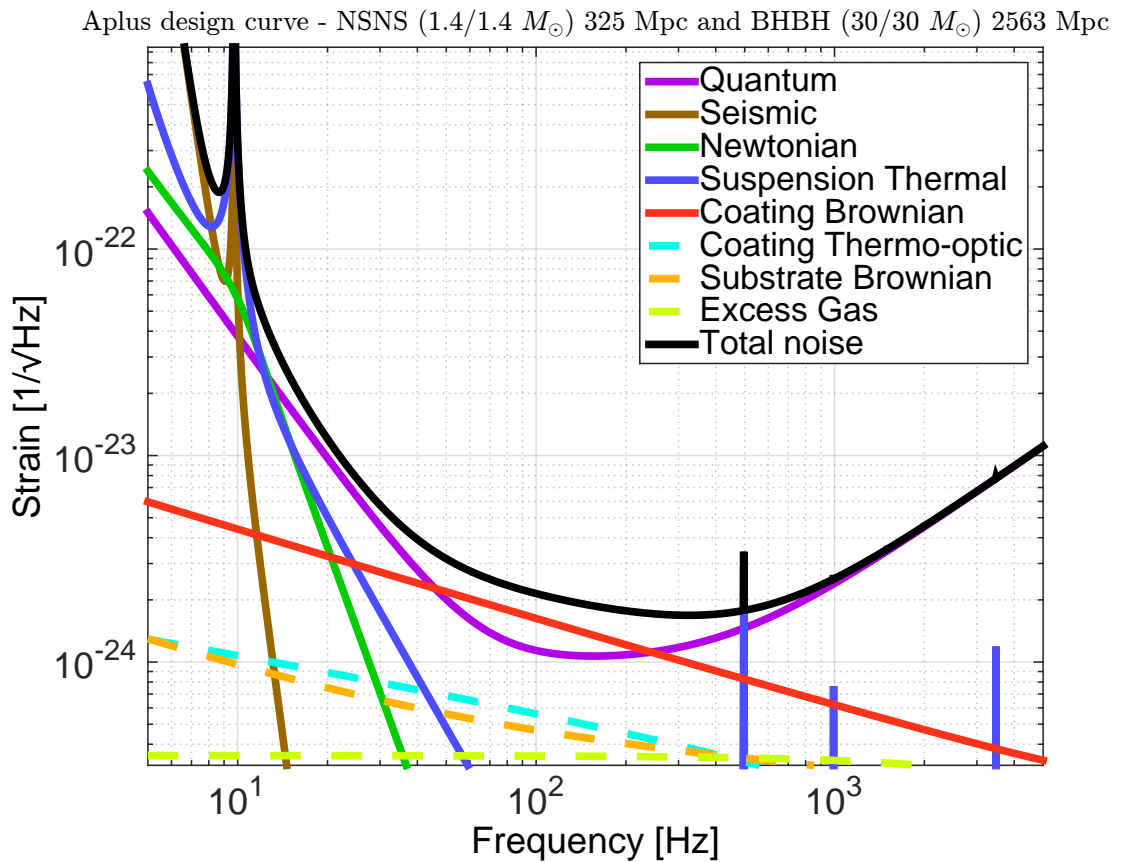


Figure 1.5: A+ Design Curve with noise curves for all the fundamental noise sources in LIGO and Seismic noise. [44].

1.4.1 Thermal noise

Thermal noise in a Michelson-interferometer gravitational-wave detector arises from dissipation (mechanical loss) in multiple subsystems. Through the fluctuation–dissipation theorem, any mechanism that dissipates energy under an applied force produces corresponding thermal fluctuations in the associated displacement readout. In the test masses, the dominant internal contributions are typically substrate Brownian noise, substrate thermoelastic noise, and coating Brownian noise, all driven by microscopic mechanical and thermodynamic fluctuations in the mirror materials [89, 88, 70].

In Advanced LIGO, the leading internal thermal-noise contribution in the most sensitive band ($\sim 50\text{--}200$ Hz) is coating Brownian noise in the high-reflectivity dielectric multilayers deposited on the test-mass surfaces [89, 88, 82]. This dominance reflects the fact that the interferometer readout is directly sensitive to thermally driven surface displacement within the beam spot, and the relevant mechanical loss angle of the coating stack is typically higher than that of the low-loss fused-silica substrate.

In addition to test-mass internal noise, thermal noise is also associated with the suspension system. Suspension thermal noise includes the Brownian motion of the pendulum modes of each stage as well as the violin modes of the suspension fibres [146, 20]. Below $\sim 20\text{--}30$ Hz, suspension thermal noise is often the dominant *thermal* contribution, motivating the use of low-loss fused-silica fibres, multi-stage pendulum geometries, and careful optimisation of fibre geometry and attachment design in Advanced LIGO [1].

Looking ahead, planned upgrades aim to further suppress thermal noise, particularly coating Brownian noise. In the A+ upgrade, the coating-thermal-noise contribution is targeted to be reduced by approximately a factor of two, while more ambitious next-step concepts (often referred to as “A#”) seek order-of-magnitude reductions through advanced coat-

ing materials and related design changes. Third-generation observatories such as the Einstein Telescope pursue even larger thermal-noise improvements by combining new low-loss coatings with cryogenic operation, with the goal of pushing displacement-noise levels well below those planned for second-generation detectors [70, 132, 22].

1.4.2 Quantum noise

Quantum noise in interferometric gravitational-wave detectors arises from vacuum fluctuations of the optical field that couple into the readout at the antisymmetric port. In the conventional picture, it can be separated into (i) *shot noise*, an imprecision noise associated with photon counting statistics that dominates at high frequencies, and (ii) *quantum radiation-pressure noise*, a back-action force from fluctuating photon momentum that drives test-mass motion and dominates at low frequencies.

In the simplest (power-recycled) Michelson picture, the high-frequency quantum-limited strain sensitivity is set by photon shot noise, the corresponding strain amplitude spectral density may be written as

$$h_{\text{SN}}(f) = \frac{1}{L} \sqrt{\frac{\hbar c \lambda}{2\pi P}}, \quad (1.6)$$

where L is the interferometer arm length, λ the laser wavelength, c the speed of light, \hbar the reduced Planck constant, and P the (circulating) optical power relevant for the readout. At low frequencies, for a test mass of mass m , the associated strain amplitude spectral density can be expressed as

$$h_{\text{RP}}(f) = \frac{1}{mf^2 L} \sqrt{\frac{\hbar P}{2\pi^3 c \lambda}}, \quad (1.7)$$

making explicit the characteristic $1/f^2$ rise of radiation-pressure noise and its increase with optical power.

Increasing the circulating power reduces shot noise but enhances radiation-pressure noise, leading to the free-mass *standard quantum limit* (SQL) for a continuous displacement measurement. For a test mass of mass m measured at angular frequency ω , the SQL displacement amplitude is

$$x_{\text{SQL}}(\omega) = \sqrt{\frac{2\hbar}{m\omega^2}}, \quad (1.8)$$

which encapsulates the trade-off between measurement imprecision and quantum back-action [58, 63, 52].

Modern detector configurations exploit quantum correlations to reshape this limit and to reduce quantum noise over a broad band. The most widely deployed technique is the injection of *squeezed vacuum* states, which can lower the phase-quadrature fluctuations and hence reduce shot noise; however, frequency-independent squeezing typically increases radiation-pressure noise at low frequencies. Broadband improvement therefore motivates *frequency-dependent squeezing*, in which the squeezed quadrature is rotated as a function of frequency (e.g. using a low-loss filter cavity), enabling simultaneous reduction of radiation-pressure noise at low frequencies and shot noise at high frequencies. Squeezing has now been implemented in operating gravitational-wave detectors, delivering quantum-enhanced strain sensitivity, and frequency-dependent squeezing has been demonstrated and deployed as part of the pathway toward next-generation performance [43, 170, 114, 74].

1.4.3 Seismic noise

Seismic noise refers to ground motion (from natural and anthropogenic sources) that couples into the interferometer as residual displacement and angular motion of the suspended optics. It is not a fundamental noise source but a technical noise source in LIGO. It is most problematic at low frequencies, where the amplitude of ground motion is large and where maintaining optical resonance and alignment requires significant control forces, which can themselves re-inject motion through sensing and actuation paths. In addition,

ground motion can couple through scattered-light paths, producing excess phase noise that contaminates the gravitational-wave readout. For these reasons, seismic disturbances set key practical limits on low-frequency performance and on robust interferometer operation [1, 112].

Modern detectors mitigate seismic coupling using multiple cascaded stages of isolation: external active platforms and in-vacuum inertial isolation, followed by multi-stage pendulum suspensions that provide strong passive attenuation above their resonances. In Advanced LIGO this strategy combines active inertial control with carefully designed mechanical isolation to achieve many orders of magnitude reduction of ground motion in the detection band [112, 1]. However, at sufficiently low frequencies an irreducible contribution arises from gravity-gradient (Newtonian) noise, where fluctuating gravitational forces from moving ground and atmosphere directly accelerate the test masses and therefore cannot be screened by mechanical isolation; proposed mitigations include site selection (e.g. underground facilities), environmental sensor arrays, and subtraction techniques [86].

Other Technical Noise Sources

Technical noises are non-fundamental noise contributions arising from imperfections in the interferometer, its controls, and its environment, which couple auxiliary degrees of freedom (laser, alignment, seismic, electronics) into the differential arm length readout. In current generation GW detectors, the major contributor to technical noise is at low frequencies and it comes from auxiliary controls such as length and alignment controls [54]. We will not be covering these other noise sources in detail in this thesis.

1.5 Suspension System in LIGO

Ground-based interferometric gravitational-wave detectors rely on multi-stage pendulum suspensions to isolate their test masses from ground motion. Instruments such as GEO 600 and Advanced LIGO suspend the main optics as triple and quadruple pendulums, respectively, combining low-frequency mechanical resonances with careful control of thermal-noise and alignment requirements [128, 139, 99]. There are auxiliary suspensions in these detectors that also hang their optics through a multi-stage pendulum, some double and some triple. Above the pendulum resonances, each stage provides an approximate $1/f^2$ attenuation of transmitted motion, so cascading multiple stages yields the broadband seismic isolation required in the gravitational-wave observation band.

These suspensions are inherently multi-degree-of-freedom systems: each stage supports six rigid-body coordinates (three translations and three rotations, see section 2.3 in the next chapter), and the normal modes of a full chain are generally mixtures of these degrees of freedom. As a result, practical suspension design must account not only for the placement of resonance frequencies, but also for cross-couplings between nominally independent motions that can complicate sensing and control [92, 169]. In addition to providing passive isolation, LIGO suspensions are actively damped using local sensors and actuators so that the interferometer can be acquired and held in lock, which motivates modelling approaches that connect the mechanical plant directly to digital control design [152, 161].

The Advanced LIGO suspensions (see Figure 1.6) are four-stage pendulums whose design builds on the concepts developed for the GEO 600 prototype. The key elements of the suspension system can be grouped into the main pendulum chain, the surrounding support structure, and various auxiliary components like sensors and actuators. The support structure encloses both the main chain and the parallel reaction chain. It protects the suspension from mechanical failure, provides convenient mechanical interfaces, and facil-

itates installation, maintenance and recycling of suspension components. Multiple stages are employed so that horizontal seismic isolation is achieved through the cascaded pendulum stages, while vertical isolation is provided by soft blade springs at several levels in the chain.




Figure 1.6: Drawing of Advanced LIGO quadruple pendulum with a quadruple reaction chain hanging behind it with its surrounding structure [39].

Passive isolation alone is not sufficient to operate a gravitational-wave interferometer at its design sensitivity. Each pendulum stage has resonance modes with high quality factors; if left undamped, these modes can lead to large residual motions and cause the interferometer laser to fall out of phase-lock during operation. Advanced LIGO therefore employs local active control using shadow sensors known as BOSEMs (Birmingham Optical Sensor and Electro-Magnetic actuators) [56][40]. The BOSEMs provide both displacement shadow sensing and coil-magnet actuation: the actuators apply forces to the test-mass chain by pushing against the reaction chain, which serves as a relatively quiet mechanical reference. This arrangement allows the suspension resonances to be actively damped while preserving the high passive isolation at frequencies relevant to gravitational-wave detection.

1.6 Shadow Sensors

Shadow sensors are compact, non-contact displacement sensors widely used for local sensing and damping of suspended optics in ground-based gravitational-wave detectors. In their standard implementation an LED illuminates a photodiode while a “flag” attached to the moving mass partially occludes the beam; small displacements of the flag changes the transmitted optical power and hence the photodiode current, providing an approximately linear displacement signal around a chosen operating point. This is illustrated in detail Figure 1.7. The same module typically integrates a coil–magnet actuator so that the sensor output can be used in local feedback to damp suspension resonances and to provide robust acquisition and alignment support while avoiding the complexity of interferometric readout at low frequencies. The ultimate performance is limited by a combination of fundamental photodetection noise (e.g. shot noise), electronic noise, and practical effects such as LED relative-intensity noise and temperature-dependent responsivity, which become particularly important when extending these sensors to upgraded (A+) requirements and to cryogenic environments for future third generation GW detectors [144, 165, 61].




Figure 1.7: When the flag moves by a distance x , the amount of light falling on the photodiode changes. This is registered as V_{out} .

The remainder of this thesis makes extensive use of suspension modelling and mode analysis in the context of cryogenic interferometry. For readability, the introduction presents only this high-level overview. A detailed treatment of single- and multi-stage pendulum dynamics, state-space formulations, and the modal metrics used throughout this work is deferred to Chapter 2, where we develop the specific models used for the GCIF common-platform suspension and discuss their implications for sensing, damping, and cross-coupling in a cryogenic setting.

1.7 Future upgrades to Advanced LIGO

The second-generation ground-based network has evolved through a sequence of observing campaigns interleaved with commissioning periods, during which targeted hardware and control upgrades are implemented to improve astrophysical reach. In this context, the LIGO detectors progress from the Advanced LIGO configuration toward the A+ upgrade, while longer-term post-O5 concepts (often referred to as A#) and more substantial redesigns (e.g. Voyager) explore how far the existing facilities can be pushed before the transition to third-generation observatories. This section summarises the principal upgrade directions and the corresponding noise-reduction goals.

1.7.1 The A+ upgrade

A+ is an incremental but high-impact upgrade programme aimed at improving broadband strain sensitivity primarily by reducing quantum noise and coating thermal noise for the next observing run O5. A commonly used reference for quantitative comparisons is the A+ design curve model, which modifies the Advanced LIGO noise budget by (i) lowering coating thermal noise (implemented in the model as a reduction of coating mechanical loss) and (ii) applying frequency-dependent squeezing using a dedicated filter cavity to

obtain quantum-noise reduction across a wider frequency band (see Figure 1.5) [44, 114]. In the A+ design-curve parameter set, the frequency-dependent squeezing implementation assumes a ~ 300 m filter cavity and improved optical losses in the squeezing and readout paths, together with increased input power relative to earlier configurations [44].

A major design change for the detector during A+ is the incorporation of Balanced Homodyne Detection Readout scheme (BHD).

1.7.2 LIGO India

LIGO–India is a planned kilometre-scale laser-interferometric gravitational-wave observatory that will add a major new site to the global network of ground-based detectors. Following nearly a decade of planning, the project has moved into the construction phase, with the Indian government authorising full-scale implementation and preparation of the site infrastructure. The observatory is to be built on a remote ~ 174 -acre tract of land at Aundha in Maharashtra’s Hingoli district, selected to support the stringent environmental requirements of precision interferometry. [130]

A central motivation for LIGO–India is improved GW sky localisation which relies on a combination of time-of-arrival triangulation and the differing antenna responses of separated detectors. Adding a LIGO-class instrument in India lengthens the network baselines and improves the distribution of detector orientations, which can substantially reduce typical localisation areas and strengthen multi-messenger follow-up by enabling faster and more targeted electromagnetic observations. Beyond localisation, a better-distributed network helps break parameter-estimation degeneracies (e.g. between luminosity distance, inclination, and polarisation) and increases the overall duty-factor resilience of the network [142].

Technically, LIGO–India is planned as a 4 km Michelson interferometer with Fabry–Pérot arm cavities, closely aligned with the design philosophy of the A+ LIGO instruments, and is being realised through a partnership between Indian funding agencies and institutions and the LIGO Laboratory. The project therefore strengthens both the scientific reach of the international network and the regional capacity for large-scale precision instrumentation, commissioning, and long-term operations in gravitational-wave astronomy [59].

1.7.3 Post-O5 concepts: A#

Beyond A+, the LSC Post-O5 Study Report recommends pursuing a coordinated set of upgrades, collectively referred to as *LIGO A#* (“A-sharp”), as a post-A+ configuration intended to extend the scientific reach of the existing 4 km facilities after O5 and to serve as a technology stepping stone toward third-generation observatories (notably Cosmic Explorer and the high-frequency Einstein Telescope concept). In the baseline A# framing, the detector remains a room-temperature, 1 μm interferometer with fused-silica test masses, but incorporates four major hardware advances: (i) substantially heavier optics (of order $40 \rightarrow 100$ kg) with improved suspensions, (ii) enhanced seismic isolation, (iii) increased circulating arm power (targeting ~ 1.5 MW) together with improved squeezing performance (targeting ~ 10 dB at high frequencies), and (iv) further reduction of coating thermal noise via improved low-loss coatings. The overarching motivation is to deliver broadband sensitivity gains beyond A+ while simultaneously de-risking key technologies and subsystem choices that are relevant for next-generation detector designs [106][42].

A more transformative concept is the LIGO Voyager upgrade, which explores the sensitivity achievable in the existing 4 km facilities using advanced technologies that go beyond the A+/A# room-temperature paradigm. A defining feature of Voyager is the adoption of crystalline silicon test masses operated at cryogenic temperature (commonly discussed around ~ 123 K) together with a longer-wavelength, high-power laser (around the ~ 2 μm band), enabling substantial reductions in thermal noise and changes in optical absorp-


tion/thermal lensing constraints. The Voyager concept is motivated as a route to significantly extend astrophysical reach within the current observatory footprint, while serving as a technology bridge toward third-generation facilities (e.g. through cryogenics, new coatings, and long-wavelength high-power laser development) [22, 102].

1.8 Third-generation gravitational-wave detectors


Third-generation (3G) gravitational-wave observatories aim to extend ground-based sensitivity by roughly an order of magnitude over the advanced-detector era and to improve low-frequency performance down to a few hertz, thereby increasing detection rates and enabling precision tests of strong-field gravity and cosmology with large event populations. Achieving this requires simultaneous reductions in quantum noise, thermal noise (especially coating and suspension thermal noise), and low-frequency disturbances (seismic and Newtonian noise), motivating new facility-scale design choices such as longer baselines, improved vacuum systems, and quieter sites [132, 110, 84].

The *Einstein Telescope* (ET) is a European 3G concept based on an underground, triangular interferometer with 10 km arms, designed to operate as three colocated interferometers with multiple baselines and polarisations for robust signal reconstruction. A key element of the ET strategy is the use of cryogenic operation and dedicated interferometers optimised for different frequency ranges (often discussed as “xylophone” configurations), combining low thermal noise at mid frequencies with improved low-frequency sensitivity enabled by a seismically quiet underground environment [132, 110]. These design features target substantial reductions in coating and suspension thermal noise and mitigate low-frequency environmental limitations relative to surface facilities [110].

In the United States, *Cosmic Explorer* (CE) pursues a complementary approach based on very long-baseline, L-shaped interferometers (tens of kilometres) to increase the gravitational-wave signal response and relax several technical constraints that scale with arm length. By combining longer arms with advanced quantum-noise reduction (including squeezing), improved optics/coatings, and upgraded suspension and isolation technologies, CE aims to deliver broadband sensitivity gains and a dramatic increase in accessible volume [84]. At lower frequencies than ground-based instruments, the *Laser Interferometer Space Antenna* (LISA) will open the milli-hertz band using a constellation of spacecraft in a heliocentric orbit to form million-kilometre-scale laser interferometers. This enables observations of massive black-hole binaries, extreme-mass-ratio inspirals, and Galactic compact binaries, providing a low-frequency counterpart to ground-based 3G detectors and enabling multi-band gravitational-wave astronomy for suitable sources [25, 67].



(a) An artist's rendering of ET showing 10 km arms triangular configuration.



(b) An artist's rendering of CE showing 40 km. (Credit: Eddie Anaya, California State University Fullerton)

Figure 1.8: Artist's impression of the two third-generation GW detectors.




Figure 1.9: Amplitude spectral density of strain noise, $\sqrt{S_h(f)}$, for representative current and proposed ground-based gravitational-wave detectors. Shown are measured sensitivities from the third observing run (e.g. LIGO and Virgo), together with design or target sensitivity curves for Advanced LIGO, Advanced Virgo, KAGRA, LIGO A+, LIGO Voyager, Cosmic Explorer (40 km) and the Einstein Telescope (ET-D). The curves illustrate the characteristic low-frequency limitation from seismic and Newtonian noise, the mid-band thermal-noise floor, and the rise at high frequency dominated by quantum shot noise [103] [68, 91] [83, 133].

Figure 1.9 compares the strain sensitivity of the current detector network with projected upgrade and third-generation concepts across the \sim Hz–kHz band. The familiar “bucket” shape reflects the transition from low-frequency limitations (seismic and Newtonian/gravity-gradient noise), through a mid-band floor dominated by thermal noise, to a high-frequency rise set primarily by quantum shot noise. The upgrade trajectories (e.g. A+ and Voyager) show incremental improvements within the existing facility scale, whereas third-generation concepts (Cosmic Explorer and Einstein Telescope) target order-of-magnitude sensitivity gains and significantly improved performance below \sim 10 Hz [68, 91]. Cosmic Explorer sensitivities are often presented as a *family* of curves: CE1 and CE2 denote staged implementations, with CE2 sometimes split into “silica” and “silicon” variants to represent different technology paths (room-temperature fused-silica versus cryogenic silicon/longer-wavelength operation) [68, 135, 83].

1.9 Detector prototypes and technology demonstrators

A defining feature of gravitational-wave instrumentation is that many of the techniques required for future observing runs cannot be fully de-risked on the kilometre-scale detectors without unacceptable commissioning risk. Instead, dedicated prototypes and technology demonstrators provide controlled environments in which new interferometer configurations, sensing and control schemes, high-power optical effects, and quantum-noise reduction methods can be developed to maturity before deployment. These facilities also serve as platforms for end-to-end integration tests, allowing subsystem interactions (e.g. optomechanical dynamics, scattered-light couplings, and control noise reinjection) to be studied in regimes that resemble full-scale operation.

Historically, prototype interferometers have played a key role in bridging conceptual designs to operational observatories. The Caltech 40 m prototype has been used as a system-level testbed for advanced readout and control developments, including early demonstrations of DC readout concepts relevant to the transition from RF to DC/balanced-homodyne readout architectures [178]. Likewise, the University of Glasgow 10 m prototype provided an early experimental platform for suspended Fabry–Pérot Michelson operation, sensitivity characterisation, and suspension/laser developments that informed later generations of detectors [137, 154]. More recently, purpose-built facilities such as the AEI 10 m prototype in Hannover have been designed explicitly to explore quantum-limited interferometry and sub-SQL measurement techniques, combining ultra-high-vacuum infrastructure with advanced seismic isolation concepts and suspended optical systems representative of next-generation requirements [80].

Alongside full interferometer prototypes, specialised test facilities address limiting noise sources and scaling challenges expected for future upgrades. High-optical-power effects—such as thermal lensing, angular/parametric instabilities, and power-dependent scattered-light couplings—are difficult to study safely on observatories, motivating dedicated high-power facilities such as the Gingin High Optical Power Test Facility [184]. In parallel, prototypes have been central to demonstrating quantum-noise reduction methods: squeezed-light injection was first shown to yield a clear sensitivity improvement on a suspended-mirror prototype interferometer, providing an early validation of techniques now routinely used on operating detectors [77]. Audio-band frequency-dependent squeezing, realised using an external filter cavity, subsequently established the practicality of broadband quantum-noise reduction in a detector-relevant band and informed the implementation pathway toward advanced-detector upgrades [121].

For space-based detection, where in-flight validation is essential, *LISA Pathfinder* served as the principal technology demonstrator for the key measurement problem: placing freely falling test masses in near-perfect geodesic motion and measuring their relative displacement with exquisite precision. The mission demonstrated sub-femto-*g* differential free-fall and achieved performance at (and later beyond) the required acceleration-noise levels for LISA in relevant frequency bands, substantially reducing risk for the full observatory and providing a mature experimental foundation for milli-hertz gravitational-wave detection [37, 36].

1.9.1 Glasgow Cryogenic Interferometer Facility (GCIF)

Interferometric gravitational-wave research in Glasgow dates back to the establishment of Ronald Drever’s group in 1970 and the first operation of an interferometer prototype in the 1980s. In the early 2000s the prototype facility was relocated to a clean-room laboratory and rebuilt with an upgraded layout comprising two 10 m arms and a parallel 5 m arm, supported by a vacuum system of nine 1 m-diameter tanks and a modern control and data-acquisition system comparable in architecture to that used in Advanced LIGO.

This upgraded Glasgow 10 m facility operated from 2003 to 2024 and hosted a broad programme of experiments, including speed-meter interferometry, studies of optical springs to mitigate radiation-pressure effects, development of diffractive optics for GW detectors, direct measurements of coating thermal noise at room temperature, and investigations of sensing and control schemes for dual-recycled Michelson interferometers [81][79, 108, 93, 160, 182].




Figure 1.10: GCIF Cryostat in the JIF lab (Kelvin Building, University of Glasgow). A section of the 10 m reference cavity beam tube is seen behind the cryostat in the photograph.

The Glasgow 10 m prototype underwent a major refurbishment in 2024–25 to incorporate a cryostat housing a short (~ 10 cm) optical cavity suspended using crystalline material. Following this upgrade, the facility is now referred to as the Glasgow Cryogenic Interferometer Facility (GCIF). The ~ 1 kg test masses for the 10 cm cavity are suspended using a triple-stage chain that includes two stages of vertical blade springs, providing a platform to demonstrate and characterise multi-material crystalline suspension concepts at cryogenic temperature. A primary scientific objective is to measure coating thermal noise in the 10 cm cavity, enabling comparative studies of candidate low-loss mirror coatings relevant to third-generation cryogenic detectors. In addition, GCIF is designed to operate at both 1550 nm and 2000 nm, supporting investigations of long-wavelength laser architectures for future observatories. More broadly, the facility provides an opportunity to develop and validate sensing and control of silicon suspensions under cryogenic conditions and to study practical issues such as ice formation on cold silicon optics [81].

1.10 Overview of thesis

In this chapter we introduced the basic concepts required for the rest of the thesis: the astrophysical sources of gravitational waves, the operating principles of kilometre-scale interferometric detectors, and the main noise sources that limit their sensitivity. We also reviewed the role of multi-stage pendulum suspensions and local sensing schemes in current detectors, with an emphasis on how these ideas motivate the developments presented in later chapters.

Chapter 2 focuses on the design of a triple-stage *common platform suspension* for the Glasgow Cryogenic Interferometer Facility (GCIF). Using a Mathematica state-space model, we investigate the dynamics of a split-pendulum configuration in which two double-pendulum chains are suspended from a shared top stage. The chapter explores the resulting forest of normal modes, the implications for length noise, local control, and common-mode rejection, and outlines design choices that make the suspension mechanically and cryogenically feasible for a multi-material silicon–sapphire suspension system.

Chapter 3 describes the development of a cryogenic-compatible shadow sensor intended for local control of the GCIF suspensions. We characterise a range of infrared LEDs, quadrant photodiodes, and flag geometries down to cryogenic temperatures, and build a detailed noise and performance model of the sensor. The resulting prototype meets the combined requirements on sensitivity, dynamic range, and heat load.

Chapter 4 reports on the design, fabrication, and characterisation of the A+ LIGO Output Mode Cleaner (OMC) cavities. The chapter details the optical and mechanical design of the monolithic bow-tie cavities, the screening and metrology of coated optics, the assembly and bonding procedures, and the techniques used to optimise cavity throughput performance. The measured performance of the first set of A+ OMCs demonstrates transmission in excess of 98%, meeting the requirements for squeezed-light operation in the upgraded detectors.

Chapter 5 develops and tests a polarisation-based balanced homodyne detection scheme that uses a single OMC to filter both the interferometer signal and the local oscillator. We present the theoretical framework for this configuration, analyse the impact of cavity birefringence and waveplate settings on noise coupling, and describe a table-top experiment that validates the concept and identifies an operating point with suppressed coupling of OMC length noise to the homodyne readout.

Finally, the appendices collect supporting material that would otherwise interrupt the flow of the main text, including analytical details of special suspension cases and extended mode tables for the common-platform suspension models.

Glasgow Cryogenic Interferometer Facility: Common Platform Suspension Design

2.1 Introduction

The Glasgow Cryogenic Interferometer Facility (GCIF) is designed to directly measure coating thermal noise at cryogenic temperatures. This will be achieved by observing the displacement thermal noise of a Fabry–Pérot cavity with a round-trip length of 20cm, formed by two suspended silicon test masses of mass 1kg each. A further objective of the GCIF is to demonstrate and characterise a fully crystalline cryogenic suspension system. The facility will operate in two phases: an initial phase at 123K and a later phase at 18K. Both temperature regimes coincide with a zero crossing of silicon’s coefficient of thermal expansion and are therefore of particular interest as potential operating temperatures. This happens for sapphire at 123K making it a likely candidate too. Operating at these temperatures reduces the overall thermal noise of the experimental system, enabling a direct measurement of coating thermal noise, and provides a testbed for the development and characterisation of materials and technologies relevant for third-generation gravitational-wave detectors.

Ground-based gravitational-wave detectors employ multistage pendulum suspensions to provide passive seismic isolation for the cavity mirrors. Increasing the number of suspension stages improves the level of isolation, with each pendulum stage contributing an attenuation that scales as $1/f^2$ above its resonance frequency. Current-generation detectors, such as Advanced LIGO, typically use four stages of pendulum isolation for the test masses. For the Glasgow prototype, we intend to implement a three-stage suspension system. The suspension design must simultaneously satisfy the stringent requirements on both thermal noise and seismic isolation in order to achieve the sensitivity necessary for coating thermal noise measurements.

In this chapter, we present the suspension modelling of common-platform triple-stage suspensions in Mathematica, including the development of the model, the resulting performance studies, and possible future directions [47]. We also worked on implementation of a Python notebook designed to enable faster computation of the suspension model, motivated by the fact that evaluating the state-space matrices becomes computationally expensive for the common-platform suspension configuration. A comparison of double- and triple-stage pendulum suspensions was done in the past [81] and concluded that a triple-stage pendulum configuration was preferable. In this chapter, I will compare the triple-stage suspension model with the common-platform triple-stage model.

2.2 Pendulum Dynamics: State Space

Multi-stage pendulum suspensions are the standard solution for providing seismic isolation of test masses in ground-based interferometric gravitational-wave detectors such as GEO 600 and Advanced LIGO [99][113]. In these instruments the main optics are suspended as double, triple, or quadruple pendulums in order to achieve sufficient attenuation of ground motion in the observation band, while simultaneously satisfying stringent requirements on thermal noise and controllability [128, 139]. In GEO 600, for example, the main mirrors are suspended as triple pendulums incorporating fused-silica fibres and cantilever

springs to provide horizontal and vertical isolation, and the Advanced LIGO test masses extend this concept to a quadruple suspension [128, 139]. Increasing the number of suspension stages improves the level of isolation, with each pendulum stage contributing an attenuation that scales as $1/f^2$ above its resonance frequency.

The dynamic behaviour of these suspensions is inherently multi-degree-of-freedom: each stage has six rigid-body coordinates (three translations and three rotations, refer to figure 2.1), and the normal modes of the full system are in general mixtures of these degrees of freedom. A detailed six-degree-of-freedom model of a triple pendulum were already developed for GEO 600 to support suspension design, optimisation of attachment geometry, and prediction of resonance frequencies and transfer functions [92, 128, 169]. In this approach the equations of motion are first derived in terms of generalised coordinates, usually starting from a Lagrangian formulation and linearising about the equilibrium configuration, and then written in the standard second-order matrix form

$$\mathbf{M}\ddot{\mathbf{q}} + \mathbf{C}\dot{\mathbf{q}} + \mathbf{K}\mathbf{q} = \mathbf{F}, \quad (2.1)$$

where \mathbf{q} collects the translational and rotational coordinates of the suspended mass, \mathbf{M} is the mass/inertia matrix, \mathbf{C} the damping matrix, \mathbf{K} the stiffness matrix and \mathbf{F} the vector of generalised forces and torques.

For LIGO suspensions it is not sufficient to model only the passive isolation, the suspensions are actively controlled using local sensors and actuators to damp suspension modes so that the interferometer can be kept locked. This has led to use of state-space models, in which the second-order equations are recast as a set of coupled first-order differential equations,

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{B}(t)\mathbf{u}(t), \quad (2.2)$$

$$\mathbf{y}(t) = \mathbf{C}(t)\mathbf{x}(t) + \mathbf{D}(t)\mathbf{u}(t). \quad (2.3)$$

Equation 2.2 is the state equation and equation 2.3 is the output equation.

The state vector \mathbf{x} that typically contains the coordinates \mathbf{q} and their velocities $\dot{\mathbf{q}}$, \mathbf{y} is the output vector, \mathbf{u} is the input or control vector. \mathbf{A} , \mathbf{B} , \mathbf{C} and \mathbf{D} are state space matrices. The state variables we select must sufficiently describe the behaviour of the system. In the context of multi-stage pendulum suspension modelling, the inputs are linear or angular displacements of the structure and forces or torques on the pendulum coordinates. The outputs are the pendulum coordinates and back reaction forces on the structure. The state coordinates are the pendulum coordinates.

$$\mathbf{A} = \begin{pmatrix} 0 & I \\ \Re(-\mathbf{M}^{-1}\mathbf{K}_{\text{eff}}) & \frac{\Im(-\mathbf{M}^{-1}\mathbf{K}_{\text{eff}})}{2\pi f} \end{pmatrix}, \quad (2.4)$$

$$\mathbf{B} = \begin{pmatrix} 0 & 0 \\ \Re(-\mathbf{M}^{-1}\mathbf{C}_{XS(\text{eff})}) & \mathbf{M}^{-1} \end{pmatrix}, \quad (2.5)$$

$$\mathbf{C} = \begin{pmatrix} I & 0 \\ \Re(-\mathbf{C}_{SX(\text{eff})}) & 0 \end{pmatrix}, \quad (2.6)$$

$$\mathbf{D} = \begin{pmatrix} 0 & 0 \\ -\mathbf{S}_{\text{eff}} & 0 \end{pmatrix}. \quad (2.7)$$

Here \mathbf{M} is the mass matrix, \mathbf{K} is the stiffness matrix for the pendulum with the structure held fixed, $\mathbf{C}_{SX(\text{eff})}$ is the coupling matrix from the structure to the pendulum (it's a symmetric matrix), and \mathbf{S}_{eff} is the structure stiffness matrix with the pendulum held fixed at its nominal position.

This formalism is particularly convenient for designing digital control systems (for example using modal control and state estimation), and for integrating the mechanical plant, sensors and actuators into a single consistent model. State-space techniques for suspension control have been developed and applied in the context of Advanced LIGO quadruple suspensions and related prototypes, including modal damping and state-estimation-based controllers [152, 161]. Within the LIGO collaboration, Mark Barton has developed a suite of numerical tools that generate state-space models for a variety of suspension configurations (single, double, triple, quadruple stages and split pendulum systems) from a common parameter set describing the geometry and material properties [45]. In this framework, the Mathematica suspension models first determine the resonance modes by solving the undamped generalised eigenvalue problem of the linearised equations of motion. The resulting dynamics are then recast into a state-space representation, which is used for time-domain simulations and exported to Matlab, where the models are incorporated into Simulink for further numerical and control-system studies.

These models provide the \mathbf{M} , \mathbf{C} , and \mathbf{K} matrices, and hence the corresponding state-space matrices \mathbf{A} , \mathbf{B} , \mathbf{C} , and \mathbf{D} , and have been extensively used for controller design and noise studies. More recently, such state-space suspension models have been used in the analysis and reduction of control noise in current detectors, where the full control loop (suspension dynamics, filters, and sensing) is simulated in a unified framework [179].

2.3 Dynamics of a Single Stage Pendulum

Every finite-mass mechanical system exhibits resonance modes. In the case of a mass suspended by a wire, the system forms a pendulum whose normal modes of vibration are referred to as *pendulum modes*. Owing to the six mechanical degrees of freedom, these modes can be classified into six types: longitudinal, bounce, side, roll, yaw, and pitch. These are illustrated in Fig. 2.1.




Figure 2.1: Illustration of the six mechanical degrees of freedom of a single-stage pendulum: (a) longitudinal translation, (b) side translation, (c) vertical (bounce) motion, and rotations about the three orthogonal axes corresponding to (d) pitch, (e) roll, and (f) yaw.

A single-stage pendulum can be viewed as the simplest building block of the full multi-stage suspension. Writing down its equations of motion in matrix form, and then converting them to state-space form, provides a transparent way to understand how the suspension geometry and material parameters enter the mass, damping, and stiffness matrices, and how couplings between degrees of freedom (for example between longitudinal and pitch, or sideways and roll for a four-wire suspension) appear as off-diagonal terms.

A detailed analysis of a single stage pendulum hung using four wires can be found in Calum Torrie's thesis [169]. The equations of motion for the pendulum are given in equations 2.8 and 2.9.




Figure 2.2: Schematic of a rigid body suspended from a stationary platform by four wires. The parameters n_1 and n_2 denote the lateral distances of the wire attachment points from the z -axis along the centre-of-mass (COM) line of the body. The quantity d is the vertical offset of the attachment plane from the COM, l is the wire length, and s is the axial separation of the front and back attachment points.

$$\frac{d^2}{dt^2} \begin{bmatrix} x \\ y \\ z \\ \theta_{\text{pitch}} \\ \theta_{\text{roll}} \\ \theta_{\text{yaw}} \end{bmatrix} = \mathbf{A} \begin{bmatrix} x \\ y \\ z \\ \theta_{\text{pitch}} \\ \theta_{\text{roll}} \\ \theta_{\text{yaw}} \end{bmatrix}, \quad (2.8)$$

$$\mathbf{A} = \begin{bmatrix} k_{xx} & k_{xy} & k_{xz} & k_{x\theta_p} & k_{x\theta_r} & k_{x\theta_y} \\ k_{xy} & k_{yy} & k_{yz} & k_{y\theta_p} & k_{y\theta_r} & k_{y\theta_y} \\ k_{xz} & k_{yz} & k_{zz} & k_{z\theta_p} & k_{z\theta_r} & k_{z\theta_y} \\ k_{x\theta_p} & k_{y\theta_p} & k_{z\theta_p} & k_{\theta_p\theta_p} & k_{\theta_p\theta_r} & k_{\theta_p\theta_y} \\ k_{x\theta_r} & k_{y\theta_r} & k_{z\theta_r} & k_{\theta_p\theta_r} & k_{\theta_r\theta_r} & k_{\theta_r\theta_y} \\ k_{x\theta_y} & k_{y\theta_y} & k_{z\theta_y} & k_{\theta_p\theta_y} & k_{\theta_r\theta_y} & k_{\theta_y\theta_y} \end{bmatrix}. \quad (2.9)$$

This is a generalised stiffness matrix with all the couplings.

2.3.1 Modes of a Single Stage Pendulum

The resonance modes for a single stage pendulum can be categorised as longitudinal, transverse, bounce, pitch, yaw and roll according to the dominant component of the motion. For the common case of an optic with principal moments of inertia aligned with the coordinate axes, the vertical and yaw modes are typically fairly pure, but the longitudinal and pitch modes are usually somewhat blended, as are the transverse and roll modes, as explained in the next section.

Longitudinal-Pitch (rotation about y) Modes

This mode is represented by the combination of motions illustrated in figures 2.1 (a) and (d). The coupling between longitudinal motion and pitch arises from the fact that the centre of mass is displaced by a finite distance d from the suspension point. In the limit $d \rightarrow 0$ the off-diagonal stiffness terms $k_{\theta_p x}$ and $k_{x\theta_p}$ in equations (2.12) and (2.13) vanish, and the longitudinal and pitch degrees of freedom become dynamically decoupled. However, for a realistic suspension d is non-zero, so that a longitudinal displacement of the mass produces a torque about the pitch axis, and conversely a pitch rotation produces a net longitudinal restoring force. This is explicitly seen in Eqs. (2.12) and (2.13), where both coupling coefficients scale linearly with d , $k_{\theta_p x} \propto d$ and $k_{x\theta_p} \propto d$, so that even a modest geometric offset leads to measurable mixing of the two modes.

The finite offset d also modifies the pure pitch stiffness, as reflected in the linear and quadratic d -dependence of $k_{\theta_p\theta_p}$ in Eq. (2.11). Physically, a rotation about the pitch axis shifts the centre of mass longitudinally as well as vertically when $d \neq 0$, changing the gravitational restoring torque compared to the idealised case where the COM lies exactly in the suspension plane. Together, these effects mean that the normal modes of the system are not purely longitudinal or purely pitch, but instead mixed modes whose degree of coupling is controlled by the suspension geometry through the parameter d .

$$\frac{d^2}{dt^2} \begin{bmatrix} \theta_{\text{pitch}} \\ x \end{bmatrix} = \mathbf{A}_{x\theta_p} \begin{bmatrix} \theta_{\text{pitch}} \\ x \end{bmatrix}, \quad \mathbf{A}_{x\theta_p} \equiv -\mathbf{M}_{x\theta_p}^{-1} \mathbf{K}_{x\theta_p}^{(\text{stiff})}. \quad (2.10)$$

where elements of the stiffness matrix are given by

$$k_{\theta_p\theta_p} = -\frac{mgs^2}{I_y h l^2} (n_2 - n_1)^2 - \frac{mgd}{I_y} - \frac{mgd^2}{I_y h} - \frac{2ks^2 h^2}{I_y l^2}, \quad (2.11)$$

$$k_{\theta_p x} = \frac{mgd}{I_y h}, \quad (2.12)$$

$$k_{x\theta_p} = \frac{gd}{h}, \quad (2.13)$$

$$k_{xx} = -\frac{g}{h}. \quad (2.14)$$

$$f_{\text{pitch/longitudinal}} = \frac{1}{2\pi} \sqrt{\text{abs}(\text{eig}(\mathbf{A}_{x\theta_p}))} \quad (2.15)$$

where `abs` returns the absolute value of the eigenvalues `eig`($\mathbf{A}_{x\theta_p}$) of the matrix $\mathbf{A}_{x\theta_p}$.

In the limits where $n_2 = n_1$ and $d = 0$ i.e. the wires are straight and the attachment points are in the centre of mass plane, the stiffness matrix elements reduce to

$$k_{\theta_p \theta_p} = -\frac{2ks^2}{I_y}, \quad (2.16)$$

$$k_{\theta_p x} = 0, \quad (2.17)$$

$$k_{x \theta_p} = 0, \quad (2.18)$$

$$k_{xx} = -\frac{g}{l}. \quad (2.19)$$

These are very familiar forms we are used to seeing. The eigenfrequencies then become

$$f_{\text{pitch}} = \frac{1}{2\pi} \sqrt{\frac{2ks^2}{I_y}}, \quad (2.20)$$

$$f_{\text{longitudinal}} = \frac{1}{2\pi} \sqrt{\frac{g}{l}}, \quad (2.21)$$

Transverse-Roll (rotation about x) Modes

The equations of motion for this mode which is a combination of translation along y and roll i.e. rotation about x axis (see figures 2.1 (b) and (e)) are given by:

$$\frac{d^2}{dt^2} \begin{bmatrix} \theta_{\text{roll}} \\ y \end{bmatrix} = \mathbf{A}_{y\theta_r} \begin{bmatrix} \theta_{\text{roll}} \\ y \end{bmatrix}, \quad \mathbf{A}_{y\theta_r} = \begin{bmatrix} k_{\theta_r \theta_r} & k_{\theta_r y} \\ k_{y \theta_r} & k_{yy} \end{bmatrix}, \quad (2.22)$$

where

$$k_{\theta_r\theta_r} = \frac{mgn_2(n_2 - n_1)}{hI_x} - \frac{mgd}{I_x} - \frac{mgd^2h}{l^2I_x} - \frac{mgn_2^2(n_2 - n_1)^2}{l^2hI_x} - \frac{2kd^2(n_2 - n_1)^2}{l^2I_x} + \frac{2mgdn_2(n_2 - n_1)}{l^2I_x} - \frac{4kdn_2h(n_2 - n_1)}{l^2I_x} - \frac{2kn_2^2h^2}{l^2I_x}, \quad (2.23)$$

$$k_{\theta_y} = \frac{mgdh}{l^2I_x} - \frac{mgn_2(n_2 - n_1)}{l^2I_x} + \frac{2kd(n_2 - n_1)^2}{l^2I_x} + \frac{2kn_2h(n_2 - n_1)}{l^2I_x}, \quad (2.24)$$

$$k_{y\theta_r} = -\frac{gdh}{l^2} + \frac{gn_2(n_2 - n_1)}{l^2} - \frac{2kd(n_2 - n_1)^2}{ml^2} - \frac{2kn_2h(n_2 - n_1)}{ml^2}, \quad (2.25)$$

$$k_{yy} = -\frac{gh}{l^2} - \frac{2k(n_2 - n_1)^2}{ml^2}. \quad (2.26)$$

$$f_{\text{transverse/roll}} = \frac{1}{2\pi} \sqrt{\text{abs}(\text{eig}(\mathbf{A}_{y\theta_r}))} \quad (2.27)$$

where `abs` returns the absolute value of the eigenvalues `eig`($\mathbf{A}_{y\theta_r}$) of the matrix $\mathbf{A}_{y\theta_r}$.

The transverse–roll coupling has a similar geometric origin as the longitudinal–pitch coupling discussed above: it arises because the suspension wires do not act through the centre of mass and are laterally angled with respect to the z axis. In the transverse case the two fibres are attached at positions displaced by $(n_2 - n_1)$ in the horizontal direction and by a finite vertical offset d from the centre of mass. A pure transverse displacement y therefore changes the lengths, and hence the tensions, of the left and right fibres by different amounts. This imbalance in fibre tension generates a net torque about the roll axis, which appears as the off–diagonal stiffness terms k_{θ_y} and $k_{y\theta_r}$ in Eq. (2.22).

Conversely, a pure roll rotation θ_{roll} tilts the mass so that the fibre attachment points move transversely as well as vertically whenever the lever arms d and $n_2 - n_1$ are non-zero. The resulting change in the horizontal components of the fibre tensions produces a net transverse restoring force. This is reflected in the explicit dependence of the coupling

coefficients on d , n_2 and $(n_2 - n_1)$ in Eqs. (2.24) and (2.25). In the symmetric limit where the left and right fibres are straight ($n_2 - n_1 \rightarrow 0$) and the attachment offsets vanish ($d \rightarrow 0$), these off-diagonal terms go to zero and the transverse and roll degrees of freedom decouple, yielding normal modes that are purely translational or purely rotational:

$$k_{\theta_r \theta_r} = -\frac{2kn_2^2}{I_x}, \quad (2.28)$$

$$k_{\theta_r y} = 0, \quad (2.29)$$

$$k_{y \theta_r} = 0, \quad (2.30)$$

$$k_{yy} = -\frac{g}{l}. \quad (2.31)$$

and the eigenfrequencies become

$$f_{\text{roll}} = \frac{1}{2\pi} \sqrt{|k_{\theta_r \theta_r}|} = \frac{1}{2\pi} \sqrt{\frac{2kn_2^2}{I_x}}, \quad (2.32)$$

$$f_{\text{transverse}} = \frac{1}{2\pi} \sqrt{|k_{yy}|} = \frac{1}{2\pi} \sqrt{\frac{g}{l}}. \quad (2.33)$$

Bounce mode (vertical translation)

For a pendulum mass supported by four inclined fibres at angle Ω (so $\cos \Omega = h/l$), a small vertical displacement z changes each fibre length by

$$\Delta l = z \cos \Omega + \frac{mg}{2k} \left(\frac{z \tan^2 \Omega}{l} \right), \quad (2.34)$$

where k is the axial spring constant of a single fibre. To first order, the vertical equation of motion is

$$m\ddot{z} = -4k\Delta l \cos\Omega \approx -\left(4k\cos^2\Omega + \frac{mg\sin^2\Omega}{l\cos\Omega}\right)z. \quad (2.35)$$

Thus the bounce resonance is a simple harmonic mode with frequency

$$f_{\text{bounce}} = \frac{1}{2\pi} \sqrt{\frac{4k\cos^2\Omega}{m} + \frac{g\sin^2\Omega}{l\cos\Omega}}. \quad (2.36)$$

The first term is the *elastic vertical stiffness* of the fibres projected along z ; the second is a *geometric gravitational stiffness* that appears because the fibres are inclined ($\sin\Omega \neq 0$) so a vertical motion changes the vertical projection of the suspending forces.

In the limit $\Omega \rightarrow 0$, Eq. (2.36) reduces to $f_{\text{bounce}} = \frac{1}{2\pi} \sqrt{4k/m}$, which corresponds to the standard mass–spring resonance for four vertical springs in parallel.

Yaw mode (rotation about z)

With two fibre pairs attached at transverse/lateral offset s and longitudinal offsets n_1, n_2 , a small yaw rotation σ produces the total restoring torque

$$I_z \ddot{\sigma} = -\frac{mg}{h} \left(s^2 \left(\frac{h}{l} \right)^2 + n_1 n_2 \right) \sigma - \frac{4ks^2(n_2 - n_1)^2}{l^2} \sigma, \quad (2.37)$$

which yields the yaw resonance

$$f_{\text{yaw}} = \frac{1}{2\pi} \sqrt{\frac{mg}{I_z h} \left(s^2 \left(\frac{h}{l} \right)^2 + n_1 n_2 \right) + \frac{4ks^2(n_2 - n_1)^2}{I_z l^2}}. \quad (2.38)$$

The first (gravity) term arises from the *lever-arm geometry*: when the suspension points are offset from the COM, a yaw rotation lifts one side and lowers the other, creating a gravitational restoring torque. The second term is a *differential elastic (torsional) stiffness*: if the two fibre pairs are not longitudinally symmetric ($n_2 \neq n_1$), a yaw rotation differentially stretches the left/right fibres and adds an elastic restoring torque proportional to their axial stiffness k .

For $n_2 = n_1$, the elastic-asymmetry term vanishes and Eq. (2.38) becomes

$$f_{\text{yaw}} = \frac{1}{2\pi} \sqrt{\frac{mg}{I_z l} (s^2 + n_1^2)}. \quad (2.39)$$

Takeaway

It is useful to pause here and compare the angular eigenmodes. The pitch and roll eigenfrequencies, Eqs. (2.20) and (2.32), take closely analogous forms: in both cases the restoring torque is set by the elastic response of the suspension fibres, and the mode frequency scales linearly with the relevant wire separation (i.e. the lever arm in the direction of the tilt). This reflects the fact that pitch and roll motion involves differential stretching and compression of the fibres, so the effective stiffness enters explicitly in the square-root dependence.

In contrast, the yaw eigenfrequency, Eq. (2.39), does not arise from fibre extension but from torsional motion, and consequently its restoring term is dominated by gravity rather than the fibre spring constant. The yaw frequency also depends on the wire separations in both transverse directions through the geometric factors appearing in Eq. (2.39).

It is generally desirable to minimise coupling between nominally independent suspension modes, since such couplings complicate both the interpretation of transfer functions and the design of robust control filters. In particular, mixing between longitudinal and pitch, and between transverse and roll, can lead to situations in which a controller designed to act on one degree of freedom inadvertently excites motion in another. This is already a concern for room-temperature multi-stage suspensions such as those used in GEO 600 and Advanced LIGO, where careful optimisation of the suspension geometry and tuning of local damping loops are required to keep modal cross-couplings within acceptable limits [128, 139, 161].

In cryogenic suspensions these issues become even more pronounced. Cooling introduces additional asymmetries through differential thermal contraction of structural elements, changes in wire or fibre tension, and possible non-uniform temperature gradients along the suspension chain. All of these effects can perturb the nominal geometry and material properties, effectively shifting the elements of the mass and stiffness matrices \mathbf{M} and \mathbf{K} , and thereby enhancing off-diagonal terms that mix the different degrees of freedom. Experience from cryogenic prototype systems and from the KAGRA cryogenic interferometer indicates that such thermo-mechanical drifts can lead to slow changes in resonance frequencies, increased modal coupling and, in the worst case, degraded controllability or excess control noise [26, 155].

Besides the cross-couplings between independent DOFs studied above (longitudinal-pitch and transverse-roll), like in any real system, small imperfections such as unequal wire lengths, slight differences in wire diameter or tension, and machining tolerances on attachment points further break the ideal symmetry of the suspension, introducing additional off-diagonal elements in both \mathbf{K} and \mathbf{M} that mix longitudinal/pitch and transverse/roll motion with yaw and bounce [92, 47]. For a cryogenic system, these imperfections are designed to be small at room temperature, but the cooling cycle can amplify them through differential contraction. It is therefore essential to provide sufficient local sensing and actuation authority to monitor and compensate for these drifts over the lifetime of the experiment. Shadow sensors, optical levers and interferometric position sensors, combined

with appropriately designed state-space controllers, offer a route to actively suppress unwanted couplings and maintain a well-behaved set of modes [152, 179]. In the next chapter, we describe the development of a cryogenic-compatible shadow sensor intended specifically for this role in the GCIF suspension system.

Transfer Function

The transfer function from longitudinal motion of the suspension point to longitudinal motion of the test mass for a representative single-stage pendulum hung using four vertical wires with the attachment points 10 mm above the COM plane is shown in Figure 2.3.




Figure 2.3: Transfer functions for a representative single-stage pendulum. The orange trace shows the magnitude of the transfer function from longitudinal motion of the support to longitudinal motion of the suspended mass, $|H_x|$. Below the pendulum resonance the mass follows the support motion with approximately unity gain, while near the resonance the response is amplified before rolling off with an asymptotic $1/f^2$ dependence at higher frequencies, providing seismic isolation in the observation band. The green trace shows the corresponding cross-coupled transfer function from support longitudinal motion to optic pitch, $|H_{x \rightarrow \text{pitch}}|$. The blue trace displays the ratio $|H_{x \rightarrow \text{pitch}}|/|H_x|$, which quantifies the longitudinal-to-pitch coupling: it peaks near the coupled resonance and tends to a constant at high frequency, reflecting the fixed geometrical ratio between translational and rotational response in the inertial regime.

2.4 Triple Stage Pendulum

In contrast to a single stage pendulum, a triple-stage pendulum comprises three such suspended masses (or stages), each with the same six degrees of freedom. Consequently, a triple-stage pendulum exhibits a total of 18 normal modes (six per stage), corresponding to the full set of coupled longitudinal, bounce, side, roll, yaw, and pitch motions of the three-stage system. It is crucial to design the suspensions such that the resonance mode frequencies are as low as possible, since noise above the pendulum resonance frequency is attenuated in that degree of freedom. In a triple-stage suspension each stage has its own set of pendulum modes, giving a total of 18 modes. Each successive mode provides an additional $1/f^2$ of isolation, so it is important to place these modes at sufficiently low frequencies in order to maximise isolation for each degree of freedom. Among these, the longitudinal motion is the most critical, as it couples directly to the cavity error signal. In practice, there is usually cross-coupling between different resonance modes. This cross-coupling depends on the number of fibres used and on geometric parameters such as the fibre attachment points, as depicted in Figure 2.6. Even for an ideal, perfectly constructed suspension there is an inherent coupling of a small fraction of vertical motion into the longitudinal direction. For example, in LIGO the 4-km arm length combined with the curvature of the Earth requires the suspensions at each end station to be tilted slightly with respect to the local surface. This geometric effect leads to a coupling of vertical motion into longitudinal motion at the level of $\sim 0.1\%$ [139, 128]. For the GCIF cavity, this geometric coupling is expected to be significantly smaller than in the LIGO case [81].

Adding additional stages to the pendulum offers further isolation. In figure 2.4 we see how a two stage pendulum and three stage pendulum offer further isolation.

The multi-stage pendulums used in current detectors can be generalised to a *split pendulum* configuration, in which two independent lower chains are suspended from a single common top mass. A schematic of such a configuration is shown in Figure 2.5. Each lower chain can carry its own optic (for example a silicon and a sapphire test mass), while




Figure 2.4: Longitudinal transfer functions from support motion to test-mass motion, $|H_x|$, for single-, double-, and triple-stage pendulums. At low frequency all three systems follow the support motion with approximately unity gain, while around the pendulum resonances the response is amplified. Above the highest resonance the attenuation improves with each additional stage, approaching asymptotic slopes of $1/f^2$, $1/f^4$, and $1/f^6$ for the single-, double-, and triple-stage configurations respectively, illustrating the enhanced seismic isolation provided by multiple cascaded stages.

the shared top stage acts as a mechanical reference that supports both. In the context of the GCIF common-platform suspension, this opens up the possibility of exploiting the common-mode rejection between the two chains to reduce the effective coupling of platform motion into the cavity length, while still allowing independent control of the individual optics. In the following chapter we develop a state-space model of such a split pendulum and investigate how the choice of geometry and mass distribution determines the normal modes, the level of common-mode rejection, and the residual cross-coupling relevant for interferometric sensing and control.




Figure 2.5: Split pendulum configuration with common top mass.

2.4.1 Some Definitions

Cross Coupling

The mode-shape vectors reported in this work are Euclidean normalised [45]. For a mode eigenvector \mathbf{q} formed by stacking all translational and rotational DOF components of all suspended bodies, the normalisation condition is

$$\|\mathbf{q}\|_2 \equiv \sqrt{\sum_i q_i^2} = 1. \quad (2.40)$$

To quantify modal cross-coupling from the eigenvector listings, let's consider a six DOF eigenvector vector for a single stage pendulum for a given mode,

$$\mathbf{q} \equiv (q_x, q_y, q_z, q_{yaw}, q_{pitch}, q_{roll})^T, \quad (2.41)$$

where the components are taken directly from the eigenvector table output in Mathematica [45][47]. We define the *dominant DOF* as the index of the largest-magnitude component,

$$j_{\text{dom}} \equiv \arg \max_j |q_j|. \quad (2.42)$$

The cross-coupling from the dominant DOF into another DOF j is then defined as the relative amplitude

$$C_j \equiv \frac{|q_j|}{|q_{j_{\text{dom}}}|}, \quad j \neq j_{\text{dom}}, \quad (2.43)$$

and the maximum cross-coupling for the mode is

$$C_{\text{max}} \equiv \max_{j \neq j_{\text{dom}}} C_j. \quad (2.44)$$

Mode Visibility

The resonance spectrum of a multi-stage suspension contains many closely spaced modes distributed across the six rigid-body degrees of freedom (DoFs) of each stage. For suspension and control design it is therefore not sufficient to list eigenfrequencies alone: one must also quantify how strongly each mode appears at a location of interest (e.g. the common platform) relative to the optic, and in which DoF the mode is primarily expressed. To do this, we compute a *mode visibility* metric that compares the modal amplitude of a given stage to the modal amplitude of the optic.

For each eigenmode k , we obtain the corresponding eigenvector $\mathbf{q}^{(k)}$ from the linearised equations of motion. The eigenvector is partitioned into the six DoFs of each rigid body. We identify the *dominant DoF* of mode k at the optic as the coordinate with the largest magnitude component,

$$j_{\text{dom}}^{(k)} \equiv \arg \max_j \left| q_{\text{opt},j}^{(k)} \right|, \quad (2.45)$$

where $q_{\text{opt},j}^{(k)}$ denotes the optic component of the eigenvector in DoF $j \in \{x, y, z, \text{yaw}, \text{pitch}, \text{roll}\}$.

The mode visibility of the top (common platform) mass for mode k is then defined as the ratio of the top-mass amplitude to the optic amplitude in this dominant DoF,

$$Q_{\text{top}}^{(k)} \equiv \frac{\left| q_{\text{top}, j_{\text{dom}}}^{(k)} \right|}{\left| q_{\text{opt}, j_{\text{dom}}}^{(k)} \right|}. \quad (2.46)$$

By construction, $V_{\text{top}}^{(k)} \ll 1$ indicates a mode that is largely localised to the optic (poorly visible at the top mass) and would need to be controlled at the middle mass, whereas $Q_{\text{top}}^{(k)} \sim 1$ indicates a mode with comparable participation at the top mass and the optic, which is generally more accessible to local sensing and damping at the top mass as desirable. The same definition is used to compute visibilities for other stages (e.g. intermediate masses) by replacing the numerator in Eq. (2.46) with the corresponding stage component.

This visibility metric provides an intuitive measure of how strongly each mode is expressed at potential local-control locations relative to the optic motion that is most directly relevant to cavity length. In addition, the eigenvector components can exhibit cross-coupling between DoFs; therefore, alongside $Q^{(k)}$ we also record the dominant DoF and the largest cross-coupled DoF using the coupling metric defined in equation 2.43. The resulting mode-visibility tables are used throughout this thesis to compare suspension geometries, identify modes that are likely to be controllable from the common platform, and highlight configurations in which cryogenic asymmetries may introduce additional modal mixing.

2.5 Split Triple Stage Pendulum: Common Platform Suspensions

The common-platform suspension is modelled as a split pendulum in which two suspension chains share a common upper mass. The primary motivation for this architecture is the prospect of common-mode rejection between the two chains that form the cavity, since a significant fraction of the support motion is transmitted through the shared top stage. By mechanically correlating the excitation of the two chains, the common top mass can reduce the differential motion that directly appears as cavity length noise, while retaining the multi-stage isolation of a cascaded pendulum.

To assess the performance of this configuration, we compare several figures of merit, including the resonance-mode visibility at relevant stages, the longitudinal isolation inferred from the suspension transfer functions, and mode cross couplings all while keeping mechanical feasibility in mind.

2.5.1 Common-Platform vs. Conventional Triple Suspensions

In a conventional three-stage *single-chain* suspension [81], the resonance structure is close to separable: each stage contributes six canonical rigid-body resonances (longitudinal x , transverse y , bounce z , yaw, pitch, and roll), giving a tidy set of $3 \times 6 = 18$ modes whose dominant motion is largely confined to one DOF at one stage. Weak geometric couplings do exist—notably longitudinal–pitch and transverse–roll—but they are small enough that modes can be classed cleanly as longitudinal, transverse, bounce, yaw, pitch, and roll; observed degeneracies mostly reflect the symmetric participation of the input and end chains (I/E) rather than true mixing. Sensing and control for separate triple-chain

suspensions are therefore straightforward: modes of the lower stages that are sufficiently visible at the top mass are sensed and damped by the shadow sensor and coil–magnet actuators on the top mass, while modes that are not visible there can be sensed and damped by the sensors and actuators on the middle mass.

By contrast, the *common–platform* suspension is based on a split–pendulum design and introduces systematic, symmetry–breaking couplings that reorganise the resonances into clusters with mixed DOF content. The shared top mass ties transverse motion of the bottom stages to platform rotation, producing measurable transverse–roll modes with occasional platform–yaw leakage even when the optic motion is nominally transverse. Likewise, the split geometry makes top–stage pitch tilt the chains and drive off–phase vertical (bounce) motion in the lower stages, so the “bounce family” contains a mixed z +pitch resonance rather than a perfectly isolated pitch motion of the top mass. As a result, several modes that would be distinct in a single chain collapse into near–degenerate groups, and the I/E ordering often swaps within a cluster.

These differences have practical consequences:


1. Sensing and actuation must be designed with the mixed families in mind, since platform readouts can be driven by non–platform DOFs and some modes are not visible at the top mass but can be controlled at the intermediate stage;
2. Deliberate design choices (e.g. introducing small asymmetries, increasing top–stage yaw/pitch moments of inertia, or detuning vertical stiffnesses) may be required to reduce cross–couplings and separate troublesome coincidences.

Besides the suspension parameters indicated in the diagram in Figure 2.6 there are additional parameters that can be explored to further understand the dynamics of the suspension, these parameters are summarised in Tables 6, 7 and 8 in Appendix C.


2.5.2 Design Requirements and Considerations

When the circulating laser field in a resonating optical cavity interacts with coatings on the surface of the cavity mirrors, Brownian thermal noise in the coatings produces fluctuations in the optical phase. These appear as apparent cavity length fluctuations and are registered as noise in the cavity error signal. For the GCIF cavity the target coating thermal noise performance is a displacement sensitivity of $1 \times 10^{-17} \text{ m}/\sqrt{\text{Hz}}$ at 100 Hz at 123 K operation, and $4 \times 10^{-18} \text{ m}/\sqrt{\text{Hz}}$ at 100 Hz at 18 K operation, with an approximately $1/\sqrt{f}$ dependence above 100 Hz [81]. The suspension system must therefore be designed such that its contribution to the cavity length noise budget remains comfortably below these targets above 100 Hz, which is the frequency band of interest. The common platform suspension introduces an additional set of design constraints beyond those of a single chain: the common platform must provide sufficient isolation for both optical chains, while keeping cross-coupling between the two as small as possible.

Besides the length noise requirements discussed in the previous paragraph, we want to understand the requirements in the modelling phase of the design of the suspensions. When designing the suspensions it is important to understand the parameter space and what we want to optimise for those parameters. These parameters are highlighted in the schematic in figure 2.6. In particular, the normal modes that couple strongly to the cavity length must lie at low enough frequencies, with adequate attenuation above the science band (i.e. observation frequency band), and the differential motion between the two chains must be controllable without exciting large common-mode motion of the platform.



(a) Top view of the common platform showing suspension point geometry for the common platform. The six blade spring locations used to hang the common platform from the cryostat platform is shown.



(b) Triple-stage common platform suspension model (front and side views).

Figure 2.6: Triple-stage common platform suspension model. The input and end chains are identical double pendulums suspended from a rigid common platform or upper mass. The upper mass is hung from the top plate, also referred to as cryostat platform, using 6 metal wires and 6 blade springs as shown in (a). The upper blade tips lie on the top plate on a circle of radius 392.3 mm, and they one metal wire hanging from them each. The left two chains show these head on and the right set shows side view. The three masses m_1 , m_2 and m_3 represent the upper, intermediate and optic masses of the cryogenic suspension with succeeding letter “I” or “E” indicating Input or End chain respectively. The parameters l_1 – l_3 denote the pendulum wire lengths, d_0 – d_4 give the vertical separations between wire attachment points and COM planes for the respective masses; n_2 – n_5 indicate horizontal separation between wire attachment points when viewed head on. The s_i and s_l parameters indicate the wire separation along the cavity axis. We keep the wires vertical in the model to minimise cross couplings.

2.5.2.1 Length Noise and Mode Placement

From the point of view of the cavity error signal, the relevant requirement is that the total cavity noise in the longitudinal degree of freedom is below the coating thermal noise target across the observation band. For Advanced LIGO, the quadruple suspensions were designed such that the suspension displacement noise at each test mass is of order $10^{-19} \text{ m}/\sqrt{\text{Hz}}$ at 10 Hz, well below the overall design noise curve [139, 1]. A similar approach is adopted for the common-platform suspension model in Mathematica/Python. The suspension mechanical parameters and blade-spring parameters are tuned to obtain a design that meets the longitudinal noise requirements. Specifically, the target is cavity length isolation below $4 \times 10^{-18} \text{ m}/\sqrt{\text{Hz}}$ at 100Hz at 18K [81].

The low-frequency resonance modes which couple directly to longitudinal motion should lie low so that the passive $1/f^n$ attenuation of a pendulum reduces seismic motion to an acceptable level in the observation band [138, 72].

We also require the visibility of the resonance modes to be large enough so that the local sensors on the top mass and the intermediate mass are able to detect and damp them.

2.5.2.2 Local Control and Sensing Requirements

In addition to passive isolation, the suspensions must support the active control system that keeps the cavity on resonance. In GEO 600, Advanced LIGO and KAGRA the interferometer lengths must be held within a small fraction of the laser wavelength to maintain lock and achieve design sensitivity [139, 72, 152]. The RMS motion of the cavity after damping should be smaller than the laser linewidth for the PDH locking to keep the interferometer under lock. We also do not want large swinging of the optics as this makes it hard for the control loop keep the interferometer under lock.

2.5.2.3 Mechanical and Cryogenic Feasibility

As the suspensions are cooled, a number of mechanical parameters change by significant amounts. Structural elements and blades contract, changing wire lengths and blade pre-curvedness (work in progress); and the Young's moduli and loss angles of the materials change with temperature. In room-temperature fused-silica suspensions, such as those used for the Advanced LIGO test masses, the vertical compliance of the silica fibres is large enough that dedicated blade springs are not required at the lowest stage [1]. In contrast, sapphire and silicon fibres are significantly stiffer, so vertical isolation must be provided by carefully designed blade springs at the top stages and other isolation techniques [24, 98].

For the common-platform design this leads to several mechanical feasibility constraints:

- The stresses in all fibres and blade springs must remain below a limit that depends on the strength of the fibres at the lowest operating temperature, with a comfortable safety factor against fracture [176, 24]. This is also an on-going work in the Institute for Gravitational Research, Glasgow [168][120].
- Vertical and horizontal resonance frequencies involving the platform and chains must be kept within ranges that are compatible with both isolation and controllability (e.g. avoiding very high vertical modes that up-convert seismic noise, and avoiding very low modes that are difficult to damp).
- Differential contraction between different materials (e.g. blades, maraging steel, silicon or sapphire fibres, bonding interfaces) must not over-stress joints or introduce excessive misalignment during cool-down [24, 98].
- The suspension must provide sufficient thermal conductance to cool down to 18K and to further extract the absorbed laser power from the test mass while maintaining the required operating temperature, which strongly constrains the fibre cross-section and thus the achievable pendulum thermal noise [24, 70].

The silicon masses and the top stage should also be designed such that the parametric internal modes involving the masses do not cause instabilities in the control system [69]. The masses also need to provide suitable attachment points for the sapphire or silicon fibres [62]. The design parameters indicated in the Figure 2.6 should also be practical such that the variations in the geometries of the fibres or attachments due to thermal drift and machining tolerances do not lead to unwanted asymmetries in the suspensions 2.10.

2.5.3 Resonance Mode Forest

As introduced earlier, Advanced LIGO-style suspensions comprise a *main* chain and a *reaction* chain. The reaction chain is a pendulum suspension that acts as a recoil frame for the main suspension chain and provides inertial reference for the displacement sensors. In our Mathematica framework we consider two reaction-chain idealisations: (i) a *stiff* reaction structure, and (ii) a *pendulum* reaction chain mirroring the main chains. The current GCIF conceptual design omits blade springs at the lowest stage—fabricating silicon blade springs at that location is presently impractical—so the bottom stage is modelled without blades.

This section surveys the resonance modes and their *mode visibility* (i.e. which degrees of freedom are observable at the common platform and, when platform motion cancels, at the intermediate mass). Two platform configurations to isolate how the equations of motion and inter-stage couplings change with recoil implementation can be considered:

1. a common platform supporting **two** double-pendulum chains (main only), with the recoil realised as a rigid structural cage [38, 174, 119].
2. a common platform supporting **four** double-pendulum chains (main + reaction for two optics) [39][139];

In this chapter we consider only the first model as this is the nominal plan for the GCIF cryogenic suspensions design.

Accurate modal modelling is essential for commissioning; when a resonance rings up, we must identify the dominant degree of freedom and any cross-DOF leakage. Minimising such cross-couplings simplifies sensing and damping design, reduces controller spillover between channels, and ultimately improves robustness of the local control loops.

2.6 Model A - Cage like Reaction Mass

The suspension dynamics were modelled in Mathematica using the Advanced LIGO suspension modelling toolkit developed by Barton [45]. In this work I used the *Common Platform* model (`CommonPlatformTriple_6B6W_2x4B4W4B4W`), for which the toolkit formulates the Lagrangian, derives the linearised equations of motion, and solves the resulting eigenvalue problem to obtain the normal-mode eigenfrequencies and eigenvectors.

To enable rapid exploration of the numerical parameter space, the symbolic state-space matrices for Model A were exported directly as Python code, avoiding repeated symbolic evaluation inside Mathematica. The corresponding state-space export assumes a fixed parameter vector; the complete list of parameters used for this model is given in Table 2.1.

Table 2.1: Key mechanical parameters used in the suspension model of Sec. 2.6. The listed d -values are *effective* attachment offsets: the corresponding physical attachment distances are obtained by subtracting the relevant flexure length from these effective values.

Symbol	Description	Units	TOP	IM	TST
m	Mass	kg	8.483	0.996937	0.996937
I_{xx}	Roll moment of inertia (about x)	kg m^2	0.537952	1.17252×10^{-3}	1.17252×10^{-3}
I_{yy}	Pitch moment of inertia (about y)	kg m^2	0.538075	8.64773×10^{-4}	8.64773×10^{-4}
I_{zz}	Yaw moment of inertia (about z)	kg m^2	0.12227	8.64773×10^{-4}	8.64773×10^{-4}
I_{xy}	Product of inertia (x - y)	kg m^2	3.41396×10^{-5}	0	0
I_{yz}	Product of inertia (y - z)	kg m^2	0	0	0
I_{zx}	Product of inertia (z - x)	kg m^2	0	0	0
$d_{\text{top}}/d_1/d_3$	(upper) wire vertical attachment distance	mm	0.05	1	1
$d_0/d_2/d_4$	(lower) wire vertical attachment distance	mm	1	1	1
$n_0/n_2/n_4$	(upper) wire transverse attachment half-spacing	mm	30	45	45.5
$n_1/n_3/n_5$	(lower) wire transverse attachment half-spacing	mm	40	45.5	45.5
$s_u/s_i/s_l$	front-back wire attachment half-spacing	mm	10	22.5	10
$l_1/l_2/l_3$	Stretched wire length	m	0.298167	0.298167	0.298167
k_w	Net longitudinal (axial) wire stiffness	N/m	1396.07	1396.07	186218
k_{bz}	Net vertical blade stiffness	N/m	2490	201.51	N/A
Y	Young's modulus (wire material)	Pa	2.12×10^{11}	2.12×10^{11}	1.67326×10^{11}
$M_{\text{dir}1}$	Second moment of area (direction 1)	m^4	3.06796×10^{-19}	3.06796×10^{-19}	8.76241×10^{-15}

This model comprises two double pendulum chains suspended from a common platform. The top stage is intended to be fabricated from aluminium, while the lower two stages are silicon. The optic stage is planned to be suspended using silicon/sapphire fibres from the intermediate mass, with the top and middle stages supported by metal wires. The masses belonging to the bottom two stages are made of silicon. This makes the bottom two stages in the Figure 2.6b (quasi-)monolithic. The sensors and actuators for the lower masses are mounted on a safety/actuation/sensing cage extending down rigidly from the top mass, rather than on independently suspended recoil masses.

The common platform mass and overall height budget are constrained by the cryostat geometry, following an initial mechanical design by Russell Jones. The reaction cage is assumed to have negligible mass and moment of inertia (MOI) in the dynamical model.

Two levels of blade springs are included: one at the top stage and one at the upper mass stage. The silicon test masses have a mass of 1 kg, limited primarily by cost considerations, and are designed with diameter approximately equal to thickness in order to maximise the internal mode frequencies.

The suspension fibres are sapphire with a diameter of 650 μm , chosen to satisfy thermal load constraints. The fibres are attached 1 mm above and below the centres of mass (COMs) of the intermediate mass (IM) and the optic, respectively, resulting in small effective d parameters and therefore reduced pitch–longitudinal cross-coupling. The model exhibits 30 normal modes (Table 2.2). In the following, we identify these modes, highlight any significant cross-couplings, and evaluate their mode visibility (Figure 2.7). Tolerances of 1 mm are assumed for the d parameters; comparable tolerances are also assumed for the s and n parameters. The impact of suspension asymmetry is deferred to future sections.




Figure 2.7: Mode visibility for the common–platform triple suspension. The model corresponds to a configuration in which the recoil masses are realised as a cage structure integrated with the top mass.




Figure 2.8: Split Pendulum diagram indicating the notation used for the degrees of freedom for each stage used in Table 2.2

Table 2.2: Suspension mode list for the two-chain (I/E) common–platform suspension model with reaction cage. Refer to the diagram in Figure 2.8 for guidance on the DOF notations. The mode images can be found in Appendix D

#	Frequency (Hz)	Mode (dominant DOFs)
1	0.641210	x3E, x3I, x2E
2	0.641280	y3I, y3E
3	0.723078	y3E, y3I, y2E
4	0.728643	x3I, x3E
5	1.085460	x3I, x3E, pitch1, x1, pitch2E
6	1.086016	y3I, y3E, roll1, roll2I, roll2E, roll3I
7	1.190957	yaw3I, yaw3E
8	1.204083	yaw3E, yaw3I
9	1.677232	yaw3I, yaw3E
10	1.761336	pitch3E, pitch3I, pitch2E, pitch2I
11	1.763599	y2I, y2E, roll3I, roll3E, roll2I
12	1.769449	x2I, x2E
13	1.847908	yaw3I, yaw3E, yaw1
14	2.381061	pitch3E, pitch3I, pitch2E
15	2.406700	pitch3I, pitch3E, pitch2I
16	2.500904	z3I, z3E, z2I
17	2.620410	pitch3I, pitch3E, pitch2I
18	2.658471	roll3I, roll3E, roll2I, roll2E
19	2.907044	yaw2I, yaw2E
20	2.916330	yaw2E, yaw2I
21	3.093916	pitch3I, pitch3E, pitch2I, pitch2E
22	4.019618	roll3I, roll3E, roll2I
23	4.048402	roll3E, roll3I, roll2E
24	4.787872	z1, z3E, z3I, z2E
25	66.101819	pitch2E, pitch2I, pitch3E
26	66.101820	pitch2I, pitch2E, pitch3I
27	194.577838	z2E, z2I, z3E
28	194.577840	z2I, z2E, z3I
29	258.154297	roll2I, roll2E, roll3I
30	258.154297	roll2E, roll2I, roll3E

The calculated spectrum naturally organises into frequency clusters that reflect the split-pendulum couplings described above. The 30 resonance modes comprise 11 longitudinal-pitch modes, 8 transverse-roll modes, 7 yaw, and 4 bounce modes. This in contrast to the 6,6,3,3 structure observed for a traditional triple stage suspension. Since the split pendulum has 5 masses, we expect 10 longitudinal-transverse modes, 10 transverse-roll, and 5 of yaw and bounce.

Cross-couplings outside the usual families.

Several modes exhibit secondary components outside the standard longitudinal-pitch and transverse-roll pairings:

- **Transverse-yaw mixing (Mode 3).** Mode 3 ($f \simeq 0.723\text{Hz}$) shows appreciable simultaneous y and yaw content at the lower stages/optic, indicating a transverse-yaw family rather than a pure transverse-roll mode. Mode 13 also exhibits this behaviour although the mixing is lower.
- **Bounce leakage into longitudinal-pitch modes (e.g. Mode 5).** Mode 5 ($f \simeq 1.085\text{Hz}$) is longitudinal-pitch dominated but has a noticeable bounce component at the optic. This corresponds to coupling of $\approx 5.7\%$. This is substantially larger bounce leakage than Mode 1. Other modes that exhibit this are Mode 10, 17 and 21

It is important to understand the downsides of this model before we explore the parameter space to improve the model. Mode 3 has significant transverse-yaw coupling, the eigenvectors for these modes are shown in Table 2.3 and the mode shape in Figure 2.10. The plot in Figure 2.9 shows the relevant transfer function summarising the behaviour. Other cross couplings that naturally occur, either due to the geometry of the suspensions: longitudinal to vertical, or the finite effective d values: longitudinal-pitch coupling, are plotted in Figure 2.11.

Table 2.3: Mode eigenvector components for the translational (x, y, z) and rotational (yaw, pitch, roll) DOFs for each stage for Mode 3 at 0.723078 Hz.

	x	y	z	yaw	pitch	roll
Mass U	0	0	0	0.146099	0	0
Mass II	0	-0.268244	0	0.192529	0	-0.0374406
Mass IE	0	0.268244	0	0.192529	0	0.0374406
Input Optic	0	-0.564069	0	0.243573	0	-0.0374485
End Optic	0	0.564069	0	0.243573	0	0.0374485

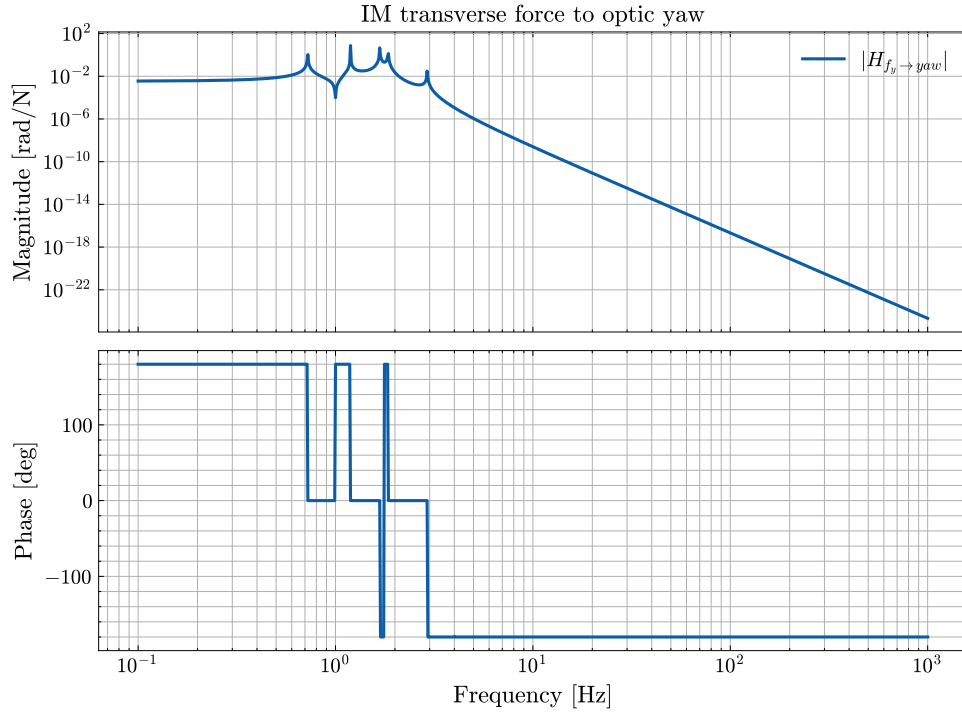


Figure 2.9: Transfer function from a transverse force disturbance applied at the intermediate mass to the resulting yaw displacement of the optic, $H_{F_y \rightarrow yaw}(f)$, shown as magnitude and phase versus frequency.


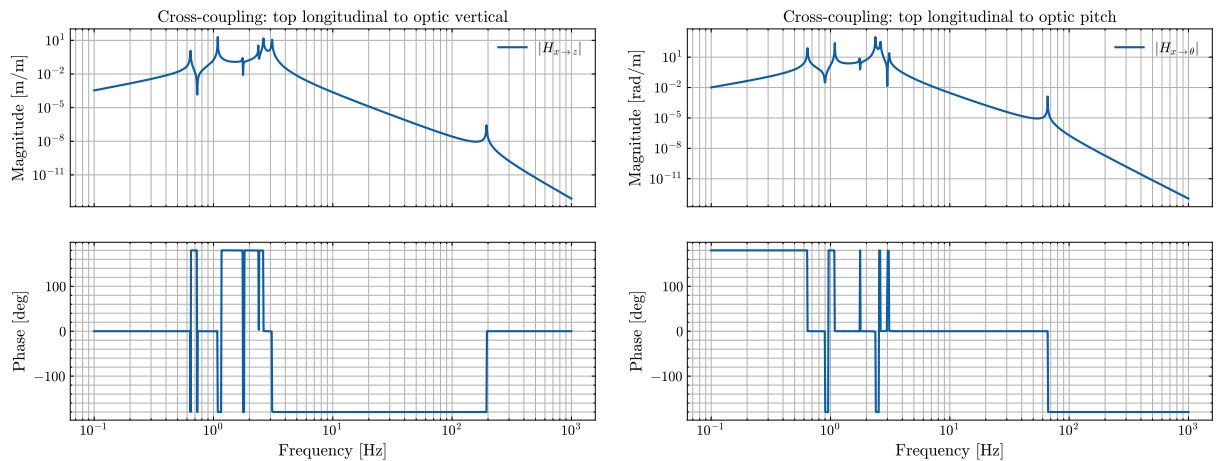


Figure 2.10: Top view of the suspension model in Mathematica with common platform mass shown as semi-transparent for clarity. Resonance mode 3 at 0.723078 Hz. Here we see the dominant optic stage transverse motion coupling to top mass yaw motion that in turn couples to yaw motion of the bottom stages.

From the mode-visibility plot we see that 8 out of 30 modes are not visible enough at the top mass and requires controlling at the intermediate stage. In the next section we will see if moving the COM of the top mass back to the centre improves cross coupling.



(a) Longitudinal \rightarrow vertical transfer function, $H_{x \rightarrow z}(f)$, showing the magnitude and phase of optic vertical motion per unit longitudinal excitation at the suspension point.

(b) Longitudinal \rightarrow pitch transfer function, $H_{x \rightarrow \theta}(f)$, showing the magnitude and phase of optic pitch rotation per unit longitudinal excitation at the suspension point.

Figure 2.11: Cross-coupling transfer functions from longitudinal excitation at the suspension point to (left) optic vertical displacement and (right) optic pitch rotation. Each panel shows magnitude and phase as a function of frequency, evaluated for the common-platform suspension model.

2.7 Moving the COM to the centre of top mass

In the previous model in Figure 2.7, the COM of the common platform is below the geometric centre of the top mass, this is due to the presence of the reaction cage.

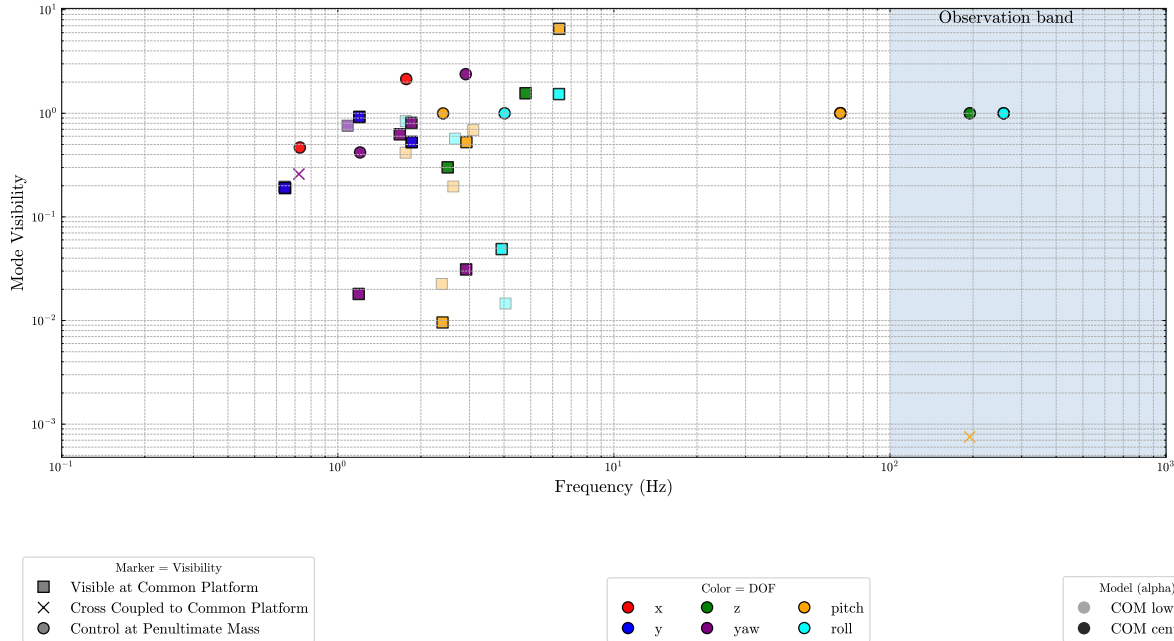


Figure 2.12: Mode visibility for the common-platform triple suspension with COM of the common platform at the centre of the top mass.

The average mode visibility of this configuration is higher than that of the model with the centre of mass (COM) located lower. This improvement arises primarily from the reduced moment of inertia: with a smaller inertia, the top mass is more readily rotated by the suspension forces, leading to larger angular components in the eigenvectors and therefore improved visibility of the corresponding modes in the top-mass degrees of freedom.

Raising the COM also alters the geometric interpretation of the offset parameters. In this model, the values of d are defined relative to the shifted COM plane, meaning the attachment points lie effectively closer to the COM. Consequently, the effective lever arm responsible for pitch-longitudinal coupling is reduced, and we observe weaker cross-coupling between these degrees of freedom.

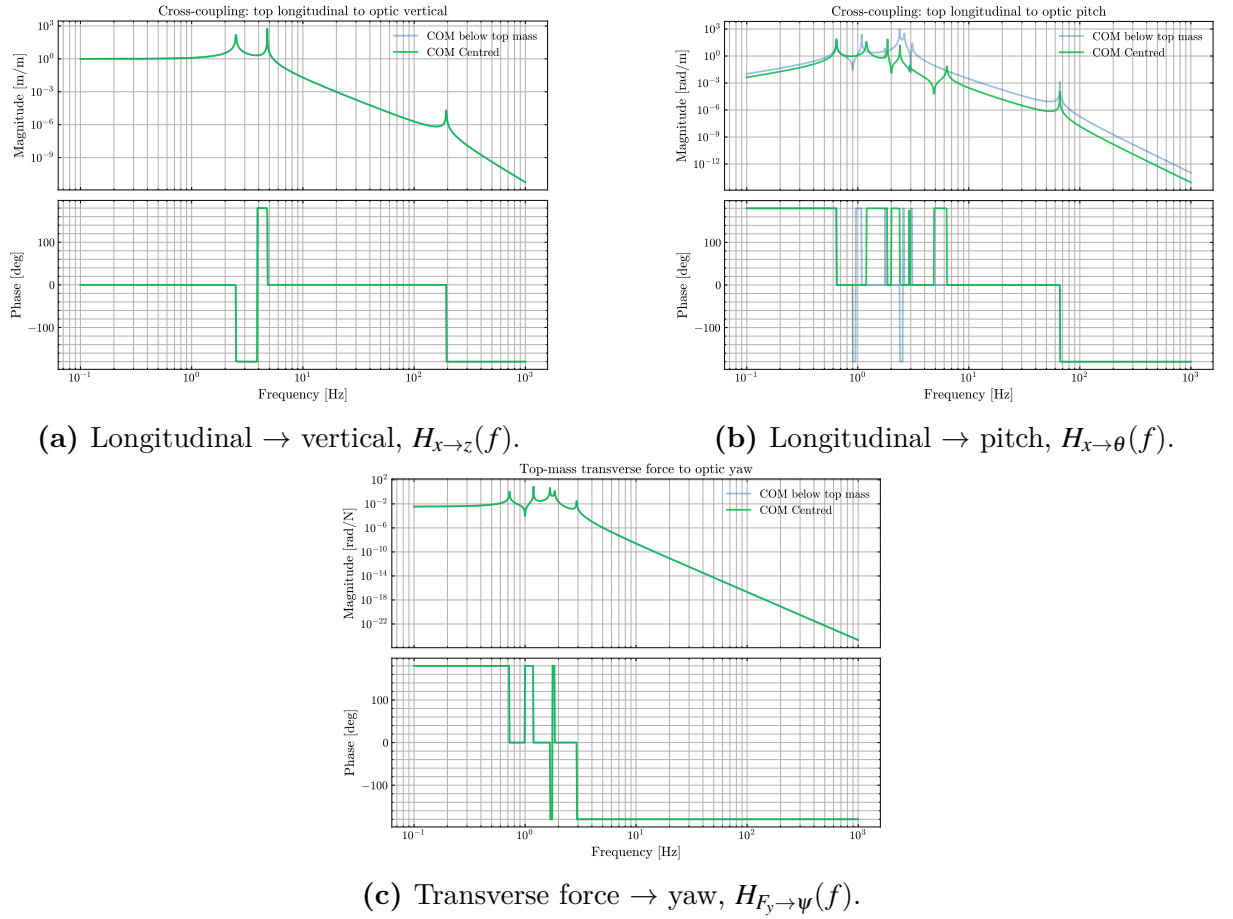


Figure 2.13: Comparison of cross-coupling transfer functions for the common-platform suspension model with the top-mass centre-of-mass (COM) *centred* and with the COM *lowered* relative to the top-mass reference plane. In each panel, the magnitude and phase are overlaid for the two configurations (baseline COM-centred model and COM-lowered model), illustrating how the COM shift modifies the coupling pathways from the applied disturbance to the optic motion for pitch motion [improving the performance].

2.7.1 Takeaway

Moving the top-mass COM closer to the geometric reference plane of the top mass is desirable from both a mode-visibility perspective (Figure 2.12) and in terms of reducing longitudinal-to-pitch (Figure 2.13). In practice, this implies incorporating compensating masses on the top mass to offset the COM shift due to the recoil structure located below the common platform.

2.8 Making the final stage longer- More Crystal, less Metal

We observe four modes, two bounce and two roll modes in the observation band (see Figure 2.7 and Table 2.2). The filter for the active control system is rolled off at frequencies below the observation band so that the sensor and actuator noise doesn't trickle down to the final stage of the suspensions. So to be able to damp these modes we want to reduce their frequencies. One of the ways this can be done is by increasing the length of the final stage relative to the other stages. In Figure 2.15 we show how increasing the length changes the frequency of the bounce and roll modes. Although this is what is generally desired, for our purposes we can damp these modes interferometrically.




Figure 2.14: Overlay of mode-visibility scatter plots for a set of lower stages length configurations. Opacity indicates the length setting for each configuration (from lowest length to highest length for the bottom stage), the total length of the suspension is maintained constant.




Figure 2.15: Variation of the roll and bounce mode frequencies in the observation band as a function of the final suspension-stage length, with the upper-stage lengths held fixed. The observation band contains two degenerate roll modes and two degenerate bounce modes. As the final-stage length increases, the out-of-phase bounce mode becomes increasingly cross-coupled to pitch motion of the upper mass, and the annotation in the plot indicated the mode-visibility value for this mode.

Beyond the overall decrease in the roll and bounce resonances, several mode-specific trends emerge in Figure 2.14. Most resonances improve; however, Mode 8— a pure yaw mode at 1.204 Hz with zero visibility at the common platform— shows a decrease in its visibility at the middle mass as the final-stage length l_3 increases.

Conversely, Mode 27, a high-frequency bounce mode, begins to cross-couple into platform pitch as l_3 is increased as shown in figure 2.15. The mode at 2.907 Hz, which is a pure yaw mode, is not visible at the common platform but its visibility at the intermediate stage increases drastically with increase in l_3 as seen in figure 2.14. This happens to a longitudinal-pitch mode at 1.769 Hz.

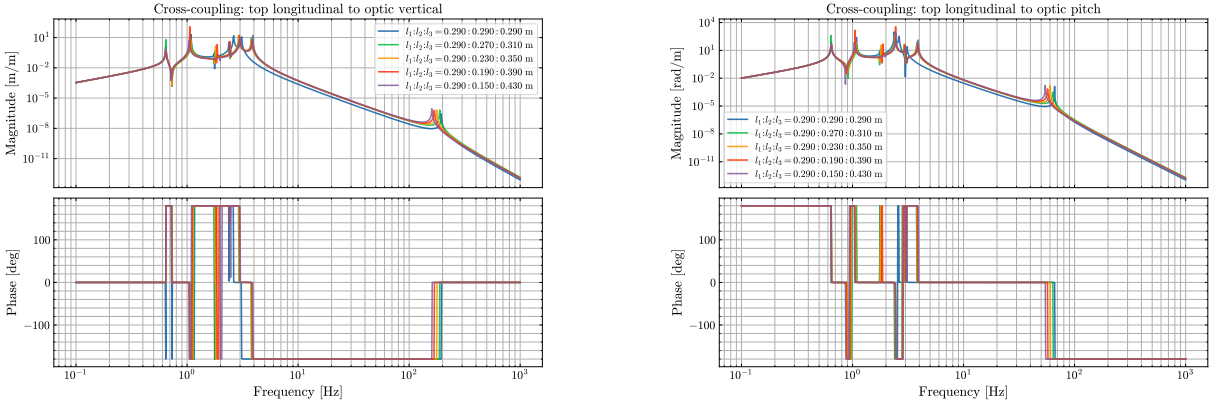
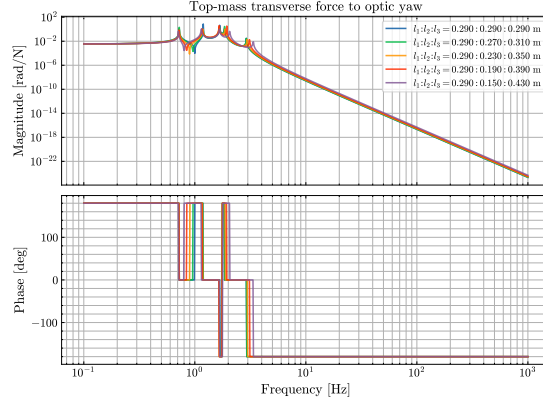
(a) Longitudinal \rightarrow vertical, $H_{x \rightarrow z}(f)$.(b) Longitudinal \rightarrow pitch, $H_{x \rightarrow \theta}(f)$.(c) Transverse force \rightarrow yaw, $H_{F_y \rightarrow \psi}(f)$.

Figure 2.16: Comparison of cross-coupling transfer functions for the common-platform model as the lower-stage lengths are varied. Each panel overlays magnitude and phase for five length sets $(l_1 : l_2 : l_3)$: 0.290 : 0.290 : 0.290, 0.290 : 0.270 : 0.310, 0.290 : 0.230 : 0.350, 0.290 : 0.190 : 0.390, and 0.290 : 0.150 : 0.430 m. The transfer functions shown are (a) Structure longitudinal excitation to optic vertical displacement, (b) Structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation.

Looking at the cross-coupling transfer functions (Figure 2.16) we do not see any trend. In the top-mass transverse force disturbance to optic yaw rotation transfer function we see a decrease in the height of the peaks with increasing l_3 .

2.8.1 Takeaway

From a mode-visibility point of view, a longer lower stage is preferable: on average, the mode-visibility values increase across the resonance-mode forest. From a cross-coupling point of view, there is little difference between the models. Hence, subject to mechanical constraints, using more crystal and less metal is clearly favourable.

2.9 Massive Top Mass

It is likely that the final common–platform (top–mass) design will be appreciably heavier than the main chain suspension chains beneath it for the main cavity optics. To anticipate the impact of this change, we carry out a parametric study of the *resonance–mode forest* as the platform mass is increased. The overlay in Figure 2.17 illustrates how the mode frequencies and their visibilities at the platform evolve with mass, highlighting which families shift, merge, or remain robust.




Figure 2.17: Overlay of mode–visibility scatter plots for several common–platform mass configurations. Each point marks the visibility of a normal mode at its eigenfrequency; colour encodes DOF and marker encodes where the mode is sensed (platform, cross–coupled, or intermediate mass). Opacity monotonically maps the mass setting from lightest (most transparent) to heaviest (least transparent), enabling direct visual comparison of frequency shifts and visibility changes across configurations.

As the platform mass m_1 is increased, the mode visibility at the top mass generally decreases. The primary mechanism is the increase in the platform moments of inertia, which raises the rotational/translation impedance of the top stage and makes platform DOFs harder to excite. This trend, although beneficial for mitigating cross-couplings, is not ideal for other modes.

From the mode-visibility plot we identify a cross-coupled mode at ~ 3 Hz: a bounce resonance that couples into pitch at the top mass. For $m_1 = 8.4$ kg, this resonance is in fact classified as a pitch mode, but the underlying bounce–pitch coupling still persists, as also evidenced by the transfer-function comparisons (Figure 2.18). Increasing the common-platform mass increases its moment of inertia, making the platform harder to rotate; correspondingly, the mode-visibility plot shows that this cross-coupling decreases as m_1 is increased. The same increase in mechanical impedance, however, reduces the direct visibility of several useful mode families at the platform—most notably the yaw modes between 1 and 3 Hz and the roll and bounce modes between 3 and 5 Hz—so that most modes appear weaker at the top mass in Figure 2.17.

Two exceptions stand out. A roll mode and a bounce mode between 2 and 3 Hz exhibits an increase in platform visibility as m_1 grows. By contrast, the *high-frequency* pitch, bounce, and roll doublets near ~ 66 Hz, ~ 195 Hz, and ~ 258 Hz are essentially *insensitive to m_1* and show negligible movement or visibility change across the mass sweep.

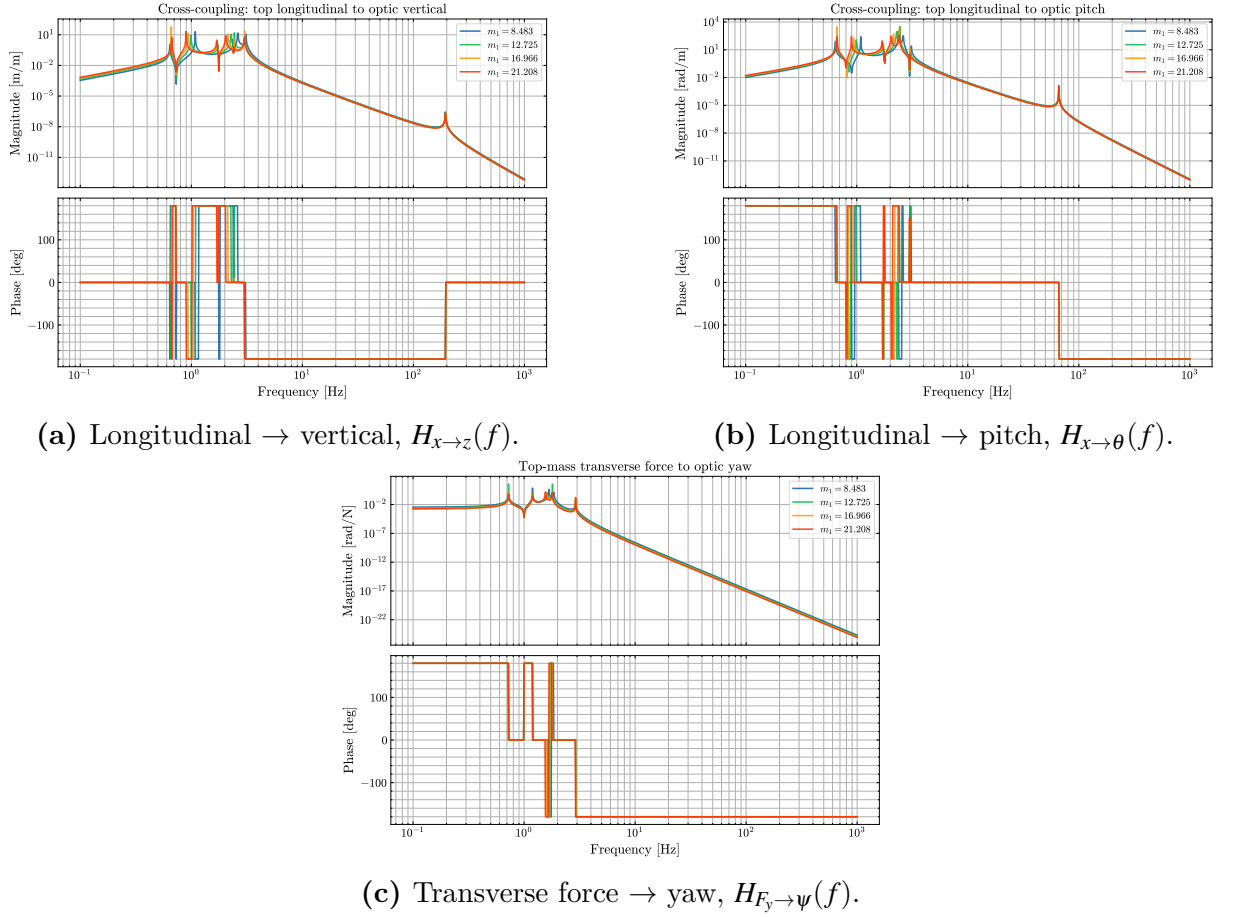


Figure 2.18: Comparison of cross-coupling transfer functions for the top-heavy common-platform configurations, parametrised by the common-platform mass m_1 (8.483, 12.725, 16.966, and 21.208 kg). Each panel overlays the magnitude and phase for the corresponding coupling path: (a) Structure longitudinal excitation to optic vertical displacement, (b) structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation.

Looking at the cross-couplings (Figure 2.18), the zeros in the longitudinal-to-vertical transfer function remain essentially unchanged and the resonance peaks maintain comparable heights. However, we can control where these peaks sit in frequency by tuning the mass of the common platform. The > 10 Hz features in both the longitudinal-to-vertical (Figure 2.18a) and longitudinal-to-pitch (Figure 2.18b) transfer functions show little dependence on the platform inertia, which is consistent with these resonances being dominated by motion of the bottom stages. For the transverse-force-to-yaw cross-coupling (Figure 2.18c), we see a more consistent improvement as the mass is increased; this is expected because increasing the top-mass moment of inertia raises its rotational impedance and makes it harder to excite yaw motion.

2.9.1 Takeaway

From a mode-visibility point of view, a lower common-platform mass is desirable. From a cross-coupling point of view, a more massive top mass is preferable. This therefore represents a trade-off between mode visibility and cross couplings. Given that we will use local control at the middle stage in any case, it may be reasonable to compromise on mode visibility and increase the top-mass mass to reduce cross couplings.

2.10 Asymmetric chains

We introduced an asymmetry in the d_4 values (see Fig. 2.19) between the two chains to investigate its effect on the resonance modes. We chose to vary d_4 because it corresponds to the monolithic stage: the technology for ear-bonding to the silicon mass is still under investigation, unlike the top stages where we will use metal wires as in existing GW detectors. This implies that the wire attachment point at the optic for the input chain may differ from that of the end chain and it is worth investigating how this asymmetry affects the suspension dynamics. We also want to see if introducing asymmetry breaks certain degeneracies with modes involving the lower stages.

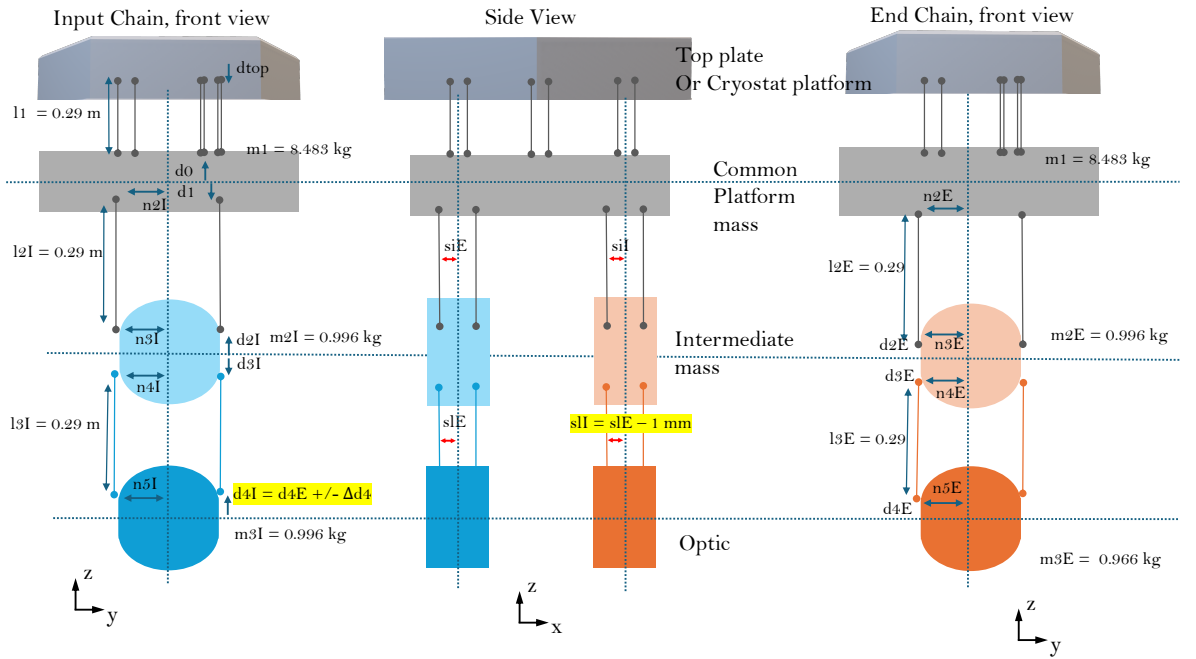


Figure 2.19: Common Platform Model highlighting asymmetry in d_4 between the chains.

We investigated three asymmetric cases with difference in d_4 between the two chains:

1. $\Delta d_4 = -1$ mm,
2. $\Delta d_4 = +1$ mm,
3. $\Delta d_4 = +5$ mm,
4. $\Delta d_4 = -1$ mm and $\Delta s_l = -1$ mm.

Interestingly, from the transfer-function comparisons in Fig. 2.20, the $\Delta d_4 = +5$ mm case is very similar to the symmetric scenario, whereas the $\Delta d_4 = \pm 1$ mm cases deviate more noticeably. From an overall behaviour point of view, however, all cases remain broadly similar. The main benefit is improved common-mode rejection, as illustrated in Fig. 2.21.

We also observe that for $\Delta d_4 = +5$ mm the pitch and roll modes begin to localise to either optic, which compromises common-mode rejection, as shown in Fig. 2.22. This is expected with the asymmetry. Asymmetry in d values at the optic stage doesn't seem to affect the mode-visibility of the modes in general.

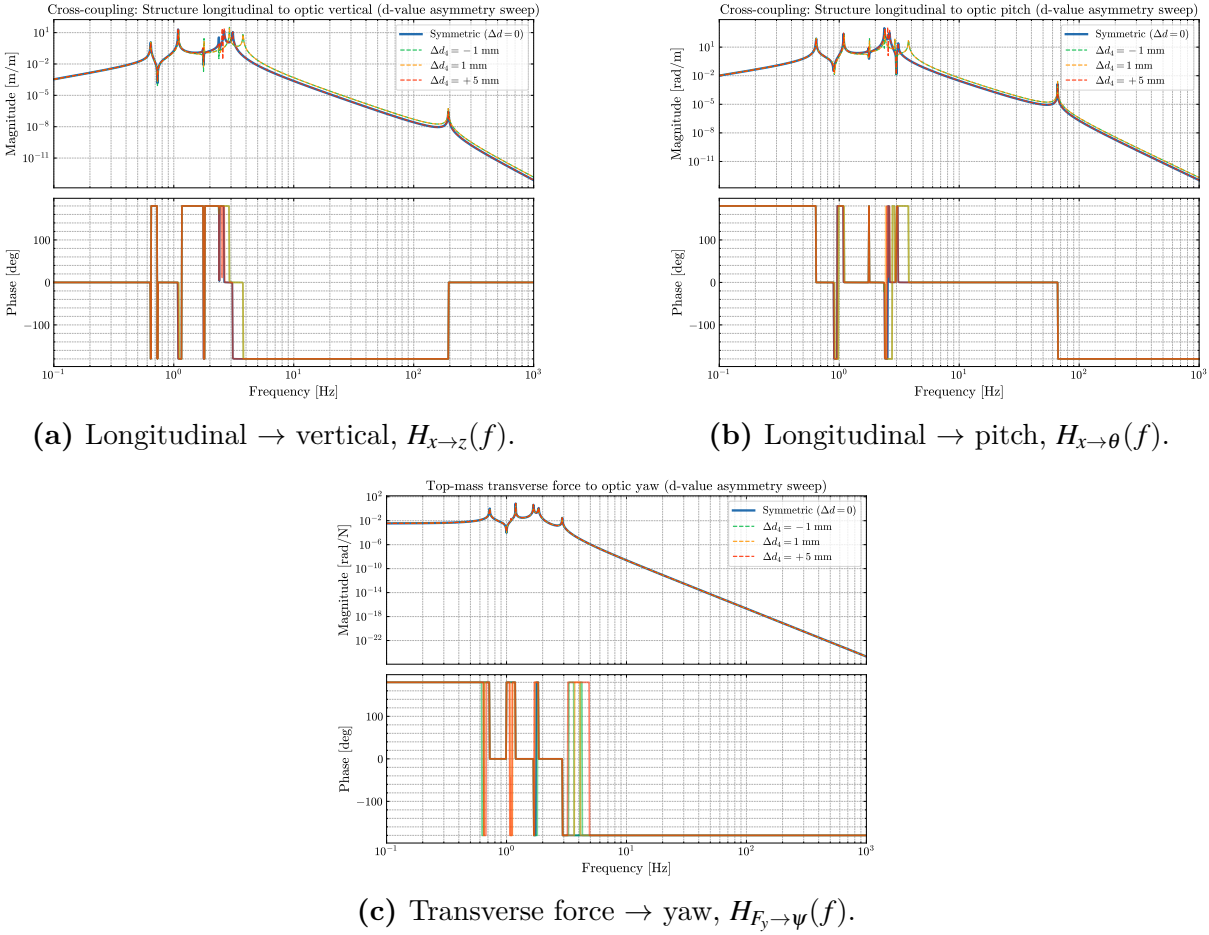


Figure 2.20: Comparison of cross-coupling transfer functions for the common-platform suspension model under d_4 -value asymmetry between the two suspension chains. The symmetric baseline ($\Delta d_4 = 0$) is compared against three asymmetric cases: $\Delta d_4 = -1$ mm, $\Delta d_4 = +1$ mm, and $\Delta d_4 = +5$ mm. Each panel overlays the magnitude and phase for (a) structure longitudinal excitation to optic vertical displacement, (b) structure longitudinal excitation to optic pitch rotation, and (c) top-mass transverse force disturbance to optic yaw rotation. The transfer functions are taken to the end chain in (a) and (b) and input chain in (c).




Figure 2.22: Example of mode-shape localisation for $\Delta d_4 = +5$ mm, where pitch/roll motion becomes more strongly associated with one optic chain, reducing the degree of common-mode rejection.




Figure 2.21: Differential (common-mode rejected) cryostat seismic length-noise projection for the common-platform model under Δd_4 asymmetry, overlaid with the CryoTriple large-upper-blade (LUB) reference model [81].

We also observed splitting of the high-frequency pitch doublet (modes 25 and 26; see Table 2.2) in the case where both d and s asymmetries are present, with a separation of approximately 7 Hz. This splitting does not significantly affect the cross-coupling transfer functions, nor the seismic-noise common-mode rejection.

From the cryostat seismic length-noise in Figure 2.21, we find that the common-mode rejection, even for d and s asymmetries that lie within the assumed tolerances, offers 4 orders of magnitude extra isolation over the Large Upper Blade spring triple chain model. We also observe that the change in common-mode rejection is not significant for increasing asymmetry.

2.10.1 Takeaway

The mode visibility at the top mass is not strongly affected by asymmetries between the two suspension chains. We also find that the cross-coupling transfer functions remain very similar to the symmetric case. The primary effect is the lifting of degeneracies (mode splitting), which is beneficial in practice. Overall, these results suggest that modest chain-to-chain asymmetry is acceptable, and can even be advantageous.

2.11 Conclusion and Future Development

This chapter presented the modelling and design study of a common-platform (split-pendulum) triple-stage suspension for the Glasgow Cryogenic Interferometer Facility (GCIF). The motivation for the architecture is to exploit mechanical common-mode rejection between the two cavity chains while retaining the seismic isolation benefits of a multistage pendulum. Using the established Mathematica modelling framework [47], supported by a Python-based workflow for faster evaluation, we assessed the resonance-mode structure, mode visibility at the common platform and the intermediate stage, and the cross-coupling transfer functions relevant for longitudinal length noise and control spillover. The results were benchmarked against a conventional triple-stage model [81].

The baseline two-chain model with a cage-like recoil structure exhibits a reorganised mode spectrum relative to a conventional triple chain: the shared platform introduces systematic couplings that cluster modes into mixed families, with the expected longitudinal-pitch and transverse-roll pairings accompanied by occasional off-family mixing (e.g. transverse-yaw and bounce leakage into longitudinal-pitch modes). For the baseline configuration, a subset of modes are insufficiently visible at the top mass and therefore require sensing and

damping at the intermediate stage. This directly informs the local-control architecture: platform sensors/actuators alone are not sufficient to robustly damp the full resonance-mode forest, and intermediate-stage instrumentation is required for commissioning margin.

A set of parametric studies was then used to explore practical design levers:

- **Top-mass COM location.** Shifting the common-platform centre of mass closer to the geometric reference plane improves mode visibility at the platform and reduces the longitudinal-to-pitch coupling. In practice, this motivates incorporating compensating masses to offset COM shifts introduced by the recoil structure.
- **Longer final stage (more crystal, less metal).** Increasing the lower-stage length improves the average mode visibility across the resonance-mode forest. The cross-coupling transfer functions are broadly similar across the length sweep, with only modest changes in specific peaks, indicating that the primary benefit is improved controllability rather than a systematic reduction in coupling.
- **Top-heavy platform.** Increasing the platform mass (and hence its moments of inertia) generally reduces platform mode visibility, but improves certain cross-coupling pathways—most clearly for transverse-force to yaw coupling—by increasing the rotational impedance of the platform. This highlights a real trade-off between *visibility* (controllability at the platform) and *cross-coupling* (isolation against off-axis disturbances). Given that local control at the intermediate stage is required in any case, a modest increase in platform mass may be acceptable if it provides a meaningful reduction in problematic coupling.
- **Chain asymmetry.** Introducing modest chain-to-chain asymmetry in the monolithic-stage geometry (d_4 , and one combined d and s case) does not strongly affect top-mass mode visibility, and the cross-coupling transfer functions remain similar to the symmetric baseline. The main dynamical effect is the lifting of degeneracies (mode splitting), including a measurable splitting of the high-frequency pitch doublet when both d and s asymmetries are present. In the cryostat seismic length-noise projection, however, the common-mode rejection can degrade substantially even for




Figure 2.23: Preliminary CAD rendering of the cryostat showing the current work-in-progress integration of the common-platform suspension assembly. The model is primarily used for envelope, clearance, and access checks, and does not represent a final mechanical design.

asymmetries within assumed tolerances; importantly, the degradation appears to plateau rather than worsen indefinitely with increasing asymmetry. Despite this, the common-platform concept still offers additional isolation margin relative to the conventional triple suspension.

In Figure 2.23 the CAD integration of the cryostat with the common-platform suspension assembly is illustrated.¹ This model is used to verify overall envelope constraints, available height, and mechanical clearances between the suspension chain, cryostat internals, and access ports. The integration design is ongoing, and detailed interfaces will be iterated as the mechanical design matures and component-level drawings are finalised.

1. Cryostat mechanical design CAD developed by Russell Jones.

Overall, the modelling supports the feasibility of the common-platform approach for GCIF and clarifies the key design trade-offs: (i) maximising resonance-mode visibility at the platform favours a lighter, more easily rotated top stage with a well-centred COM, while (ii) minimising cross-coupling favours a more massive, higher-inertia platform. These constraints must be balanced alongside cryogenic practicality (thermal conduction through fibres, blade feasibility at the lowest stage, and tolerance-driven asymmetries) and the sensing/actuation plan for commissioning.

Future development

Several extensions follow naturally from this work:

- **Control-oriented modelling.** Incorporate the measured sensor/actuator noise and realistic control filters in the Simulink model to quantify loop noise injection and stability margins for both platform and intermediate-stage control.
- **Damping layout and heat load.** Optimise the placement of shadow sensors and actuators across stages to ensure observability of low-visibility modes while minimising cryogenic heat load and wiring complexity.
- **Manufacturing tolerances and asymmetry survey.** Extend the asymmetry study to the full set of relevant geometric parameters (including s and n variations), and map tolerance-driven degradation of common-mode rejection to explicit alignment and bonding tolerances.
- **Reaction-chain realism.** Revisit the four-chain (main+reaction) configuration to assess additional mode families and coupling pathways once a mechanically realistic recoil implementation is available.
- **Prototype validation.** Use early suspension prototypes to validate the predicted mode families, mode visibility trends, and cross-coupling transfer functions, and feed measured parameters back into the model for iteration.

Development of a Cryogenic Shadow Sensor

3.1 Introduction

In the Introduction chapter we discussed the importance of local control in suspensions. The objective of this project is to characterise light-emitting diodes (LEDs), quadrant photodiodes (QPDs), and various flag geometries to develop a high-sensitivity, low-noise, and high-dynamic-range cryogenic shadow sensor with minimal thermal load on the cryostat. The sensor design must satisfy all these requirements simultaneously. A high sensitivity (or equivalently, a high signal-to-noise ratio, SNR) is essential to achieve lower, measurable suspension noise levels. A wide dynamic range is particularly important for cryogenic suspensions, as small differences in the thermal expansion coefficients and wire lengths of the suspension elements can lead to misalignments during cryogenic cooldown.

KAGRA employs photoreflective displacement sensors with a dynamic range of approximately 10 mm, which is sufficient to accommodate the expected thermal contractions during cooldown [174]. In cryogenic interferometers such as KAGRA, thermal contraction in the suspensions during cooldown is expected to occur particularly in the vertical and yaw degrees of freedom (on the order of 1 mrad). Pitch motion also occurs, though typically to a lesser extent (around 150 μ rad). Additionally, the suspension wires may

twist as they thermally contract, producing angular displacements of a few milliradians [173]. For the multi-material suspension planned for the Glasgow Cryogenic Interferometer Facility (GCIF), calculations are currently in progress to establish the corresponding displacement requirements. It is equally important to quantify the heat load from the LED, ensuring that the sensor can be properly incorporated into the overall thermal budget of the cryostat (see Figure 3.14).

Alternative approaches for cryogenic displacement sensing include capacitive sensors [35, 156] and Linear Variable Differential Transformers (LVDTs) [166]. While we investigated capacitive sensing for displacement readout, we ultimately adopted shadow sensing, as it can be integrated more straightforwardly into the existing suspension architectures used in ground-based gravitational-wave detectors.

The prototype Glasgow Optical Sensor and Electromagnetic Actuator, or GOSEM, consists of the following key components:

1. **Optical Sensor:** A quadrant photodiode (QPD), model **FCI InGaAs Q3000**.
2. **Emitter:** An infrared LED (IRLED) with a central wavelength of 1200 nm. Devices tested include **LED1200L** and **LED2100E** from Thorlabs, and **L13072-0120L** and **L13072-0120K** from Hamamatsu.
3. **Flag:** A mechanical element that partially occludes or transmits light to the QPD, forming the shadow pattern used for displacement sensing.

For the first prototype the mount design for the shadow sensor is similar to the BOSEM design [41][40], but parts are modified to suit cryogenic requirements (Section 3.5).

3.1.1 Photodiodes

Quadrant photodiodes (QPDs) are position-sensitive optical detectors widely used in precision metrology, beam alignment, and displacement sensing applications [143, 96]. A QPD consists of four photodiode elements arranged in a quadrant configuration on a single substrate, each generating a photocurrent proportional to the incident optical power in its respective region. By comparing the photocurrents from the four quadrants, both the position and displacement of a light spot or shadow on the detector surface can be determined with high spatial resolution.

The InGaAs photodiode employed in this work is a heterostructure device supplied by OSI. In such photodiodes, incident radiation is absorbed within the semiconductor material through interaction with electrons, resulting in an electrical signal. This signal arises from a redistribution of the electronic energy states. The response of these detectors exhibits a selective wavelength dependence, typically expressed as responsivity in units of amperes per watt (A/W) of incident optical power [165, 144].

Noise Sources in photodiodes

Several fundamental and practical noise sources contribute to the overall performance of a photodiode and its readout chain [140, 144]. In the context of optical displacement sensing, these noise terms limit the achievable signal-to-noise ratio and, depending on operating conditions, different mechanisms can dominate:

1. **Shot noise:** Fluctuations in the detected photocurrent arising from the discrete arrival of photoelectrons (and hence the quantised nature of charge), producing a white current-noise floor that scales with the square root of the average photocurrent [144].

2. **Photon–electron (conversion) noise:** Additional statistical fluctuations associated with the stochastic conversion of incident photons into collected charge carriers, which can include effects of non-unity quantum efficiency and carrier-collection statistics [140, 144].
3. **Gain noise:** In photodiodes that exhibit internal multiplication (e.g. avalanche photodiodes), the amplification process introduces excess noise due to fluctuations in the multiplication gain; for unity-gain PIN photodiodes this term is typically negligible [144, 140].
4. **Receiver (electronic) noise:** Noise generated by the front-end electronics (e.g. Johnson–Nyquist noise of resistors, amplifier voltage and current noise, and $1/f$ noise), which is shaped by the transimpedance and filtering stages and can dominate at low photocurrent or outside the measurement band [144].
5. **Background noise:** Photocurrent fluctuations produced by unwanted optical power incident on the detector (e.g. stray light and thermal background), which adds both a DC offset and an associated shot-noise contribution proportional to the background-induced photocurrent [140].
6. **Dark-current noise:** Noise associated with the detector dark current, arising from mechanisms such as thermal generation–recombination, tunnelling, and surface leakage currents; this contributes an additional shot-noise-like term even in the absence of illumination and generally increases with temperature [165, 140].

In this project, we use FCI InGaAs Q3000 quadrant photodiodes (QPDs) sourced from OSI. These devices are employed in KAGRA’s photoreflexive displacement sensors, providing a well-established starting point with prior demonstration of reliable operation at cryogenic temperatures.

3.1.2 Light Emitting Diodes

A Light Emitting Diode or an LED is a p-n junction device that generates monochromatic incoherent light under forward bias. In early 1900s, a British scientist Henry Joseph Round, discovered that when current was passed through silicon carbide crystals it generates light. What happens is recombination of electrons and electron holes in a process called electroluminescence.

Noise Sources in LEDs

At the device level, LED intensity noise arises from several carrier-dynamics processes that imprint power fluctuations on the spontaneous emission rate:

1. **Shot (quantum) noise:** Poisson statistics of photon emission produce a white floor that, when referred to absolute power, scales with \bar{P} and with detection bandwidth; in relative form, the shot limit is constant for fixed \bar{P} [78].
2. **Generation–recombination noise:** Stochastic trapping/de-trapping and recombination events modulate the carrier population and radiative rate, often giving *Lorentzian* PSD components with corner frequencies set by trap lifetimes. A distribution of lifetimes produces quasi- $1/f$ behaviour [123, 76].
3. **$1/f$ (flicker) noise:** In III–V LEDs (LEDs made from Group 3-5 elements) this is widely observed at low frequencies (Hz–kHz) and is linked to defects, dislocations, and interface inhomogeneities that modulate current flow and quantum efficiency. Its level is strongly device- and bias-dependent and is a known reliability/ageing indicator [141, 49].
4. **Driver/electronics coupling:** Residual current noise from the LED driver transfers to optical power with a roughly linear electro-optical gain; good current regulation and filtering are therefore essential in low-Relative Intensity Noise (RIN) setups.

Lowering the temperature alters carrier lifetimes, trap kinetics, and nonradiative pathways in III–V LEDs, so the mechanisms that drive optical–power fluctuations change with T . At the same time, stochastic trapping/detrapping (generation–recombination) processes slow as traps become less thermally activated, shifting Lorentzian corner frequencies and often enhancing the relative prominence of generation–recombination/ $1/f$ components at low f [118, 124]. In parallel, external quantum efficiency and emission characteristics are themselves temperature dependent, so the mapping from drive-current fluctuations to optical-power fluctuations is modified as T changes [126]. Empirically, cryogenic investigations on commercial LEDs for space/astrophotonics report substantial changes in brightness and operating behaviour across $\sim 77\text{--}300\text{ K}$, reinforcing that RIN should be measured directly as a function of temperature rather than inferred from room-temperature values [125, 57].

We tested several LED models to assess their suitability for use in GOSEMs: the Thorlabs LED1200L and LED1200E, and the Hamamatsu L13072-0120L and L13072-0120K. While KAGRA employs the LED1200E in their photoreflective displacement sensors, we were motivated to investigate alternative devices with smaller divergence angles and higher optical efficiency.

3.2 GOSEM Electronics

3.2.1 Light Emitting Diode

In a GOSEM, a 1200 nm light-emitting diode (LED) serves as the primary near-infrared (NIR) illumination source. The 1200 nm wavelength was chosen due to the known availability of cryogenically compatible LED’s and the availability of InGaAs QPDs that can operate at this wavelength and at low cryogenic temperatures. We initially employed the Thorlabs LED1200L to perform testing, which is a near IR diode with a spectral emission

peak centred at 1200 nm. The device comprises compound semiconductor heterostructures formed from *Indium Gallium Arsenide Phosphide* (InGaAsP), grown on an Indium Phosphide (InP) substrate, as reported in the manufacturer’s datasheet [167]. The InGaAsP material system enables precise bandgap tuning in the 1.0–1.6 μm range when lattice-matched to InP, making it especially suitable for optoelectronic components operating in the short-wavelength infrared (SWIR) region [163, 131].

LEDs based on InGaAsP typically feature a double heterostructure (DH) configuration, where the active light-emitting layer is enclosed between wider-bandgap cladding layers. This structure enhances carrier confinement and overall internal quantum efficiency by suppressing non-radiative recombination processes [90]. These NIR LEDs are widely utilised in spectroscopy, biomedical diagnostics, and integrated photonic systems [148].

At cryogenic temperatures, particularly around 4 K, the bandgap of InGaAsP increases owing to reduced phonon interactions, leading to a blue shift in the emission wavelength [175, 21]. Moreover, lower phonon scattering at such temperatures improves carrier mobility and minimises non-radiative decay channels, which can lead to enhanced radiative efficiency [185]. These properties render InGaAsP LEDs highly applicable to low-temperature photonics and quantum optics, including experiments involving superconducting detectors.

Driver Circuit

The LED driver circuit used in the shadow sensor system is designed as a low-side current regulation topology. It employs an ADA4891 high-speed, low-noise operational amplifier, an N-channel MOSFET, and a 10 Ω current-sensing resistor. This configuration enables precise control of the current passing through the LED, which is essential for maintaining consistent light output in shadow sensing applications where accuracy and repeatability are critical.

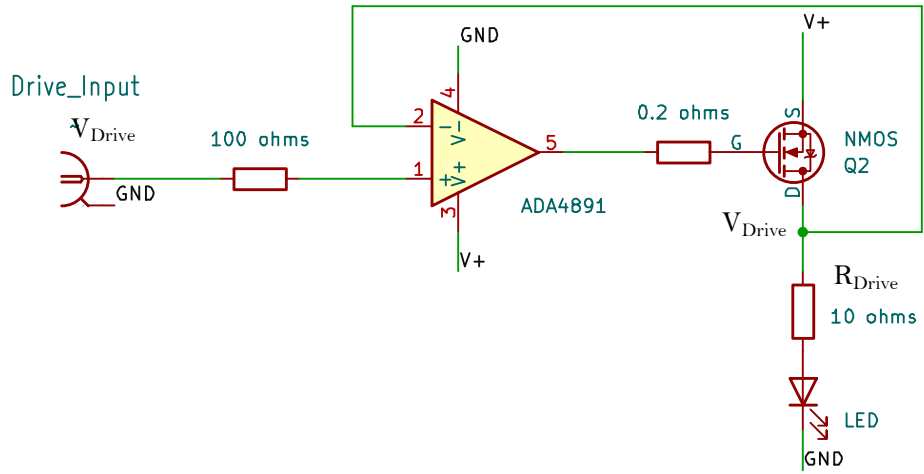


Figure 3.1: Closed-loop LED driver circuit using the ADA4891 operational amplifier and an N-channel MOSFET for low-side current regulation. The op-amp senses the voltage across the 0.2Ω current-sense resistor and adjusts the MOSFET gate drive to regulate LED current based on the *Drive_Input* control voltage. The feedback loop ensures precise LED current control independent of LED forward voltage or supply variations, making it suitable for optical sensing applications such as shadow sensors.

In this implementation of low side operation as shown in circuit diagram Figure 3.1, the LED is connected between the +5 V ground and the current-sensing resistor (R_{drive}). The other end of the resistor connects to the drain of the N-MOSFET, while the MOSFET's source is tied to positive power supply. The voltage across R_{drive} , which is proportional to the LED current, is fed into the inverting input ($-$) of the ADA4891. A drive voltage used to control the current through the LED is applied to the non-inverting input ($+$) of the op-amp. The op-amp, which compares these two voltages and adjusts its output to control the gate of the MOSFET, creates a closed-loop system that regulates the LED current according to the relationship:

$$I_{\text{LED}} = \frac{V_{\text{Drive}}}{R_{\text{drive}}}.$$

For instance, if $V_{\text{Drive}} = 1 \text{ V}$ and $R_{\text{drive}} = 100 \Omega$, the resulting LED current will be 10 mA.

The ADA4891 is particularly well-suited for this role due to its rail-to-rail output, low input offset voltage, and high slew rate of $55\text{ V}/\mu\text{s}$. These characteristics allow it to respond quickly and accurately to variations in the control voltage, ensuring stable and responsive LED current regulation. This is especially important in the shadow sensor, where small deviations in LED brightness can significantly affect detection accuracy.

Using a low-side MOSFET configuration offers several practical benefits, including simplified gate driving, accurate current sensing, and compatibility with control voltages from digital-to-analog converters (DACs) or analog signal sources. Additionally, the programmable nature of the drive voltage enables dynamic control of LED brightness, which can be valuable in adaptive sensing environments.

In summary, the combination of the ADA4891 op-amp, N-channel MOSFET, and low-side current sensing forms an efficient and precise LED driver tailored for high-performance shadow sensing. The circuit delivers excellent linearity, fast response, and reliable operation under varying conditions, contributing significantly to the overall performance and reliability of the sensor system.

3.2.2 Quadrant Photodiode

In this project, a QPD is employed to detect the motion of a shadow cast by a flag onto the photodiode surface, forming the basis of a cryogenic shadow sensor. The sensor is designed for high sensitivity, low noise, and a wide dynamic range, while maintaining minimal thermal load on the cryostat. Among commercially available devices, InGaAs-based QPDs, such as the FCI-InGaAs-Q3000 from OSI Optoelectronics, are well-suited for this application due to their high responsivity in the near-infrared range, low dark current, and excellent performance at low temperatures [122].

The QPD operates by reverse-biasing each photodiode element to improve linearity, reduce junction capacitance, and enhance response speed [85]. The sum of the photocurrents provides a measure of total incident optical power, while differential signals between opposing quadrants yield the lateral position of the shadow. This capability makes QPDs ideal for precise, contactless position detection in optomechanical and cryogenic sensing systems.

The Quadrant Photodiode (QPD) amplifier PCB is a critical component of the shadow sensor system. It is responsible for converting the small photocurrents generated by each quadrant of the QPD into usable voltage signals, enabling accurate detection of the shadow position. This circuit processes optical signals produced when a flag partially or fully intercepts light from the LED, casting a shadow onto the QPD. I designed the circuit using Autodesk EAGLE software.

The PCB is structured into three primary functional blocks: **trans-impedance amplifiers (TIAs)**, a **summing amplifier**, and a set of **instrumentation amplifiers** for signal subtraction.

The four QPD anodes are connected to a dedicated TIA. These amplifiers convert the input photocurrent (I_{photo}) into a voltage signal using a high-value feedback resistor (R_f) and a capacitor to roll off high frequency noise. The general form of the output voltage is given by:

$$V_{\text{out}} = -I_{\text{photocurrent}} \times R_f$$

The use of high-gain, low-noise op-amps with low input bias current is essential to preserve signal integrity and avoid current leakage.

A summing amplifier stage is implemented to calculate the sum of adjacent quadrants. Test pins are located on the PCB to extract these signals so that total incident light intensity can be measured and is useful for normalisation and intensity compensation. The output of the quad summing amplifier is:

$$V_{\text{sum}} = -(V_i + V_j),$$

where $i, j = A, B, C, D$ are the outputs of the quad TIA.

To detect positional changes caused by the moving shadow, the circuit includes two instrumentation amplifiers that compute the difference between opposing quadrants as illustrated in Figure 3.3. Specifically:

$$X_{\text{diff}} = (V_A + V_B) - (V_C + V_D)$$

$$Y_{\text{diff}} = (V_A + V_C) - (V_B + V_D)$$

These differential outputs allow the system to determine the lateral (X) and vertical (Y) displacement of the shadow across the sensor surface. Instrumentation amplifiers are used to provide high common-mode rejection, precision gain, and minimal offset, which are critical for extracting small signal variations against background light or noise.

Special care is taken in the PCB layout to minimise parasitic capacitance and feedback loop instability in the TIA stage. Shielding and ground planes are incorporated to reduce electromagnetic interference. Component selection prioritises low-noise performance, thermal stability, and matched resistor networks to ensure balanced subtraction and accurate summing.

In summary, the QPD amplifier PCB provides a complete analog signal conditioning path from photocurrent generation to differential position signal output. Its modular architecture—composed of TIAs, a summing amplifier, and precision subtractors—facilitates accurate and robust shadow position detection, forming the core of the optical sensing system. See Appendix Section A for PCB images.

3.2.3 Noise Modelling and Measurements

Noise in the shadow sensor consists of electronic noise from the circuitry, QPD dark noise and LED shot noise. Let us first estimate electronic noise from the QPD amplifier board. The circuit noise consists of Johnson noise from the resistor, voltage noise and voltage noise from the opamp. An LTspice simulation of the QPD amplifier circuit was run and the simulation was compared to measurements in Figure 3.2.




Figure 3.2: Comparison of measured and simulated noise spectral densities of the QPD amplifier output: The traces show noise measurements for two output channels (AB-CD and BD-AC) obtained using a Moku spectrum analyser, the modelled electronic noise from LTSpice simulations, and the baseline noise floor of the Moku. The simulation aligns well with the measured spectra above 10 Hz, validating the noise performance of the QPD circuit and we see many ground loops on top of the baseline curve.

3.2.4 Discussion

We designed and fabricated custom circuit boards comprising a low-noise, stable constant-current driver for the LED and a transimpedance amplifier to convert the QPD photocurrents into measurable output voltages. The measured electronic noise spectra show good agreement with the corresponding LTSpice simulations. In the next section, we investigate the impact of different flag geometries on sensor performance.

3.3 Python Modelling of Different Flag Geometries

In this section we will go through analytical models of different flag geometries and compare them from heat load and dynamic range point of view. These are python models. After understanding these models we built a prototype sensor based on one of the flag geometries.

Introducing the definitions first and then diving deeper into rectangular flags and circular aperture flags.

Geometry and Beam

Let the quadrant photodiode (QPD) be a disc of radius R with aperture

$$\mathcal{A} = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq R^2\}.$$

We model the optical field as a circular Gaussian with $1/e^2$ radius w , normalised so that the *total optical power* is independent of w (see figure 3.4a). Writing P for the total power in the beam plane,

$$I(x, y; w) = \frac{2P}{\pi w^2} \exp\left(-\frac{2(x^2+y^2)}{w^2}\right), \quad (3.1)$$

which satisfies

$$\iint_{\mathbb{R}^2} I(x, y; w) dx dy = P \quad \text{for all } w.$$

The total power transmitted by the bare aperture (no flag) is therefore

$$P_{\text{ap}}(w) = \iint_{\mathcal{A}} I(x, y; w) dx dy = P \left(1 - e^{-2R^2/w^2}\right). \quad (3.2)$$

Equivalently, the fraction of the beam captured by the QPD is $F(R, w) = P_{\text{ap}}(w)/P = 1 - \exp[-2(R/w)^2]$.

Flags and Transmission Masks

A shadow sensor works by sensing changes in photocurrent caused by changes in the amount of light obstructed by a flag when it moves. The flag here is attached to the moving mass whose motion we want to damp and the rest of the optics forming the sensor is attached to a fixed reference. It is important to understand the effects of different flags on the response of the sensor to identify the best geometry depending on the requirements and purpose.

The flag is centred at $(x_0, 0)$ and translated along the x -axis. Two geometries are considered: (i) a rectangular flag that blocks light and produces a shadow at the centre of the QPD in its nominal position, and (ii) a circular aperture, representing the complementary case in which light transmission rather than blocking defines the signal.

(i) Rectangular flag (vertical bar) of width a

As an *occluder* (opaque, see Figure 3.3) its transmission is

$$\chi_{\text{rect}}(x, y; x_0, a) = \begin{cases} 0, & |x - x_0| \leq a/2, \\ 1, & \text{otherwise.} \end{cases} \quad (3.3)$$

A transmissive slit uses $1 - \chi_{\text{rect}}$.

(ii) Circular aperture of diameter d

As an *aperture* (see Figure 3.6a) its transmission is

$$\chi_{\text{circ}}(x, y; x_0, d) = \begin{cases} 0, & (x - x_0)^2 + y^2 \leq (d/2)^2, \\ 1, & \text{otherwise.} \end{cases} \quad (3.4)$$

A transmissive circular aperture (hole) uses $1 - \chi_{\text{circ}}$.

Net transmission with aperture

$$T(x, y; x_0) = \chi(x, y; x_0, \cdot) \mathbf{1}_{\mathcal{A}}(x, y), \quad (3.5)$$

where χ is chosen from (3.3) or (3.4) (or their complements for transmissive cases) and $\mathbf{1}_{\mathcal{A}}(x, y)$ is a mask that is 1 for coordinates inside the QPD aperture and 0 otherwise.

Detected Powers and normalised Signal

Left/right powers (within the finite aperture) are

$$L(x_0) = \iint_{\mathcal{A} \cap \{x < 0\}} I(x, y; w) T(x, y; x_0) dx dy, \quad (3.6)$$

$$R(x_0) = \iint_{\mathcal{A} \cap \{x > 0\}} I(x, y; w) T(x, y; x_0) dx dy, \quad (3.7)$$

$$P(x_0) = L(x_0) + R(x_0). \quad (3.8)$$

The normalised QPD signal is defined as

$$S(x_0) = \frac{L(x_0) - R(x_0)}{L(x_0) + R(x_0)} \in [-1, 1]. \quad (3.9)$$

(If $S = (R - L)/(L + R)$ is preferred, all results hold with an overall sign flip.)

Transmitted fraction onto the QPD (flag centred)

Let the QPD active area be the disk $\mathcal{A} = \{(x, y) : x^2 + y^2 \leq R^2\}$. A Gaussian beam of total optical power P and $1/e^2$ radius w has

$$I(x, y; w) = \frac{2P}{\pi w^2} \exp\left[-\frac{2}{w^2}(x^2 + y^2)\right]. \quad (3.10)$$

With *no* flag, the power that lands on the QPD is

$$P_{\text{ap}}(w) = \iint_{\mathcal{A}} I(x, y; w) dx dy = P \left[1 - e^{-2(R/w)^2}\right], \quad (3.11)$$

so the corresponding fraction of the LED power that reaches (and is confined to) the QPD is

$$F_{\text{ap}}(w; R) \equiv \frac{P_{\text{ap}}(w)}{P} = 1 - e^{-2(R/w)^2}. \quad (3.12)$$

The complement,

$$F_{\text{miss}}(w; R) \equiv 1 - F_{\text{ap}}(w; R) = e^{-2(R/w)^2}, \quad (3.13)$$

is the fraction that *misses* the diode (relevant for stray-light management and cryostat heating).

General flag description

Place the flag center at $(x_0, 0)$ and, for a centered analysis, set $x_0 = 0$. Let $\tau(x, y) \in [0, 1]$ be the flag transmissivity mask (unity where light is transmitted, zero where it is blocked).

The fraction of power transmitted onto the QPD in the presence of the flag is

$$F_{\text{on}}(w; R | \tau) \equiv \frac{\iint_{\mathcal{A}} I(x, y; w) \tau(x, y) dx dy}{\iint I(x, y; w) dx dy} = \frac{1}{P} \iint_{\mathcal{A}} I(x, y; w) \tau(x, y) dx dy. \quad (3.14)$$

Raw centre slope

The geometric transfer of optical power between halves is captured by the raw slope

$$\left. \frac{d}{dx_0} (L - R) \right|_{x_0=0} \approx \frac{[L-R](+\Delta) - [L-R](-\Delta)}{2\Delta},$$

which does not include normalisation by the total . Here Δ is the infinitesimal flag displacement about the zero point. It highlights how the flag edges intersect regions of high intensity inside the aperture (Figure 3.5c). For small values of w/R the flag intercepts the high intensity section of the Gaussian beam and hence the response is high. But as the beam size gets bigger the intensity drops to conserve power so this causes the responsivity to drop. For larger flag widths the tail end of the Gaussian beam is intercepted by the edges of the flag thus the slope is small and as the beam size gets bigger the the amount of power in the tail gets bigger causing the responsivity to increase. For a fixed beam size, as the flag width increases the edges of the flag intercepts tail end of the Gaussian beam eventually making the slope smaller.

Normalised Centre Slope

The small-signal responsivity at the origin is

$$\left. \frac{dS}{dx_0} \right|_{x_0=0} \approx \frac{S(+\Delta) - S(-\Delta)}{2\Delta}, \quad \Delta > 0. \quad (3.15)$$

Dynamic Range: Physical Definition and Interpretation

As the flag translates across the detector, the normalised quadrature signal

$$S(x_0) = \frac{L(x_0) - R(x_0)}{L(x_0) + R(x_0)},$$

rises from zero near the origin (where the left/right illumination is balanced), then approaches a *plateau* where the output stops changing appreciably with further motion. Physically, this saturation occurs when one side of the split detector has become effectively fully shadowed (or fully illuminated) by the flag, so that additional displacement no longer alters the left–right power ratio.

We define the *dynamic range* (DR) as the span of flag positions over which the output lies between the two saturated regions on either side:

$$\text{DR} \equiv x_+ - x_-,$$

where x_+ (right) and x_- (left) are the first positions at which $|S|$ reaches its plateau level.

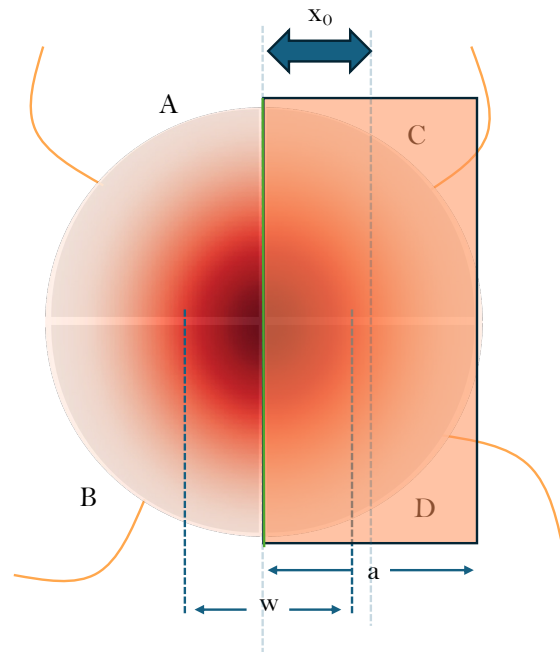


Figure 3.3: Rectangular flag moving across a QPD and intercepting the incident beam—the green highlighted edge shows the leading edge.

Geometric intuition

For an opaque flag, saturation begins when the flag's *leading edge* first reaches the part of a detector half that dominates the remaining light (see figure 3.3). Thus, wider flags enter saturation sooner and *reduce* DR. The beam size w (Gaussian $1/e^2$ radius) shapes where power resides within the finite aperture: concentrating power near the center or near the rim shifts the position at which additional motion ceases to change the left–right balance. Consequently, DR varies with both flag size (geometry) and beam size (weighting), and is independent of the chosen sign convention for S .

Relation to linear operation

Near $x_0 = 0$ the response is approximately linear with a slope $dS/dx_0|_0$. The DR measures how far one can move before the output becomes uninformative due to saturation. In design, there is a trade-off: increasing the flag size typically sharpens the slope due to decrease in the total power falling on the QPD (improving small-signal sensitivity) while shrinking the DR (reducing usable range). Both metrics, central sensitivity and dynamic range, are needed to characterise performance.

Shot-Noise-Limited Position Floor

Let $I(x_0)$ be the total photocurrent through the QPD at x_0 , with a nominal no-flag current I_0 . If we scale by the fraction of power transmitted when the flag is positioned centrally (T_{centre}), then,

$$I_{\text{center}} = I_0 \frac{P(0)}{P_{\text{ap}}(w)} = I_0 T_{\text{center}}. \quad (3.16)$$

For the normalised signal $S = (L - R)/(L + R)$, error propagation with uncorrelated shot-noise currents in L and R gives a signal amplitude spectral density (ASD)

$$S_{\text{SN}} = \sqrt{\frac{2q}{I_{\text{center}}}}, \quad (3.17)$$

where q is the electron charge. The equivalent displacement ASD is

$$x_{\text{SN}} = \frac{S_{\text{SN}}}{\left| \frac{dS}{dx_0} \right|_{x_0=0}}. \quad (3.18)$$

If $\frac{dS}{dx_0}$ is reported in units of $1/R$, then x_{SN} is in $R/\sqrt{\text{Hz}}$; multiply by the physical R to obtain $\text{m}/\sqrt{\text{Hz}}$. (If your system holds photocurrent constant, set $I_{\text{center}} = I_0$ in (3.17).)

3.3.1 Rectangular Flag

Conventional shadow sensors use rectangular flags. A similar flag geometry could be incorporated into a QPD based shadow sensor. The concept of shadow sensing with a rectangular flag moving across the QPD is shown in Figure 3.4. As the flag moves the difference signal can be plotted as in Figure 3.4b. For sensing we use the normalised intensity instead of raw intensity because this helps with removing LED power drifts as will be discussed in future sections. Understanding how flag and beam dimensions influence the sensing metrics is essential for selecting geometries that are optimal for a given application.


Heat maps for the rectangular flag

To map performance over geometry and beam size, five plots are shown in Figure 3.5 as functions of the normalised flag width a/R and Gaussian beam radius w/R , evaluated with a finite QPD aperture of radius R . The beam is modelled as


$$I(x, y; w) = \frac{2P}{\pi w^2} \exp\left[-\frac{2(x^2 + y^2)}{w^2}\right],$$

so that the *total* optical power P is independent of w and the power captured by the bare aperture (no flag) is

$$P_{\text{ap}}(w) = P\left(1 - e^{-2R^2/w^2}\right) \cdot 1q$$



(a) Rectangular flag moving across a QPD and intercepting the incident beam.



(b) Simulated normalised signal $(L-R)/(L+R)$ in solid blue and raw signal $(L-R)$ in dashed green as a function of flag centre position x_0/R . Here $L = A+B$ and $R = C+D$. The flag width $a = R$ and beam size $w = R$.

Figure 3.4: Concept and response of shadow sensing with a rectangular flag and QPD. (a) Schematic illustration of the flag partially blocking the incident Gaussian beam on a QPD. (b) Corresponding simulated normalised and raw differential signal measured from the QPD as the flag translates across the beam.

Optical power delivered to the QPD

Let the QPD be the disk $\mathcal{A} = \{(x, y) : x^2 + y^2 \leq R^2\}$ and the flag an opaque vertical strip $\{|x| \leq a/2\}$ centred at $x = 0$. A Gaussian beam of total power P and $1/e^2$ radius w has

$$I(x, y; w) = \frac{2P}{\pi w^2} \exp\left[-\frac{2}{w^2}(x^2 + y^2)\right].$$

The power that would hit the QPD with no flag is

$$P_{\text{ap}}(w) = \iint_{\mathcal{A}} I(x, y; w) dx dy = P \left[1 - e^{-2(R/w)^2}\right]. \quad (3.19)$$

The power blocked by the strip is (which corresponds to the numerator in equation 3.14)

$$P_{\text{block}}(a, w; R) = \iint_{\mathcal{A} \cap \{|x| \leq a/2\}} I(x, y; w) dx dy.$$

Writing the integral explicitly and carrying out the y-integration,

$$\begin{aligned} \frac{P_{\text{block}}(a, w; R)}{P} &= \frac{2}{\pi w^2} \int_{-a/2}^{a/2} e^{-2x^2/w^2} \int_{-\sqrt{R^2-x^2}}^{+\sqrt{R^2-x^2}} e^{-2y^2/w^2} dy dx \\ &= \frac{2}{w} \sqrt{\frac{2}{\pi}} \int_0^{\min(a/2, R)} e^{-2x^2/w^2} \operatorname{erf}\left(\frac{\sqrt{2}}{w} \sqrt{R^2 - x^2}\right) dx, \end{aligned}$$

which we record as

$$F_{\text{block}}^{\text{rect}}(a, w; R) = \frac{2}{w} \sqrt{\frac{2}{\pi}} \int_0^{\min(a/2, R)} e^{-2x^2/w^2} \operatorname{erf}\left(\frac{\sqrt{2}}{w} \sqrt{R^2 - x^2}\right) dx. \quad (3.20)$$

Then the fraction of the light power that falls on the QPD is

$$F_{\text{on}}^{\text{rect}}(a, w; R) = 1 - F_{\text{block}}^{\text{rect}}(a, w; R), \quad (3.21)$$

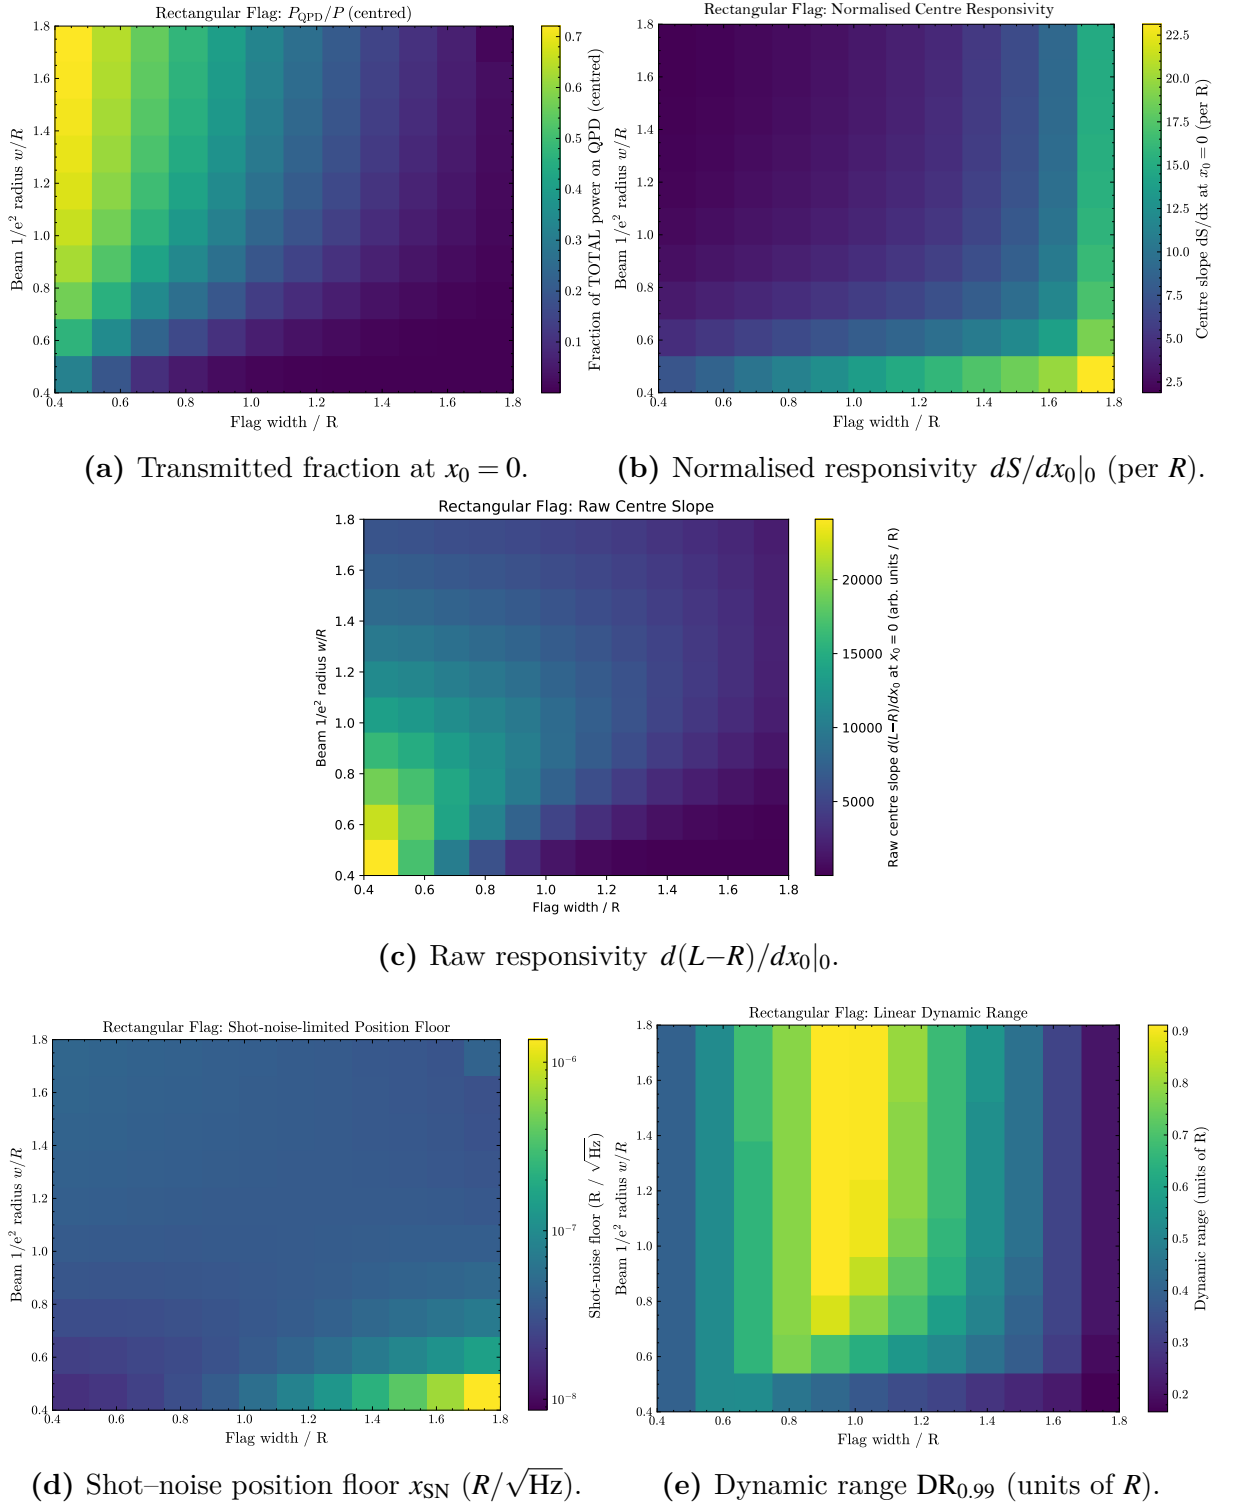


Figure 3.5: Heat map characterisation of a QPD shadow sensor with a rectangular flag. Axes are normalised to the QPD radius R and the total beam power is held fixed.

Raw centre slope

The geometric transfer of optical power between halves is captured by the raw slope

$$\left. \frac{d}{dx_0} (L-R) \right|_{x_0=0} \approx \frac{[L-R](+\Delta) - [L-R](-\Delta)}{2\Delta}, \quad (3.22)$$

which does not include normalisation by the total power. It highlights how the flag edges intersect regions of high intensity inside the aperture (Figure 3.5c). For small values of w/R the flag intercepts the high intensity section of the Gaussian beam and hence the response is high. But as the beam size gets bigger the local intensity drops as the total power in the beam is conserved so this causes the responsivity to drop. For larger flag widths the tail end of the Gaussian beam is intercepted by the edges of the flag, thus the slope is small and as the beam size gets bigger the amount of power in the tail gets bigger causing the responsivity to increase. For a fixed beam size, as the flag width increases the edges of the flag intercepts the tail end of the Gaussian beam eventually making the slope smaller.

Normalised centre slope

The small-signal responsivity of the *normalised* output,

$$\left. \frac{dS}{dx_0} \right|_0 \approx \frac{S(+\Delta) - S(-\Delta)}{2\Delta}, \quad (3.23)$$

with

$$S = \frac{L-R}{L+R}, \quad (3.24)$$

is shown in Figure 3.5b. Compared to the raw slope it is reduced where the denominator $L+R$ is large. For a fixed beam size the responsivity of the QPD increases with flag width because the total power falling on the QPD decreases.

Shot–noise floor

Assuming the total photocurrent at the centre position I_{centre} , the shot-noise ASD of the normalised signal is

$$S_{\text{SN}} = \sqrt{\frac{2q}{I_{\text{centre}}}}, \quad (3.25)$$

and the position floor is

$$x_{\text{SN}} = \frac{S_{\text{SN}}}{|dS/dx_0|_0}, \quad (3.26)$$

reported in $R/\sqrt{\text{Hz}}$ (Figure 3.7d). The best noise is obtained where the slope is high yet sufficient light still reaches the QPD. Performance worsens for narrow bars (slope too small) and very wide bars (too little light remains).

Dynamic range

We define the dynamic range as the interval of flag–centre positions x_0 over which the normalised response

$$S(x_0) = \frac{L - R}{L + R}, \quad (3.27)$$

is approximately linear, i.e. before it reaches the side plateau (or the absolute peaks when a true plateau is absent). The heat map shows that the dynamic range is largest at *intermediate* bar widths and diminishes towards both extremes. For very wide bars, one half of the QPD becomes effectively fully shadowed or fully illuminated shortly after leaving the origin, so saturation is reached early and the usable range shrinks as a/R increases. For very narrow bars the contrast is weak and the response turns over at modest $|x_0|$, again shortening the linear region. The beam size matters as well for intermediate flag widths: small w/R means the intensity of the beam is high, so a given bar width drives one half of the detector into saturation sooner, further compressing the dynamic range. But a large beam size means the beam is more uniform across the QPD and the rate of change of slope is small. In this case we get the best dynamic range. In practice, choosing $a/R \sim \mathcal{O}(1)$ together with a moderate w/R offers a good compromise between a steep central slope and a delayed onset of saturation.

3.3.2 Circular Flag

A circular flag is an aperture flag that traverses across the QPD to produce signal as shown in Figure 3.6.

Heat maps for a circular aperture flag

In Figure 3.7 we have summarised the behaviour of a QPD shadow sensor when the moving flag is a *circular aperture* (i.e. a transmissive hole) that scans laterally across the detector. As in the rectangular case, the incident beam is a circular Gaussian of $1/e^2$ radius w , and the QPD has a finite circular aperture of radius R . To compare different beam sizes fairly, the beam normalisation keeps the *total optical power* fixed; widening the beam therefore lowers the on-axis peak but spreads the power over a larger area. Both the axes in the plots are normalised to R .

Optical power delivered to the QPD


The optical power incident on the QPD is shown in Figure 3.7a)) as a thermal plot. Let the aperture be the disk $\mathcal{C}_r = \{(x, y) : x^2 + y^2 \leq r^2\}$ (with radius $r = d/2$), coaxial with the QPD disk $\mathcal{A} = \{x^2 + y^2 \leq R^2\}$. The Gaussian beam is $I(x, y; w) = \frac{2P}{\pi w^2} \exp[-2(x^2 + y^2)/w^2]$. The power that reaches the QPD through the aperture is the integral over the intersection $\mathcal{A} \cap \mathcal{C}_r$. Because the two disks are concentric, $\mathcal{A} \cap \mathcal{C}_r$ is simply the disk of radius $\min(R, r)$, so in polar coordinates

$$\begin{aligned} P_{\text{on}}^{\text{circ}}(r, w; R) &= \iint_{\mathcal{A} \cap \mathcal{C}_r} I(x, y; w) dx dy, \\ &= \int_0^{2\pi} \int_0^{\min(R, r)} \frac{2P}{\pi w^2} e^{-2\rho^2/w^2} \rho d\rho d\theta, \\ &= P \left[1 - \exp\left(-2 \frac{\min(R, r)^2}{w^2}\right) \right]. \end{aligned}$$


Dividing by the no-flag power on the diode (3.19) gives the fraction of the beam power that reaches (and is confined to) the QPD:

$$F_{\text{on}}^{\text{circ}}(r, w; R) = \frac{P_{\text{on}}^{\text{circ}}(r, w; R)}{P} = \left[1 - \exp\left(-2 \frac{\min(R, r)^2}{w^2}\right) \right]. \quad (3.28)$$

This reduces to the expected limits: $F_{\text{on}}^{\text{circ}} \rightarrow 0$ as $r \rightarrow 0$ (pin-hole closed) and $F_{\text{on}}^{\text{circ}} \rightarrow 1$ for $r \geq R$ (aperture fully covers the QPD).



(a) Circular aperture flag moving across a QPD and intercepting the incident beam.



(b) Simulated normalised signal $(L - R)/(L + R)$ and raw signal $(L - R)$ as a function of flag centre position x_0/R . Here $L = A + B$ and $R = C + D$. Here the aperture diameter is R same as the QPD radius and the beam radius w is same as the QPD radius R .

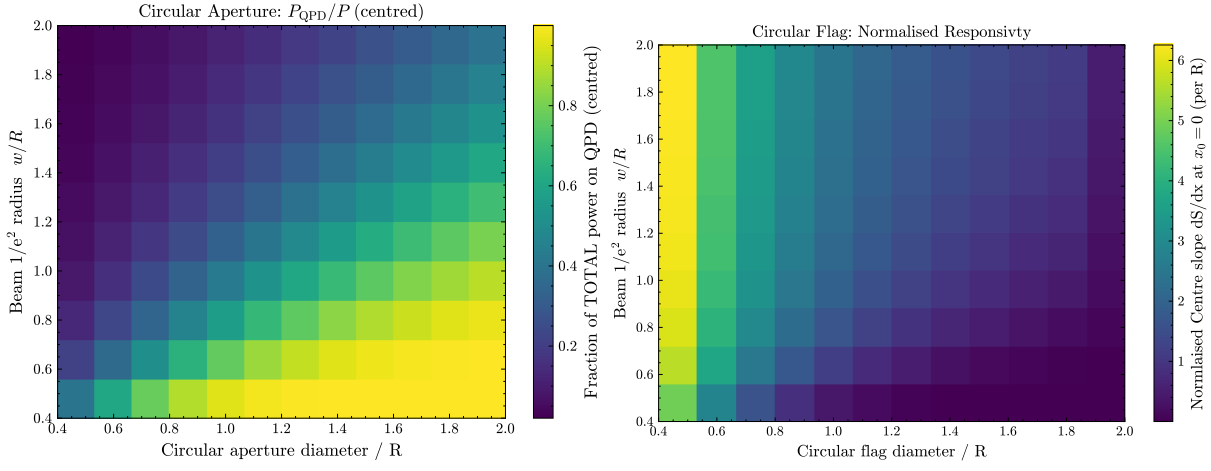
Figure 3.6: Concept and response of shadow sensing with a rectangular flag and QPD. (a) Schematic illustration of the flag partially blocking the incident Gaussian beam incident on the QPD. (b) Corresponding simulated normalised and raw differential signal measured from the QPD as the flag translates across the beam.

Raw Response


The slope or the raw response (see Figure 3.7c) of the detector is most sensitive when a small lateral motion of the aperture moves a significant amount of light from one half of the QPD to the other. This occurs when the size of the beam is comparable to the size of the aperture. And the peak responsivity occurs when the beam size is small and intense. A large flag samples the tail end of the Gaussian spot thus making the detector response at the centre smaller than that for smaller flag diameters. A larger beam is more spread out, as we want to keep the power across the detector constant to maintain a constant shot noise performance, thus making the slope smaller.

Normalised Responsivity

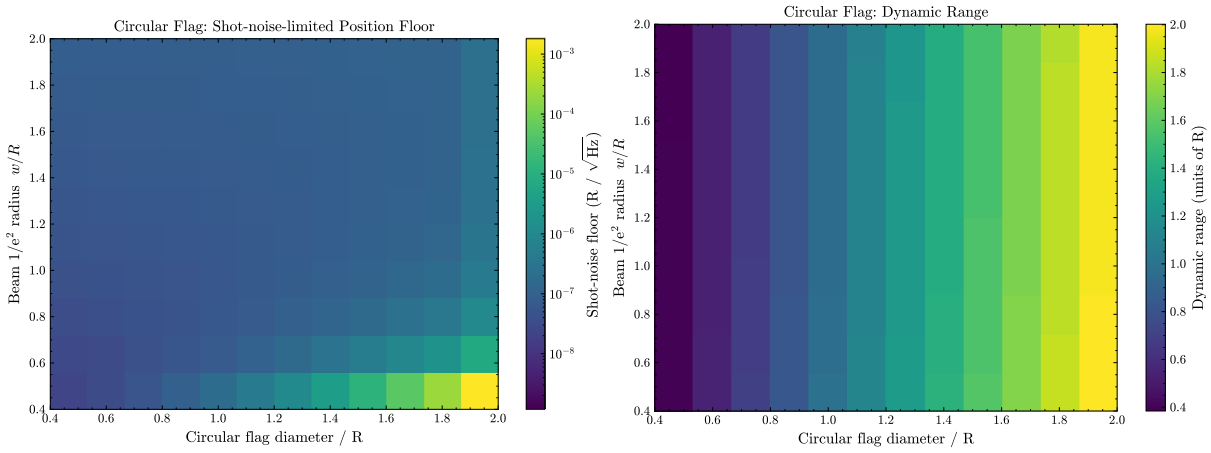
The differential signal, $S = (L - R)/(L + R)$, as function of beam radius and aperture diameter is shown in Figure 3.7(b). Sensitivity at the operating point is set by the gradient dS/dx at $x_0 = 0$. This not only depends on the slope but also the total power detected. For a given flag size, as the beam size gets bigger, power detected gets smaller and so does the slope of the raw response. But the decrease in the latter is faster thus the normalised responsivity increases with an increase in the beam size. Given a fixed beam size, the response decreases with increasing aperture size due to the increase in the amount of light falling on the detector and the fact that the flag samples the tail of the Gaussian spot.



(a) Transmitted fraction at $x_0 = 0$. (b) Normalised responsivity $dS/dx_0|_0$ (per R).



(c) Raw responsivity $d(L-R)/dx_0|_0$.



(d) Shot-noise position floor $x_{\text{SN}} (R/\sqrt{\text{Hz}})$. (e) Dynamic range $\text{DR}_{0.99}$ (units of R).

Figure 3.7: Heat map characterisation of a QPD shadow sensor with a circular aperture flag. Axes are normalised to the QPD radius R and the total beam power is held fixed.

Shot-noise floor

The shot noise limited displacement noise floor scales like

$$x_{\text{SN}} \propto \frac{\sqrt{2q/I_{\text{PD}}}}{|dS/dx|},$$

so it improves when *less photocurrent* is collected (more blocked light) and when the *raw slope* is steep (see Figure 3.7(d)). We know from (Figure 3.7c) that the centre slope is the steepest for flag diameters and beam size of similar dimensions. And from (Figure 3.7a) we see that the amount of light is high smaller flag diameters and larger beams. This is reflected well in the heat map for shot noise limited sensitivity where the shot noise floor exhibits a broad valley for where the blocked fraction is high and the slope is steep and then for large flag diameter and small beam size the slope is negligible and the amount of light detected by the detector is maximum hence the shot noise performance is at its worst.

Dynamic range

Following the same definition for dynamic range as for rectangular flags, we infer the following (see Figure 3.7(d) for the thermal plot):

- *Very small holes* reach saturation quickly: one half becomes effectively fully illuminated early in the scan. They maintain linearity over only a short distance then as soon as the aperture moves past from one side of the QPD to the other the normalised signal plateaus.
- *Very large holes*: The dynamic range is limited by the size of the aperture.

Overall, a larger aperture size will result in a larger dynamic range. But this clashes with the shot noise performance and clearly means we must accept trade-offs in the design depending on the requirements of the sensor application.


3.3.3 Design implications

The objective of this project is to evaluate photodiodes, LEDs, and flag geometries that can enable a shadow sensor with minimal thermal load on the cryostat while achieving a linear dynamic range on the order of millimetres and high sensitivity. By design, the sensitivity is limited by shot noise. In addition, it is important to consider the linearity of the sensing response, as non-linearities can introduce unwanted harmonics into the control signals.


The two primary factors influencing the choice of beam and flag dimensions are the dynamic range and the shot-noise floor. From the heat maps, we observe that the circular aperture flag (Figure 3.7e) provides nearly twice the dynamic range of the rectangular flag (Figure 3.5e). For the circular geometry, the shot-noise floor scales as $\frac{1}{\sqrt{P}}$, so higher optical power is desirable, provided that the QPD is not saturated and the cryostat heat load remains acceptable. A sufficiently large beam waist w further reduces the shot-noise floor (see Figure 3.7(d)).

For the rectangular flag, the dynamic range is maximum at intermediate flag widths (Figure 3.5e). In this region of interest, the shot-noise floor remains approximately constant, while increasing the beam size helps to reduce the heat load on the cryostat.

In summary, the optimal operating region for the circular aperture flag lies toward the middle–right edge of the heat maps, whereas for the rectangular flag the most favourable region is near the middle–top edge. Below we examine the response curves and linearity within these regions.




(a) QPD difference-signal response versus normalised flag position x_0/R . Solid lines: normalised (difference/sum); dashed: raw. Beam $1/e^2$ radius varied at fixed aperture $d = 1.6R$: $w/R \in \{0.5, 0.8, 1.0, 1.4, 1.8\}$.




(b) Response versus x_0/R at fixed $w = 1.4R$ for different aperture diameters; $d/R \in \{1.0, 1.4, 1.6, 1.8\}$. Solid = normalised, dashed = raw.

Figure 3.8: Circular aperture flag: simulated QPD response for varying beam radius w and aperture diameter d in the region of interest of Figure 3.7 (max transmission, low shot-noise floor). Normalised (solid) and raw (dashed) signals are plotted versus the flag-centre position, normalised to the QPD active radius.

Now let's look at the response curves in the region of interest for the rectangular flag. The region of interest is the portion of design space where the combination of beam and flag dimensions yields the largest dynamic range (Figure 3.5). Our selection proceeds in two steps. First, we choose the *flag width* that maximises dynamic range. For beam radii that are sufficiently large ($w \gtrsim R$), the dynamic range is governed almost entirely by the flag geometry; varying w has little influence on the span, and the associated penalty in linearity is small (Figure 3.9a). Second, with the width fixed at its optimum (here illustrated with width = $1.6R$), we tune the *beam radius* to maximise the linearity of the response (Figure 3.9b). This choice incurs a modest degradation in the shot-noise floor; however, the loss in sensitivity can be offset by increasing the LED optical power, provided QPD saturation and cryostat heat-load limits are respected.



(a) Rectangular flag: QPD difference-signal response versus normalised flag position x_0/R . Solid lines show the normalised (difference/sum) signal; dashed lines show the raw difference. The *flag width* is varied at fixed beam radius $w = 1.4R$: $\text{width}/R \in \{0.60, 0.80, 1.00, 1.20, 1.40\}$.



(b) Rectangular flag: response versus x_0/R with *flag width* fixed at $1.6R$. Curves show the effect of the *beam radius*; colours encode $w/R \in \{0.5, 0.8, 1.0, 1.4, 1.8\}$. Solid = normalised, dashed = raw.

Figure 3.9: Rectangular aperture flag: simulated quadrant-photodiode response for variations in flag width and beam radius, for operating points taken from the thermal maps (Figure 3.5). Normalised (solid) and raw (dashed) signals are plotted against the flag-centre position, normalised to the QPD active radius R .

Takeaway

Within the design space explored here, the circular aperture flag is the preferred baseline geometry for GOSEMs. It offers a substantially larger usable (approximately linear) dynamic range—nearly a factor of two compared to the rectangular flag—and provides a straightforward route to improved shot-noise-limited sensitivity by increasing optical power and/or beam waist, subject to QPD saturation and cryostat heat-load constraints. The rectangular flag, in contrast, provides a steeper central slope and higher small-signal responsivity but a reduced dynamic range.

3.4 The Prototype

The prototype sensor is mounted on an aluminium baseplate, and the LED and QPD are mounted in an aluminium cylinder as shown in Figure 3.10a. A 9 mm aspheric lens is placed inside the cylinder which helps in collimating the LED beam. A moveable circular flag is held between the QPD and the LED holders so that the responsivity curve can be measured. Refer to Figure 3.20 which shows a photograph of the inside of the cryostat, this will be described in detail in a future section.

For the purposes of the study we chose a larger aperture diameter for the flag as this takes into account the larger beam size when measured at the photodiode due to diffraction around the flag.

3.4.1 In-Air Characterisation using Circular Aperture Flag


The image in Figure 3.10 the setup that was used to characterise the LED-lens-flag-QPD system in air is shown. The figure 3.10b indicates the responses at the centre and at displaced flag position. The response at the centre is lower than at other flag position because the flag edges intercept the tail ends of the Gaussian beam spot. But as the flag moves away from the centre the edges get closer to the centre of the beam spot hence increasing the response. The shot-noise performance is independent of whether the response is taken in normalised or raw form. The equations 3.26 and 3.25 can also be written as

$$S_{\text{SN}} = \sqrt{2qV_{\text{centre}}R_f}, \quad (3.29)$$


where R_f is amplifier feedback resistance.

$$x_{\text{SN}} = \frac{S_{\text{SN}}}{\left. \frac{d}{dx_0}(L - R) \right|_{x_0=0}}. \quad (3.30)$$

From the above equations, it is easy to see that when the flag is centred the amount of light falling on the photodiode is large and the slope is small because the flag is intercepting the less intense parts of the Gaussian beam away from the centre. This makes the shot noise floor higher. In contrast, when the flag is displaced the amount of light falling on the QPD is smaller due to the new shadow, but the slope is larger as explained above; this makes the shot-noise floor lower for displaced circular flag position. Now, the contrast between flag-centred and flag-displaced is less prominent for larger beam diameter and/or larger flag apertures. But going for larger beam diameter and/or larger flag apertures comes with a compromise with the responsivity hence shot noise floor for the sensor.



(a) Experimental setup for characterising the GOSEM in air. The LED illuminates the flag on a translation stage; the transmitted beam is detected on the QPD. During the measurement, the QPD incident power was $180 \mu\text{W}$ with a beam radius of 1.89 mm. The flag aperture diameter was 2.5 mm.



(b) Measured sensing curves corresponding to (a). Black: raw QPD left–right difference. Red: normalised (difference/sum). Local linear fits give 5.92 V/mm at $x = 0$ and 7.92 V/mm at $x = -0.62$ mm.

Figure 3.10: Characterisation of the GOSEM in air. (a) Optical/mechanical setup with LED illumination, flag aperture, and QPD detection. (b) Measured raw and normalised sensing curves with local-slope responsivities.




Figure 3.11: Relative intensity noise (RIN) spectra of the GOSEM output shown in a single plot. Traces display the normalised RIN with the flag displaced (green), the raw RIN with the flag displaced (orange), the RIN with the flag centred (blue), the calculated shot-noise limits for the displaced and centred configurations (dashed lines), and the dark-noise floor of the readout electronics (black). The comparison illustrates that normalisation suppresses excess LED intensity fluctuations when the flag is displaced, lowering the apparent noise floor toward the shot-noise level. The right-hand axis converts voltage spectral density to equivalent displacement noise using the measured responsivity.

For the response curve in Figure 3.10b the RIN noise curves are presented in Figure 3.11. When the flag is centred the intensity noise is equal on both sides and cancels out leaving you with shot noise, but when the flag is off-centre the common mode cancellation is not sufficient and we are limited by the actual intensity noise of the light. However, if you normalise by the total power from all four quadrants the RIN is reduced and once again you're left with the shot noise limit even for the off-centre flag position as described in the section below.

Normalisation of the Differential Signal in Shadow Sensing

In the shadow sensing configuration, the measured signal is obtained as the difference between two opposing segments of the photodiode, which we denote as the channels AB and CD . The raw differential signal is given by

$$I_{\Delta}(t) = I_{AB}(t) - I_{CD}(t), \quad (3.31)$$

where $I_{AB}(t)$ and $I_{CD}(t)$ represent the instantaneous photocurrents in the two channels. Each of these consists of a deterministic mean intensity level and small fluctuations,

$$I_{AB}(t) = \overline{AB} + \overline{AB} \delta_{\text{RIN},AB}(t) + \delta_{\text{SN},AB}(t), \quad (3.32)$$

$$I_{CD}(t) = \overline{CD} + \overline{CD} \delta_{\text{RIN},CD}(t) + \delta_{\text{SN},CD}(t), \quad (3.33)$$

where $\overline{AB}, \overline{CD}$ are the DC photocurrents, δ_{RIN} represent relative intensity noise (RIN), and δ_{SN} denote shot-noise fluctuations.

The raw difference therefore contains both the desired signal and residual intensity noise,

$$I_{\Delta}(t) = (\overline{AB} - \overline{CD}) + \overline{AB} \delta_{\text{RIN},AB}(t) - \overline{CD} \delta_{\text{RIN},CD}(t) + \delta_{\text{SN},AB}(t) - \delta_{\text{SN},CD}(t). \quad (3.34)$$

It is evident that RIN contributions from the two halves do not cancel perfectly, since they are weighted by the relative DC powers. This leads to residual RIN in the raw differential response.

To mitigate this effect, the differential signal is *normalised* by the total detected power,

$$I_{\Delta}^{\text{norm}}(t) = \frac{I_{AB}(t) - I_{CD}(t)}{I_{AB}(t) + I_{CD}(t)} I_{\text{max}}, \quad (3.35)$$

where I_{\max} is a reference scaling factor, chosen here as the maximum of the mean intensities in the two channels. This rescaling ensures that the output is dimensionally consistent and comparable across different operating points.

Expanding the normalised expression in the limit of small noise fluctuations ($\delta \ll 1$), we obtain

$$\begin{aligned}
I_{\Delta}^{\text{norm}}(t) \approx & I_{\max} \frac{\overline{AB} - \overline{CD}}{\overline{AB} + \overline{CD}} \\
& + \frac{I_{\max}}{\overline{AB} + \overline{CD}} \left(\overline{AB} \delta_{\text{RIN},AB}(t) - \overline{CD} \delta_{\text{RIN},CD}(t) + \delta_{\text{SN},AB}(t) - \delta_{\text{SN},CD}(t) \right) \\
& - \frac{I_{\max}(\overline{AB} - \overline{CD})}{(\overline{AB} + \overline{CD})^2} \left(\overline{AB} \delta_{\text{RIN},AB}(t) + \overline{CD} \delta_{\text{RIN},CD}(t) + \delta_{\text{SN},AB}(t) + \delta_{\text{SN},CD}(t) \right).
\end{aligned} \tag{3.36}$$

The first term represents the deterministic differential response, while the remaining terms capture noise contributions. Crucially, the RIN components now appear in two distinct forms:

- A *differential contribution*, weighted by the imbalance between the two halves.
- A *common-mode contribution*, coupled via the denominator of the normalisation.

This enters only if the DC levels are imbalanced.

It is worth noting that in the perfectly balanced case ($\overline{AB} \approx \overline{CD}$, corresponding to the flag centred on the beam), both the raw and normalised difference signals are predominantly limited by shot noise, since the RIN contributions cancel to first order. Normalisation therefore offers no significant advantage at the balance point.

The advantage of normalisation becomes evident when the flag is displaced, such that $\overline{AB} \neq \overline{CD}$. In this regime, the raw differential signal acquires residual RIN terms proportional to the imbalance, leading to an elevated noise floor. The normalised response, on the other hand, suppresses this imbalance-induced RIN by dividing out the total detected intensity. As a result, the RIN floor remains low even when the flag is moved away from the symmetric position.

This property is directly observable in the measured spectra: while the raw difference exhibits a rising RIN background for displaced flag positions, the normalised response maintains a flat, shot-noise-like floor (see Figure 3.11).

3.4.2 Cryogenic Characterisation

We tested the FCI InGaAs Q3000 photodiode and the LED1200L emitter down to cryogenic temperatures of 5 K. These two components constitute the core of the shadow sensor. The prototype shadow sensor, previously characterised in air (see Figure 3.10a), was subsequently integrated into a cryostat to enable operation at liquid helium temperatures (see Figure 3.12).

It is important to test the LED and QPD independently to untangle any changes in the shadow sensor efficiency due to temperature. The setup for these tests is shown in Figure 3.12.




Figure 3.12: Cryogenic setup for component-level testing of the shadow sensor. The image shows the integration of the FCI InGaAs Q3000 photodiode (QPD1 and QPD2 systems) and the LED1200L sources inside the cryostat. Dedicated temperature monitors are installed near both QPDs and the LED system to track thermal conditions during cooldown to 5 K. QPD2 is additionally configured for tests with an LED beam injected from outside the cryostat.

3.4.2.1 LED Characterisation: I-V Curve

We first measured the current–voltage (I–V) characteristics of the Thorlabs LED1200L as a function of temperature to assess its performance under cryogenic conditions. The LED was cooled to a minimum temperature of 28 K, and the measured I–V curves are shown in Figure 3.13. The device exhibited anomalous electrical behaviour at low temperatures, with an irregular I–V response and a significant reduction in optical output power. In several trials the LED failed to turn on reliably, indicating that the device does not operate efficiently when cooled below approximately 200 K. The underlying cause remains to be investigated. In some instances the LEDs ceased emission when cooled below 198 K but recovered upon warming. For reference, KAGRA employs the LED1200E—a sister device to the LED1200L used above—which is packaged in epoxy and lacks the attached lens. The LED1200E exhibits a substantially larger output divergence than the LED1200L, making it a less efficient candidate for this application. To identify a more suitable light

source for cryogenic operation, we also tested the Hamamatsu L13072–0120L LED under the same conditions. The corresponding I–V curves, shown in Figure 3.13, demonstrate a more stable and reproducible behaviour across the temperature range, indicating improved low-temperature performance compared to the Thorlabs device.


3.4.2.2 LED Characterisation: Optoelectronic efficiency

Having established the I – V characteristics of the candidate LEDs, we now quantify their optoelectronic efficiency to inform cryogenic operation. For a shadow sensor operated in a cryostat, the electrical power dissipated in the emitter sets a significant heat load. We therefore define the efficiency


$$\eta = \frac{P_{\text{opt}}}{IV},$$

where P_{opt} is the emitted optical power and IV is the electrical input power. Measuring η as a function of drive current at fixed temperatures allows us to select a drive current that provides sufficient optical power while keeping the thermal load acceptable.

When the LED is mounted in a carrier as part of the shadow-sensor package, additional optical power is lost along the path to the photodiode due to clipping at the lens and the LED carrier aperture. There is also light lost due to flag obstruction and diffraction at the flag. Losses between the LED and lens can be reduced by selecting a lens with a large numerical aperture. However, because the source has a finite size, improving collimation generally requires limiting the numerical aperture, which reduces collection efficiency. This introduces a fundamental trade-off between (i) light captured by the lens and (ii) divergence of the resulting beam, which in turn affects diffraction around the flag and the fraction of light ultimately reaching the photodiode. At room temperature, our preliminary efforts to collimate the beam from the LED1200L LED using an 8 mm aspheric lens indicate an optical loss of $\sim 60\%$ through the packaged LED carrier assembly. To obtain ~ 5 V per quadrant at the readout, the LED is typically driven at $I \approx 2.8$ mA, corresponding to an electrical input power of $P_{\text{elec}} \approx 6.44$ mW (from the measured $P = VI$ characteristic



(a) Thorlabs LED1200L: forward I–V characteristics measured in the cryostat during cool-down. Each curve corresponds to a stabilised temperature (colour-coded, see bar). The LED shows abnormal behaviour and reduced optical output below 100 K.



(b) Hamamatsu L13072–0120L: Forward I–V characteristics measured under identical cryogenic conditions. The LED maintains regular behaviour and higher efficiency down to 18 K.

Figure 3.13: Comparison of forward I–V characteristics for two near-infrared LEDs measured during cool-down in the cryostat. The Thorlabs LED1200L exhibits abnormal behaviour and poor efficiency at low temperatures, whereas the Hamamatsu L13072–0120L performs more reliably and is better suited for cryogenic operation.

of the LED- see Figure 3.13a).

$$P_{\text{elec}} \approx 6.44 \text{ mW}. \quad (3.37)$$


The nominal optical output power of the LED is $P_{\text{opt,LED}} \approx 200 \mu\text{W}$, but measurements indicate only $P_{\text{out}} \approx 75 \mu\text{W}$ is available outside the LED carrier. The remaining electrical power is therefore dissipated predominantly as heat,

$$P_{\text{heat}} \approx P_{\text{elec}} - P_{\text{opt,LED}} \approx 6.44 \text{ mW} - 0.20 \text{ mW} \approx 6.24 \text{ mW}, \quad (3.38)$$


while an additional

$$P_{\text{rad,loss}} \approx P_{\text{opt,LED}} - P_{\text{out}} \approx 200 \mu\text{W} - 75 \mu\text{W} \approx 125 \mu\text{W}. \quad (3.39)$$

of optical power is lost within the carrier (e.g. due to geometric clipping at the lens and carrier aperture. Further losses occur downstream before the remaining light is ultimately measured on the QPD. This implies the total loss is about 6.365 mW before light reaches the flag.



(a) Thorlabs LED1200L. Optoelectronic efficiency $\eta = P_{\text{opt}}/(IV)$ versus forward current at stabilised cryogenic temperatures. Each curve is from simultaneous optical and electrical measurements in the cryostat.



(b) Hamamatsu L13072-0120L. Optoelectronic efficiency $\eta = P_{\text{opt}}/(IV)$ versus forward current at the same set of temperatures, measured with the identical setup and procedure.

Figure 3.14: Comparison of LED optoelectronic efficiency across temperature. Subfigures (a) and (b) show η for the Thorlabs LED1200L and Hamamatsu L13072-0120L, respectively.

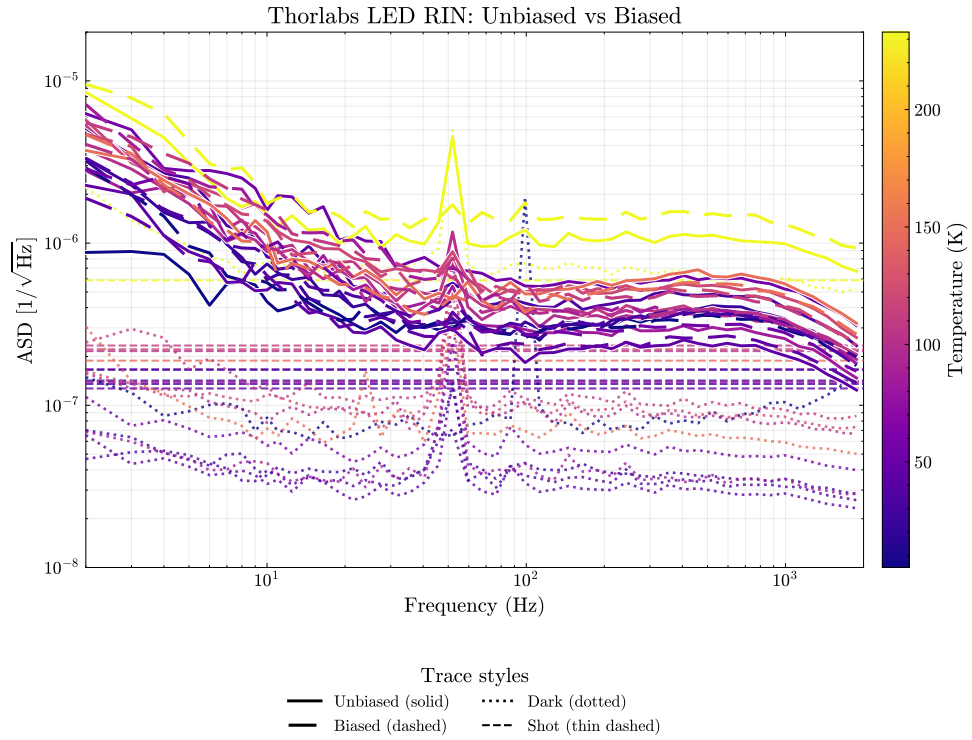
3.4.2.3 LED Characterisation: Relative Intensity Noise

We measured RIN for the Thorlabs LED and Hamamatsu LED. The light output of the Thorlabs LED during this measurement run was low. But the light output of the Hamamatsu LED was enough to produce 20 V signal on the QPD. The RIN of the Hamamatsu LED was better than that of the Thorlabs LED (Figure 3.16). In the figure 3.15b we see that the general trend is RIN increasing with temperature.

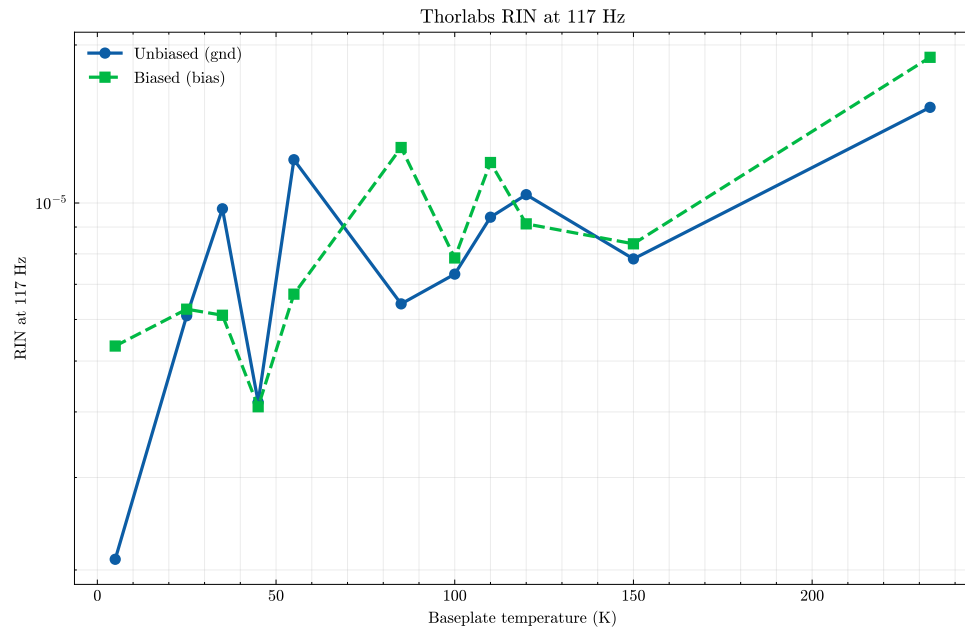
3.4.2.4 LED Characterisation: Reliability Testing

We investigated the performance of four different infrared LEDs: Thorlabs LED1200L and LED1200E, and Hamamatsu L13072-0120L and L13072-0120K. The Thorlabs LEDs with integrated lenses, as well as both Hamamatsu models, failed after the first few thermal cycles (in some cases after the initial run). This degradation is likely related to the packaging, where the LED chip reacts with the surrounding atmosphere within the package.

In contrast, the Thorlabs LED without a lens, which uses an epoxy-sealed package with no internal air, exhibited stable and repeatable behaviour over multiple cooldowns. The electrical impedance of these LEDs decreased from approximately $140\ \Omega$ at room temperature to about $40\ \Omega$ at 5 K.



(a) RIN amplitude spectral density (ASD) of the Thorlabs LED versus frequency. Solid: unbiased; dashed: biased; dotted: dark-noise floor; thin dashed: shot-noise prediction. Shot and dark noise are expressed in the same relative units, normalised to the total optical power incident on the QPD.



(b) RIN ASD at 117 Hz as a function of baseplate temperature for the same dataset. Circles: unbiased (gnd); squares: biased (bias). Values are obtained from the Welch ASD at 117 Hz (log-log interpolation) and share the same normalisation as (a).

Figure 3.15: Thorlabs LED relative intensity noise. (a) Broadband spectra across temperatures; (b) temperature trend at a representative line frequency of 117 Hz. Both panels use the same normalisation to total QPD optical power.




Figure 3.16: Comparison of relative intensity noise (RIN) of the Hamamatsu and Thorlabs LEDs measured under cryogenic conditions. Each curve corresponds to the unbiased (grounded) configuration of the shadow sensor at different baseplate temperatures. The solid lines represent the Hamamatsu LED and the dashed lines represent the Thorlabs LED, with colour indicating the corresponding baseplate temperature. The plot illustrates that the Thorlabs LED 1200L exhibits higher RIN levels across the measured frequency range.

From Figure 3.16 we see that the RIN performance of Hamamatsu LEDs is better than that of the Thorlabs LEDs. But it should also be noted that the Thorlabs LED that was tested in this configuration LED 1200L is a lens packaging and the light output was unreliable at low temperatures. A future testing of other LED 1200L or different packaging without the ball lens (which corresponds to the LED 1200E model) is necessary.

3.4.2.5 QPD Characterisation: Dark noise

We discussed in Section 3.1.1 how dark noise is one of the contributing noise sources in a photodiode. Here we will look at how the dark noise of the QPD changes with temperature. The dark noise of the QPD exhibits a slight bump at high frequencies for low temperatures, but we don't see any trend across temperatures (see Figure 3.17).

QPD Characterisation: Efficiency

In Figure 3.12 QPD2 is mounted close to the window of the cryostat and from outside a 1200 nm LED beam is injected to monitor the change in the response of the QPD with temperature. In Figure 3.18 we plot the change in responsivity of the LED with temperature. The response of the QPD seems to increase with temperature.

QPD Characterisation: Transfer function

The transfer function of QPD2 was measured by performing a swept sine. In the plot in Figure 3.19a we see that the transfer function is flat till 100 Hz and then it starts rolling off. In Figure 3.19b we see how the transfer function of the QPD varies with temperature. This agrees with the QPD efficiency trend with temperature in Figure 3.18.

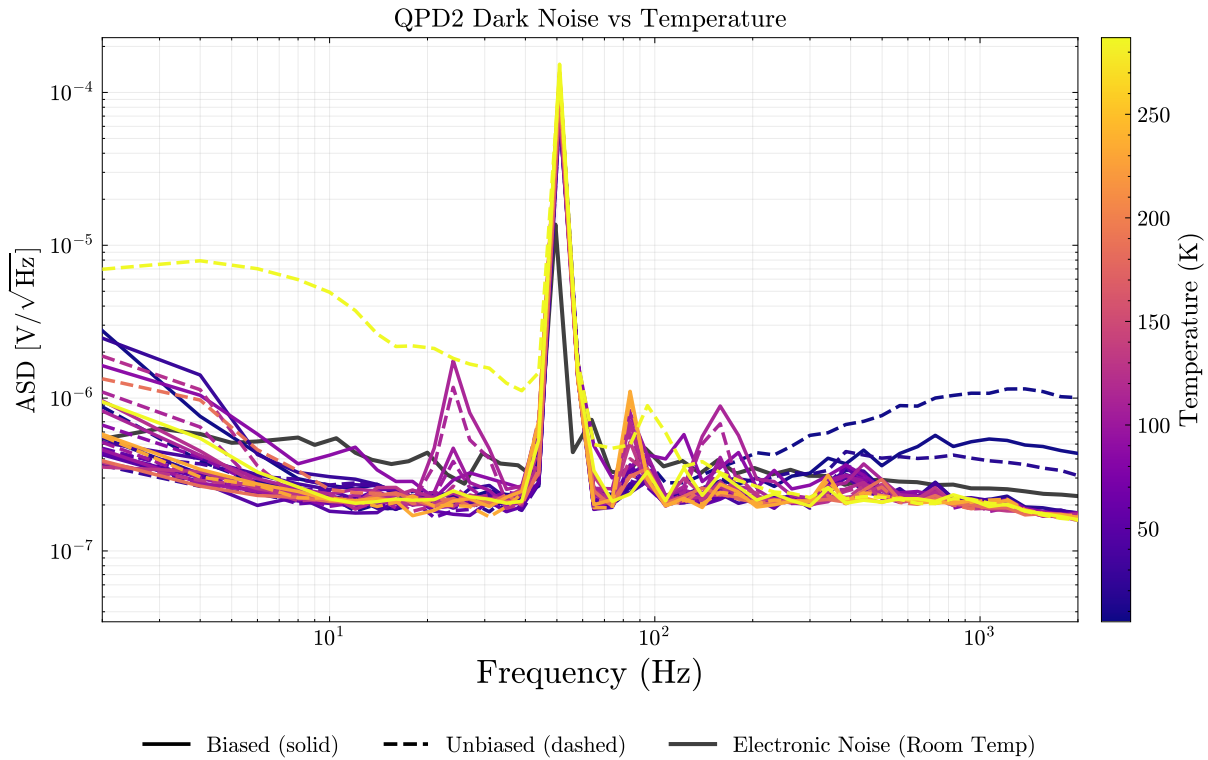


Figure 3.17: Amplitude spectral density (ASD) of the QPD differential dark noise ($AB - CD$) versus frequency for temperatures from 5 to 287 K (colour-coded; plasma map, see bar). Solid lines denote biased operation ($V_{\text{bias}} = 2.3 \text{ V}$); dashed lines denote unbiased. Spectra are estimated with a Welch periodogram (Hanning window) and then down-sampled onto logarithmic frequency bins using the per-bin median to suppress narrow spectral lines, revealing the broadband floor. The grey trace shows a room-temperature electronics-noise reference with the input shorted. The plot highlights the temperature- and bias-dependence of the low-frequency noise, while the broadband floor above $\sim 100 \text{ Hz}$ varies only moderately across temperatures.






Figure 3.18: Temperature dependence of the QPD responsivity. The quantity plotted is $R = V_{\text{QPD}}/(GP_{\text{LED}})$ with $G = 330 \text{ k}\Omega$, i.e. the QPD output voltage normalised by the LED optical power and the transimpedance gain. Each curve is then normalised to its room-temperature point so that $R(295\text{--}300 \text{ K}) = 1$. Two independent measurement runs (different LED powers) are shown; markers denote the measured points and lines guide the eye. The responsivity increases with mount temperature and approaches its room-temperature value near 300 K.



(a) Transfer function of the mounted QPD at 5 K. The response is flat over 10 Hz to 100 Hz; the dashed horizontal line marks the band-mean magnitude. Drive compensation equalises the injected source level.

Figure 3.19: Transfer-function measurements of the QPD mounted next to the cryostat window. (a) Representative 5 K sweep highlighting the flat passband used for characterisation. (continued on next page)



(b) Band-mean transfer magnitude computed over 10 Hz to 100 Hz as a function of baseplate temperature. Each point is the mean of the linear magnitude within the band (averaged across available sweeps at that temperature) and then expressed in dB. Measurements correspond to the cryogenic configuration in Figure 3.12.

Figure 3.19: (b) Temperature dependence of the band-mean magnitude (10 Hz to 100 Hz) from 5 K to 287 K.

3.4.2.6 Discussion

We tested five FCI InGaAs Q3000 quadrant photodiodes (QPDs) down to a base temperature of 5 K. During cooldown, one QPD exhibited a drop in impedance between the cathode and one of the anode legs, necessitating its replacement. In another device, the bond wires connecting the casing pins to the active area broke; however, this was the QPD from which the protective can had been deliberately removed to position the flag closer to the active surface. So mechanical shock from this process might have been the cause.

The remaining two QPDs, along with the replacement unit, operated reliably throughout the cryogenic tests. The overall responsivity of the QPDs decreased with decreasing temperature, while their transfer function remained flat over the frequency range of interest—specifically, in the band where the suspension modes are to be actively damped.

3.4.2.7 Sensor Characterisation

In the preceding subsections, we examined the performance of individual components—specifically LEDs and QPDs—under cryogenic conditions. These studies provided insight into their electrical and optical behaviour, enabling us to identify the most suitable devices for integration into a shadow sensing setup. Building on those results, the next step was to investigate how the complete sensor assembly performs when operated at liquid helium temperatures, particularly in the presence of a moving flag.

The experimental setup used for these measurements is shown in Figure 3.20. The shadow sensor was assembled inside a cryostat. The LED and QPD were mounted within a cylindrical aluminium carrier, which was itself fixed to an aluminium baseplate. Indium wires were employed to enhance thermal conduction between the aluminium baseplate and the cryostat cold plate.




Figure 3.21: Image showing photograph of the flag used in the setup in Figure 3.20




Figure 3.20: Photograph of the cryostat used for helium testing of the shadow sensor. The setup includes the QPD holder, the flag mounted on a translation stage moved via a motion feedthrough, a heater for controlled thermal loading, cold anchors for thermalisation, and temperature sensors placed near the LED and on the cold plate. The flag is a circular aperture that has an aperture diameter of 2.5 mm.

In Figure 3.10b, the shadow sensor was characterised under ambient (in-air) conditions. In Figure 3.22 we present a corresponding measurement performed at liquid-helium temperature using the same optical configuration. The slight irregularity observed between +1 mm +1.5 mm position arises from minor mechanical stiction in the motion feedthrough. The translation stage, positioned outside the cryostat, was coupled to the internal flag

via this feedthrough. It should be noted that the signal normalisation, applied to correct for LED intensity variations, was implemented in this measurement as described in Figure 3.23 and it can be seen that normalisation reduced the LED RIN to shot noise level when the flag is displaced. We can also see that the RIN levels are shot noise limited.




Figure 3.22: Cryogenic shadow-sensing response of the QPD (flag diameter 2.5 mm). The black curve shows the raw differential voltage and the red curve is normalised by the total sum signal to compensate LED intensity changes. Dashed lines indicate local linear fits (tangents) to the normalised curve: at the centre the responsivity is -4.085 V/mm, and at -1.00 mm it is -15.179 V/mm.




Figure 3.23: Relative intensity noise (RIN) spectra measured at 47 K for the shadow sensor. The orange and green traces correspond to the raw and normalised responses when the flag is displaced from the centre, respectively. The normalisation effectively suppresses LED intensity fluctuations, reducing the low-frequency excess noise towards the shot-noise limit (dashed lines). The blue curve shows the RIN with the flag centred, while the black trace represents the dark noise floor of the readout electronics. The equivalent displacement noise can be read from the right-hand axis where the voltage spectral density has been scaled using the measured responsivity of the sensor when the flag is centred. This gives a worse case upper limit for the displacement noise floor as the responsivity is lowest when the flag is centred.

3.5 Package Design

The original BOSEM package details can be found in [41]. This design cannot be used directly in a cryogenic environment because the material choices and thermal interfaces are not optimised for low-temperature operation, where high thermal conductivity is required to manage local heat loads. In addition, cooldown tests of BOSEM hardware to 77 K have reported stress concentrations at specific locations in the package [171]. Given the limited electrical-to-optical efficiency of the LED and the associated heat dissipation, it is also

desirable to select materials and interfaces that provide an effective thermal path from the LED package to the reaction cage. Besides modifications in dimensions to accommodate the new LEDs and QPD size, the BOSEM sensor package for this work therefore differs primarily in its choice of materials and its approach to thermal management.

We studied the BOSEM packaging to identify which elements could be modified with minimal disruption while meeting cryogenic requirements. In Figures 3.25 and 3.24, we see the photodiode (QPD) and LED carrier assemblies from the BOSEM design. In both carriers, the optoelectronic component is fitted in a carrier and held together with the lens using a sleeve. In the original BOSEM, the IRLED and photodiode sleeves are ceramic in order to electrically isolate the devices from the surrounding aluminium structure while still providing a mechanically robust locating interface.




Figure 3.24: IRLED carrier assembly. Left: cross-sectional view of the IRLED carrier. Right: exploded view showing the constituent parts.





Figure 3.25: Photodiode carrier assembly. Left: cross-sectional view of the photodiode carrier. Right: exploded view showing the constituent parts.

The modified BOSEM package for GOSEM prototype design is shown in Figure 3.26. The SolidWorks modelling was carried out with the assistance of *Andy Pevchikh* (summer student, 2025). We will discuss the appropriate choice of materials for the parts in detail in the following subsections.



(a) LED subassembly (exploded view). From top to bottom: retainer ring, LED (LED1200L), LED sleeve, an 8 mm-mounted aspheric lens, and the carrier. The carrier is split into two halves to minimise clamping-induced stress on the lens (reducing the risk of fracture) and to simplify integration of alternative lenses, enabling a modular optical package.



(b) QPD subassembly (exploded view). Components shown are the retainer ring, QPD sleeve, QPD, and carrier.

Figure 3.26: Exploded CAD views of the (left) LED carrier assembly and (right) QPD carrier assembly.

3.5.1 Cryogenic design drivers


The cryogenic package design is guided by the following considerations:

- *Thermal conduction:* the LED heat load should be efficiently conducted away to the reaction cage to remain compatible with the cryostat thermal budget.
- *Electrical isolation:* the LED and QPD packages must remain electrically isolated from the sensor body to meet isolation requirements, without compromising thermal performance more than necessary.
- *Differential contraction and stress:* material combinations and interfaces should minimise thermally induced stress during cooldown, consistent with observations of stress in cryogenic BOSEM cooldown tests [171].

- *Cooldown overhead*: the specific heat of the package materials contributes to cool-down time and operational cost; this motivates minimising unnecessary thermal mass while maintaining adequate thermal conduction.



(a) 3D print of the GOSEM LED and QPD subassemblies fitted into a mock-up of the BOSEM main body.



(b) Exploded view of GOSEM assembly in CAD.

Figure 3.27: GOSEM assembly.

3.5.2 Material selection and insulation strategy

For cryogenic operation, high-purity copper is an attractive candidate for the main carrier and surrounding body because of its excellent thermal conductivity at low temperatures. An aluminium alloy (e.g. Al 6063-T5) provides an alternative with good machinability and relatively low mass, but typically offers reduced thermal conductivity compared to high-purity copper under cryogenic conditions. Fasteners and small hardware can be realised in brass or bronze to provide good thermal conduction and a coefficient of thermal contrac-

tion broadly compatible with copper, while maintaining practical mechanical properties. For the moving shadow mask/flag, a material with sufficiently high heat capacity can be advantageous for thermal stability under illumination, provided that the added thermal mass remains acceptable for cooldown.

The principal packaging change concerns the electrically isolating sleeve. The ceramic sleeve used in the original BOSEM provides reliable electrical isolation but offers comparatively poor thermal conduction. For the cryogenic package, the intent is to replace this approach with an isolation scheme that preserves alignment features while improving the thermal path from the LED/QPD into the carrier body. One practical route is to use a metallic sleeve/body for heat sinking, combined with a thin electrical insulating layer (as noted in earlier design records, e.g. cigarette paper as a compliant electrical insulator) between the optoelectronic package and the carrier. This maintains electrical isolation while avoiding a thick, low-conductivity ceramic element, and also provides compliance that can help accommodate differential thermal contraction. An insulating epoxy coating may also be applied to a copper sleeve to provide the required electrical isolation while maintaining an efficient thermal path for conducting heat away from the LED. One candidate material for such a coating is EP21TDC-2LOND, which is formulated to combine electrical insulation with enhanced thermal conduction.

3.5.3 Implementation in the cryogenic package

In the proposed cryogenic carrier concept, the BOSEM alignment features (machined flat and pin-hole locator) are retained to preserve assembly repeatability, while the surrounding structure is redesigned to:

- maximise metal-to-metal contact area between the carrier and the reaction cage (or an intermediate thermal link),
- minimise thermal contact resistance at interfaces (surface finish, contact pressure, and fastener placement),

- reduce stress concentrations by avoiding over-constraint of parts that contract differently on cooldown.

This package strategy is intended to provide a straightforward integration route with existing suspension hardware, while improving thermal performance and reducing the likelihood of thermally induced stress during cooldown.

3.6 Notes and Future Development

The primary goal of this work was to characterise and evaluate the performance of light-emitting diodes (LEDs), quadrant photodiodes (QPDs), and flag geometries toward the development of a high-sensitivity, low-noise, and wide-dynamic-range cryogenic shadow sensor with minimal thermal load on the cryostat. The experiments carried out in this study have demonstrated the feasibility of operating such a sensor at liquid-helium temperatures, and have provided detailed insight into the thermal, optical, and electronic behaviour of its constituent components.

The core components identified as suitable for cryogenic operation are the Thorlabs LED1200E emitter and the FCI InGaAs Q3000 QPD. The LED exhibited consistent optical output and stable impedance characteristics through repeated cooldowns, with its resistance decreasing from approximately $140\ \Omega$ at room temperature to $40\ \Omega$ at 5 K. The QPDs performed reliably across multiple thermal cycles, maintaining a flat transfer function in the frequency band relevant for suspension damping, though with a modest decrease in responsivity at low temperatures. Additionally, comparative tests revealed that **Hamamatsu LEDs** exhibited improved relative intensity noise (RIN) performance but suffered from packaging-related reliability issues during repeated cryogenic operation. These findings highlight the importance of packaging design in determining device longevity and performance at low temperatures. We did not test LED1200E for its RIN performance and this is part of future efforts.

Analytical and numerical modelling of different flag geometries showed that the circular aperture flag offers a larger dynamic range and better linearity compared to the conventional rectangular flag, at the expense of slightly higher shot-noise floor due to low responsivity when the flag is centred. The rectangular flag, in contrast, provides a steeper central slope and higher small-signal responsivity but a reduced dynamic range. The optimal operating point for each geometry was identified from the simulated thermal maps, enabling a rational design trade-off between sensitivity, dynamic range, and heat load.

The in-air and cryogenic measurements of the prototype GOSEM confirmed the trends predicted by the models. The normalised QPD response successfully suppressed LED intensity noise to near shot-noise levels, validating the effectiveness of the normalisation scheme. The cryogenic tests at 5 K demonstrated that the sensor can achieve stable operation with measurable displacement sensitivity within the expected noise limits. Minor non-linearities in the response curves were attributed to mechanical stiction in the motion feedthrough and can be mitigated in future iterations through improved actuation and alignment mechanisms.

Looking ahead, several avenues of development are recommended:

- **LED packaging improvements:** A key limitation observed in this work was the degradation of certain LED packages under cryogenic cycling. Future studies should focus on hermetically sealed or vacuum-encapsulated LED designs to prevent interaction between the semiconductor chip and residual gases or moisture in the package. Future work may include removing the casing so that a lens can be placed close to the LED chip to enable better light collection and collimation.
- **Closed-loop operation:** Future prototypes should integrate the cryogenic shadow sensor into an active control loop for suspended masses, enabling direct assessment of displacement noise performance and feedback stability under cryogenic conditions.

- **Increased Sensitivity:** From Figure 3.23 we see that the sensor noise is shot noise limited. The sensitivity of the sensor can be improved by using a larger drive current for the LED and increasing the optical output. This will increase the responsivity and lower shot noise SNR and dark noise SNR. This should also be accompanied with a rigorous study of heat extraction because with larger drive currents comes higher heat load.

Finally, while this work focused on shadow sensing as the displacement readout scheme, it is worth noting that alternative imaging-based sensors could, in principle, capture more of the emitted light and offer spatially richer diagnostics. The Adelaide group have been working on QOSEMs which use a lens as a flag with a QPD to sense displacement of suspensions. This is however yet to be optimised for cryogenic purposes. Such approaches typically suffer from a reduced dynamic range and increased readout complexity at cryogenic temperatures. The shadow-sensing approach developed here thus remains a robust and low-heat-load solution, well-suited for implementation in future generations of cryogenic suspension systems such as those envisioned for the Glasgow Cryogenic Interferometer Facility (GCIF) and next-generation gravitational-wave detectors.

Chapter 4

A+ LIGO: Output Mode Cleaner Cavities

Over the course of a year at Caltech, I was part of a team that led the efforts to refine assembly procedures, developed optical characterisation and tighter screening of optical components to build Output Mode Cleaner (OMC) cavities with efficiency higher than 98% for O5 run. These enhanced OMCs will be installed during the A+ upgrade. In this chapter, we discuss efforts to build the first OMC cavity and the characterisation efforts of the optical components at the beginning of the build process.

The cavity design for A+ OMCs was developed by the aLIGO OMC team, which is same as that for the aLIGO OMC cavities. I made major contributions to the characterisation of all the components described in this thesis, and to the assembly and characterisation of the first cavity. The mirror scatter measurements and associated plots were produced by Dr Zhang at the LIGO Laboratory, Caltech. The piezo characterisation plots and transverse mode spacing measurement plots were produced by Dr Nakano at the LIGO Laboratory, Caltech.

4.1 Introduction

The pursuit of ever-greater sensitivity in gravitational wave detectors has driven a continuous evolution of interferometric readout schemes. As part of the A+ upgrade to Advanced LIGO (aLIGO), a key development is the transition from DC readout to balanced homodyne (BHD) readout, which promises to unlock access to quantum-noise-limited performance across a wider frequency band. Central to this transition is the deployment of enhanced Output Mode Cleaners (OMCs), which serve to filter unwanted spatial modes and radio-frequency sidebands from the interferometer’s output beam, transmitting only the fundamental TEM_{00} carrier mode with minimal loss. More details on the importance of OMCs in LIGO can be found here [71]. The OMC is a bowtie shaped monolithic optical cavity of about 1.1 m roundtrip length. In the current generation GW detectors an OMC is placed between the antisymmetric output port and DC readout photodiodes that sense DARM (Differential Arm signal) as seen Figure 4.2 [53].

In this chapter we will document the design philosophy, fabrication process, and experimental implementation of these next-generation OMCs. The cavity design is very similar to the aLIGO OMCs. The goal is to exceed the performance of existing aLIGO OMCs, whose throughput typically falls below 97%, in order to support squeezing levels of 6 dB and beyond, a target that necessitates OMC throughput above 98% [55]. Achieving this level of performance requires an enhanced approach: tighter screening of optical components, improved coatings with reduced transmission and scattering losses, more precise cavity assembly and alignment and better cleanliness standards. These efforts were made in constant consultation with the previous OMC team and their feedback was very valuable. We start with revisiting the DC Readout scheme that is employed in the existing LIGO detectors to detect GW waves [73]. In this chapter we highlight the downsides of DC readout in Section 4.1.1 and justify the necessity for BHD readout which is described in detail in Section 5.3.

4.1.1 Transition from DC Readout to BHD

A passing gravitational wave induces a differential change in the light travel time along the interferometer arms, from the beam splitter (BS) to the end mirrors (ETMX and ETMY) and back. This perturbation alters the interference condition at the BS, resulting in a portion of the light exiting through the antisymmetric (AS) port. However, operating the interferometer exactly at the dark fringe causes the readout power to depend only on the second order of the gravitational wave field amplitude. Whereas operating the interferometer with a slight offset from the dark fringe setting introduces a DC static field at the output port. The presence of this leakage field gives rise to a linear response to the GW field amplitude.

$$P_{AS} = |E_{DC}|^2 + E_{DC}^* E_{GW} + E_{DC} E_{GW}^* = |E_{DC}|^2 + 2\mathcal{R}e(E_{DC} E_{GW}^*), \quad (4.1)$$

where P_{AS} is the power output at the anti-symmetric port, E_{DC} is the leakage field at the AS port from the offset operating point of the detector, and E_{GW} is the GW signal field amplitude. But this reliance on static carrier field at the AS port means the intensity noise in LO field becomes indistinguishable from the GW signal as seen in the equation below. If we substitute $E_{DC} = (1 + \epsilon)\bar{E}_{DC}$ in equation 4.1 we get

$$P_{AS} = |E_{DC}|^2 + 2\mathcal{R}e(\bar{E}_{DC}(E_{GW} + \epsilon\bar{E}_{DC})^*), \quad (4.2)$$

ignoring second order terms in ϵ .

This leads to LO intensity noise coupling into the GW signal [159] which is a limitation of DC readout.

One other limitation of the DC readout scheme is the limited control of the phase angle between the signal field E_{GW} and the LO field E_{DC} . As seen in equation 4.1, LO and GW fields are in quadrature but some freedom in relative phase between the two can be introduced by making the circulating field amplitudes in the two fields asymmetric [73]. But the amount of phase angle range achievable from this is small and less than desired. This phase angle is called the homodyne angle.

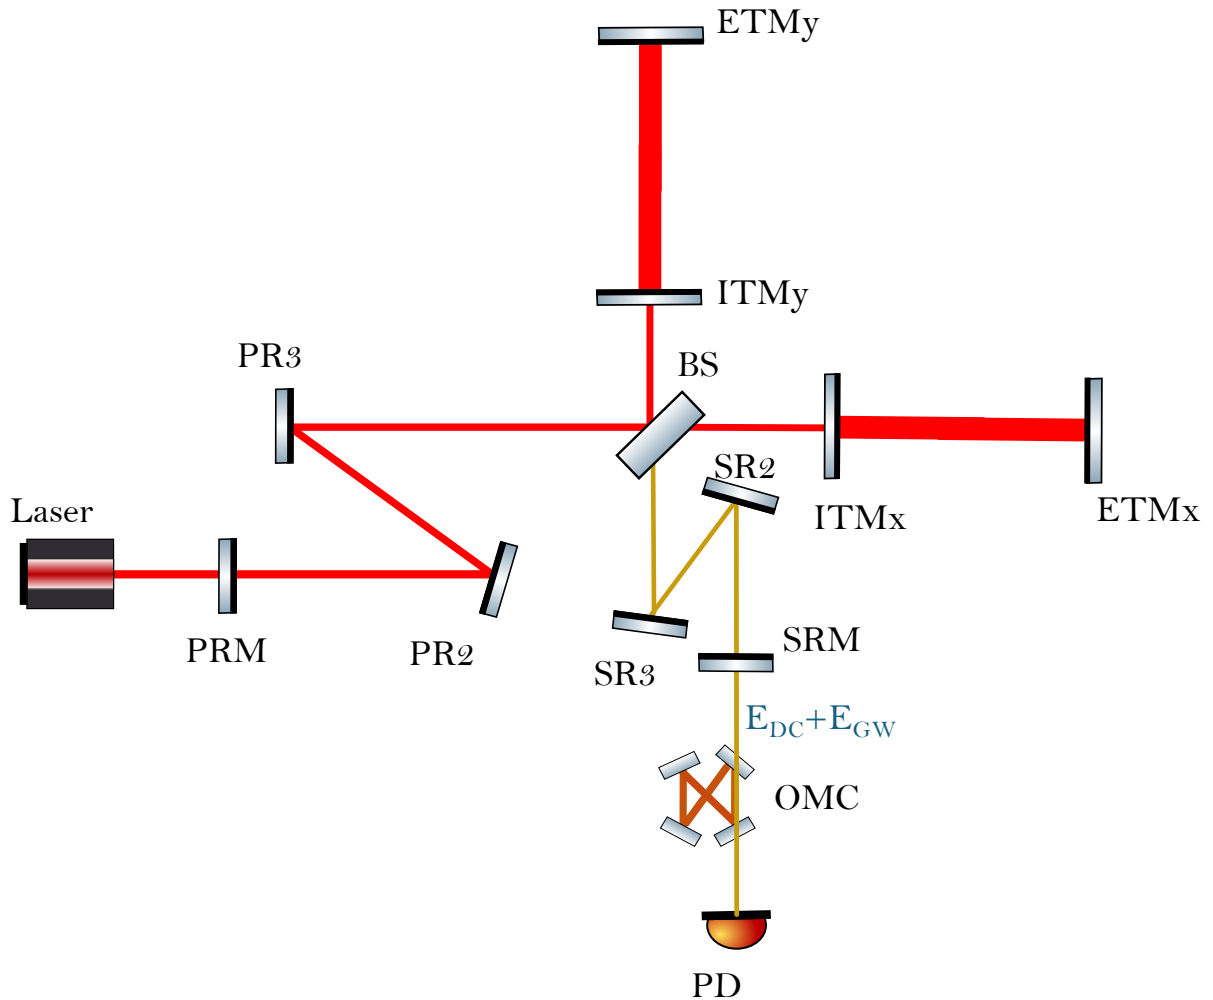


Figure 4.1: Simplified schematic of the Advanced LIGO interferometer operated with DC readout. A pre-stabilised laser is injected into the power-recycling cavity, formed by the power-recycling mirror (PRM) together with the input test masses (ITMs) and the recycling optics (PR2, PR3), to build up carrier power at the beamsplitter. The arms are Fabry–Pérot cavities formed by the ITMs and end test masses (ETMs). The differential arm-length signal exits at the antisymmetric port, where a small intentional offset provides carrier light as a local oscillator for DC readout. The output field is shaped by the signal-recycling cavity (SRM with SR2, SR3 and the ITMs), spatially filtered by the output mode cleaner (OMC), and detected at the photodiode (PD) [116, 104].

Some of the drawbacks of the DC readout scheme are carrier higher order modes and RF modulation sidebands. These do not contribute to the signal but increase the signal's shot noise contribution, thus decreasing the signal to noise ratio. LIGO uses Output Mode Cleaner, a short optical cavity of moderate finesse of about 400, to clean the field out of the antisymmetric port spatially and also filter out the RF sidebands that are used to control various cavities [71].

Balanced Homodyne Detection (BHD) is a readout scheme that is going to be employed in the next upgrade to LIGO, A+. Unlike the DC Readout scheme, in BHD the LO field does not co-propagate with the GW field. In this scheme, LO is picked off and interfered with the signal from the anti-symmetric port (Figure 4.2).

Here,

$$PD_1 - PD_2 = 2 \mathcal{R}e \left(e^{i\delta} \bar{E}_{LO} E_{GW}^* \right) \quad (4.3)$$

The presence of the complex phase $e^{i\delta}$ highlights the independence of LO field phase from GW field phase [87]. Also, the output readout is, in principle, immune to intensity fluctuations in LO field. For more details on challenges on BHD readout scheme refer [73]. In the context of measuring squeezed states of light, several technical noise aspects associated with balanced homodyne readout (BHR) have been extensively studied in quantum optics [158][115].

The LIGO detectors use frequency dependent squeezing to reach their design sensitivity. To effectively utilise squeezed states in gravitational-wave detectors, it is essential to maintain precise control over their phase, that is, their orientation relative to the readout quadrature [114] - the homodyne angle. Limited range of the homodyne angle is a major limitation of the DC readout scheme. BHD enables independent control of the homodyne angle using a phase shifter in the LO pick off path as shown in Figure 4.2. Also, the readout is in principle free of LO noise since this is, in principle, rejected as a

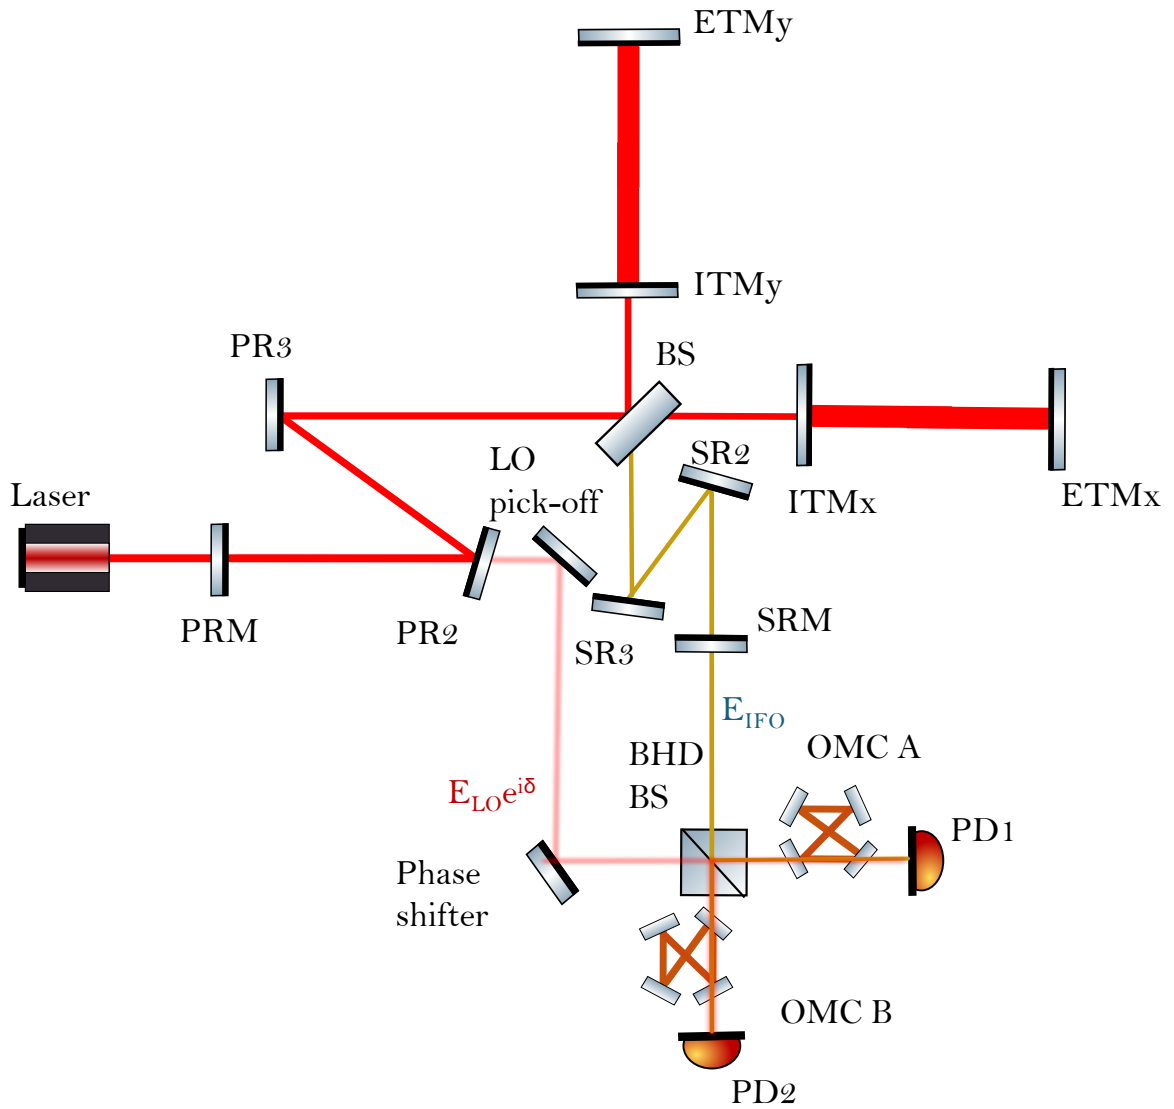


Figure 4.2: Standard balanced homodyne detection (BHD) configuration in the A+ interferometric gravitational-wave detector: This schematic we see how the local oscillator (LO) and interferometer output fields are combined at the BHD beam splitter. No vacuum ports are open, ensuring minimal quantum noise contribution. The output beams pass through orthogonal Output Mode Cleaners (OMC A and OMC B) and are detected at photodiodes PD1 and PD2. The phase shifter, which is basically a mirror, allows control over the quadrature being measured.

common mode. Unlike DC readout, since LO and GW fields are not co-propagating, BHD requires use of two OMC cavities to ensure ideal spatial overlap between the LO and GW fields, which is a crucial requirement for squeezed light technology in LIGO so that the laser modes in both fields overlap.

4.1.2 Requirements and Impact

Achieving the quantum noise reduction goals of the A+ upgrade requires a significant improvement in squeezed vacuum injection and the optical efficiency of the detection chain. While current squeezing levels have reached up to 5.2 dB at LIGO Hanford and 6.1 dB at LIGO Livingston, maintaining stable operation at these levels remains a technical challenge. To reach the A+ design target of -6 dB of quantum noise suppression [55], it is essential to improve mode-matching, implement active alignment control, and reduce optical losses, particularly in the Output Mode Cleaners (OMCs). The OMCs must exhibit a transmission efficiency greater than 98% to support this level of squeezing, surpassing the ~ 96 – 98% efficiencies seen in current aLIGO systems. Meeting these stringent throughput requirements demands improved coatings, careful fabrication and characterisation, and better cleanliness standards of the OMC cavities and associated optics. This will be discussed in detail in rest of the chapter.

The impact of A+ can be expressed either as a broadband reduction in the strain amplitude spectral density or as an increase in astrophysical range. For example, the A+ design model reports a binary-neutron-star (BNS) range of ~ 325 Mpa, compared with ~ 173 Mpa for the updated Advanced LIGO design curve, corresponding to a factor ~ 1.9 increase in BNS range (and hence an order-of-magnitude increase in surveyed volume) [117]. While the detailed commissioning trajectory depends on as-built performance and run planning, this factor-of-two-class improvement captures the intended scale of the upgrade.

4.2 Cavity Design




Figure 4.3: CAD of A+ Output Mode Cleaner Cavity: This figure highlights the bowtie geometry of the cavity. The beam enters through flat mirror 1 and exits through flat mirror 2. The grey surrounding the cavity is the bonding template used to guide the optics bonding process. The cavity consists of two flat mirrors and two curved mirrors. The curved mirror sub-assemblies are held against a couple micrometers to aid cavity alignment [94].

The OMC is a rigid optical cavity constructed in a bow-tie configuration with four mirrors to minimize retro-reflection (Figure 4.3). The even number of mirrors help make the higher order mode spectrum simpler. It is formed of two curved mirrors and two flat mirrors. The mirrors are attached to a glass breadboard that also holds the photodiodes for the DC readout of the interferometer, ensuring beam stability. This makes the cavity monolithic. The transmitted light from the OMC is split equally between two photodiodes, which allows for characterisation of correlated and uncorrelated noises, higher transmitted power, and redundancy. The OMC assembly, which includes the glass breadboard, the cavity, and the DCPDs, is suspended on the Balanced Homodyne Readout Platform Suspension (BHSS)[94] for vibration isolation and DC/slow control of its position.

Further details on the design of OMC cavities can be found in the LIGO technical document [34]. The optical testing document [32] describes all optical tests performed before the installation of the aLIGO OMCs and is used as a reference during the build process for the A+ OMCs. The A+ OMC cavities are very close in design to the current OMCs but the OMC breadboard went through some major changes and are documented in the LIGO document [162].

4.3 Testing of Cavity Optical Components

Thorough characterisation of the optical components is essential to enforce strict selection criteria and ensure that only components meeting the required specifications are used in cavity assembly. This section describes the optical tests performed as part of this qualification process. The components tested include prism flat mirrors and the components used for the curved mirror subassembly units—comprising the curved mirrors, piezoelectric ring actuators, and tombstone-style mirror mounts.

4.3.1 Characterisation of Prism Mirrors

Prism mirrors are the flat mirrors that along with two curved mirrors form an OMC cavity in a bowtie configuration. Optics bonded on the glass breadboard all have some deviation from perpendicularity but there's no pitch freedom involved once the optics are placed in the breadboard. So it's essential to measure deviation from perpendicularity and wedging of all the optics and ensure they're within the specification values. Wedging here refers to deviation of the front and back surfaces of the mirrors from parallel-ness. In section 4.3.2.2 we measured the curvature minimum position of the curved mirrors, this is essentially the wedge angle of the curved mirrors. Although the specification values on horizontal wedging is more relaxed, they are measured and ensured they're within the specification values.

4.3.1.1 Wedge angle

The prism mirrors are inclined at an angle of 0.5 degrees (equivalent to 30 arcmin). A deviation in the wedge angle from the prescribed specifications can result in unintended beam deflection making cavity alignment harder. To ascertain the correctness of the wedge angle, measurements were taken of the angle between the incident and reflected beams. An autocollimator was used to perform this measurement as described in detail in the [46]. Summary of the measurements can be found in Table 4.1.

Prt. No.	Horizontal wedge (deg)	Conform?
1	0.4788	Yes
2	0.5061	No
4	0.4856	Yes
5	0.4309	Yes
6	0.4993	Yes
7	0.4925	Yes
9	0.4651	Yes
10	0.4720	Yes
11	0.4856	Yes
12	0.4651	Yes
13	0.4651	Yes
14	0.4720	Yes
15	0.4788	Yes
16	0.4651	Yes
17	0.4788	Yes
22	0.4720	Yes
24	0.4720	Yes
26	0.4583	Yes

Table 4.1: Measured wedge angles of prism mirrors. Mirrors with wedge angles exceeding 0.5° fall outside specification and are excluded from cavity construction.

Prt. No.	Deviation from perpendicularity (arcsec)	Direction of Tilt	Conforming?
1	0	perpendicular	Yes
2	2.8	obtuse	Yes
4	8.4	obtuse	Yes
5	9.9	obtuse	Yes
6	27.0	obtuse	Yes
7	12.8	acute	Yes
9	6.8	obtuse	Yes
10	6.0	obtuse	Yes
11	5.7	obtuse	Yes
12	12.8	obtuse	Yes
13	31.5	obtuse	No
14	19.2	obtuse	Yes
15	39.2	obtuse	No
16	16.5	obtuse	Yes
17	6.3	acute	Yes
22	57.0	acute	No
24	36.8	acute	No
26	10.7	acute	Yes

Table 4.2: Prism mirror deviation from perpendicularity, as measured using a white light autocollimator. Perpendicularity better than 30 arcsec is required to minimise beam displacement due to the lack of internal pitch adjustment in the OMC cavity assembly. Deviations beyond this limit are marked non-conforming.

4.3.1.2 Perpendicularity

The pitch alignment of the OMC cavity optics lacks internal adjustment as the two flat mirrors and curved mirror sub-assemblies are directly affixed to the breadboard. A misalignment of 10 arcsec (equivalent to 2.8 rad) results in an approximate beam shift of 0.1 mm upon traversing a cavity length roundtrip. Consequently, it is essential to ensure excellent perpendicularity between the front and bottom surfaces of the prisms. To meet this criterion, a requirement of better than 30 arcsec for perpendicularity was set and should be verified. This was checked using an autocollimator (Table 4.2).

4.3.2 Characterisation of the Curved Mirror Subassemblies

The OMC is a bowtie cavity with two curved mirrors and two flat mirrors (Figure 4.3). In the previous subsections we looked at the characterisation efforts for the flat mirrors or prism mirrors. In this section we will look at the efforts to characterise the curved mirror sub-assembly. This comprises of the curved mirror, a piezoelectric transducer (PZT) and a mounting tombstone mirror (Figure 4.4). Each cavity has two curved mirror sub-assemblies, thus two PZTs. One of the PZTs is used for cavity length adjustment and the other as a shutter to prevent excessive power from entering the cavity in the case of an accident.




Figure 4.4: Components of the curved mirror subassembly: curved optic, piezoelectric transducer (PZT), and tombstone mounting prism.

4.3.2.1 Tombstone Mounting Mirrors

Just like the perpendicularity of the HR prisms was tested (section 4.3.1) the same setup was used to test the perpendicularity of mounting prisms. These prisms hold the PZT and the curved optic onto the glass breadboard as shown in the image in Figure 4.4.

Prt. No.	deviation in arcsec	Direction of Tilt	Conforming?
40	4	Obtuse	Yes
41	6	Obtuse	Yes
42	0	Perpendicular	Yes
43	26	Obtuse	Yes
44	8	Obtuse	Yes
45	19	Obtuse	Yes
46	4	Obtuse	Yes
47	9	Obtuse	Yes
48	8	Obtuse	Yes
49	8	Acute	Yes
50	0	Perpendicular	Yes
51	13	Acute	Yes
52	30	Obtuse	No
53	11	Obtuse	Yes
54	4	Acute	Yes
55	4	Obtuse	Yes
56	7	Obtuse	Yes
57	0	Perpendicular	Yes

Table 4.3: Tombstone mirror deviation from perpendicularity, as measured using a white-light autocollimator. Perpendicularity better than 30 arcsec is required to minimise beam displacement due to the lack of internal pitch adjustment in the OMC cavity assembly. Deviations beyond this limit are marked non-conforming.

4.3.2.2 Curved Mirrors

Radius of curvature

The OMC's curved mirrors were designed with a targeted radius of curvature of 2.5 m, but a measurement was required to verify this value. We used a beam profiler (BP209-VIS) and laser beam picked off from that input to the OMC test setup to perform the measurement. The picked off beam was passed through a pair of lenses to expand the beam. The

input beam projected onto the curved optic was characterised and then gaussian beam propagation formula was used to fit the beam profile at the focus of the beam reflected off of the curved optic to calculate the radius of curvature. The results of this measurement are tabulated in Table 4.5.

Mirror Prt. No.	ROC p-plane (m)	ROC s-plane (m)
2	2.507 \pm 0.004	2.528 \pm 0.002
4	2.506 \pm 0.002	2.527 \pm 0.001
5	2.504 \pm 0.003	2.530 \pm 0.002
6	2.509 \pm 0.002	2.525 \pm 0.001
7	2.504 \pm 0.005	2.534 \pm 0.002
11	2.504 \pm 0.006	2.529 \pm 0.002
12	2.529 \pm 0.006	2.529 \pm 0.01
14	2.512 \pm 0.006	2.534 \pm 0.003
16	2.500 \pm 0.002	2.524 \pm 0.001
17	2.511 \pm 0.002	2.527 \pm 0.002
18	2.507 \pm 0.002	2.526 \pm 0.002
19	2.511 \pm 0.002	2.534 \pm 0.002
20	2.512 \pm 0.003	2.526 \pm 0.001
21	2.511 \pm 0.003	2.539 \pm 0.003
22	2.499 \pm 0.007	2.536 \pm 0.003
23	2.505 \pm 0.002	2.528 \pm 0.002
25	2.466 \pm 0.024	2.520 \pm 0.004
26	2.495 \pm 0.004	2.537 \pm 0.006
29	2.507 \pm 0.003	2.528 \pm 0.002
30	2.509 \pm 0.002	2.520 \pm 0.002

Table 4.4: Radius Of Curvature (ROC) of the 20 A+ OMC curved optics.

Curvature minimum or Wedge angle

The centre of curvature of the curved mirrors is not always at its geometric centre and this causes wedging when the mirrors are glued onto a flat surface. This is particularly important to consider when the mirror is affixed to a mounting prism, using a piezoelectric transducer to attach the back surface of the mirror. Since it is not possible to adjust the mirror's pitch angle after it has been attached to the prism, it is necessary to characterise the location of the curvature minimum in order to align any misalignment along the horizontal axis. The relationship between displacement of the curvature minimum and the wedging can be written as

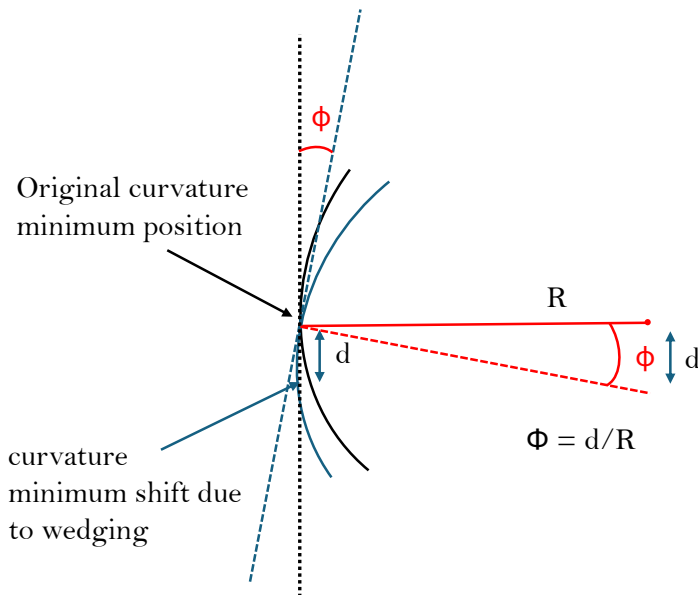


Figure 4.5: Shift in the position of the curvature minimum of the curved optic manifests as wedging of the curved optic.

$$\phi = \frac{d}{R} * \frac{180}{\pi}, \quad (4.4)$$

where ϕ represents wedge angle in degrees, d and R represent the distance of the curvature minimum from the centre of the optic and radius of curvature of the optic respectively (Figure 4.5).

We used a ZYGO Fizeau Interferometer setup to characterise the curvature minimum of the optics. The Zygo setup consisted of a flat transmission sphere with the 0.5" curved mirror mounted against a 1" flat mirror. A folded mirror mounted at a 45-degree angle was utilised to direct the Zygo beam downwards onto the table. A 3-inch flat was chosen as the reference surface, which was placed on a lens tissue parallel to the table surface. Prior to use, the reference flat was inspected with a green flashlight to ensure that no dust particles were present on the surface. Any dust was removed by blowing air using a top gun. If necessary, any remaining dust was removed using an IPA swab.




Figure 4.6: ZYGO curvature minimum measurement setup in the Downs Lauriston Labs, Caltech.

Similar cleaning and inspection methods were applied to the back surface of the curved mirror. The curved mirror was then placed on the reference flat with the fiducial marker positioned at 12:00, which was defined as the direction pointing towards the Zygo instrument. The fold mirror was adjusted in pitch and yaw until the fringes on the 3-inch reference flat were nullified. This setup allowed for a larger surface area of the reference flat to be visible. Once the fringes were nullified, the curved mirror was picked up and replaced several times to ensure that the fringe pattern remained reproducible. The data was then collected with the fiducial marker at the 12:00 position. This process was repeated at 3:00, 6:00, 9:00, and again at 12:00. The results obtained from this experimental setup were reproducible, and it was subsequently used to measure the surface profile of all the curved mirrors. Measurement of one of the mirrors can be found in Figure 4.7. Table 4.5 summarises the curvature minimum measurement data.




Figure 4.7: In the left graph we see the position vectors of curvature minimum for various clockings of the fiducial. In the graph on the right the curvature minimum positions after de-clocking is shown i.e. projecting the position vectors of various clockings to 12 o'clock. Red - 12 o'clock, green - 3 o'clock, blue - 6 o'clock, black - 9 o'clock.

Mirror Serial #	d (mm)	ϕ (deg)	Wedge angle (arcmin)
2	1.394 ± 0.06	311.015	1.903
4	0.848 ± 0.041	141.296	1.159
5	0.792 ± 0.043	112.812	1.082
6	1.266 ± 0.022	18.443	1.729
7	0.838 ± 0.058	300.549	1.144
11	0.072 ± 0.036	37.842	0.098
12	1.313 ± 0.047	269.077	1.785
14	1.104 ± 0.055	9.329	1.504
16	1.175 ± 0.057	129.497	1.608
17	0.917 ± 0.02	160.43	1.251
18	0.766 ± 0.011	257.557	1.047
19	0.594 ± 0.038	160.04	0.81
20	0.951 ± 0.031	257.958	1.298
21	0.782 ± 0.012	126.481	1.065
22	1.103 ± 0.034	162.341	1.506
23	0.916 ± 0.06	25.164	1.252
25	0.158 ± 0.036	304.796	0.218
26	0.257 ± 0.006	205.072	0.351
29	0.845 ± 0.022	214.892	1.154
30	0.849 ± 0.02	105.842	1.161

Table 4.5: Curvature minimum and wedge angle data for the curved optics.

Piezo Wedge angle

The thickness of all 18 PZTs (Noliac 2124) were measured to characterise their wedge angles (Figure 4.8). For each PZT, thickness was measured at eight points and the orientation was determined by the position of the black dot that represents cathode. A least square fitting of these eight points was done to determine the most likely PZT plane. The thickness variation is 4 μm maximum. The estimated wedge angles vary between 9 and 89 arcsec.

In the following subsection we will see how this knowledge from the characterisation efforts can be used to create a combination of curved mirror, PZT and mounting prism such that the net deviation from perpendicularity is null.

4.4 Sub-assembly matching and bonding

In the preceding sections, we have observed that PZTs and mounting tombstone prisms possess a specific vertical wedge angle resulting from machining tolerances. Furthermore, we have explored how the displacement of the curvature minimum of curved optics relative to the geometric centre translates into a wedge angle when the mirror is affixed to a flat surface.

In Figure 4.9, the bonding of a curved optic with a PZT and a mounting prism, facilitated by a bonding fixture is illustrated. However, there exists a certain degree of freedom in the placement of the PZTs and the orientation of the curved optic fiducial. These orientations significantly impact the overall vertical wedging or deviation from perpendicularity of the optic surfaces. It is imperative to ensure that the cavity axis remains well-centred on the surfaces of the optical components and parallel to the breadboard surface. Consequently, a constraint is imposed on the beam offset from the centre of the optic surfaces

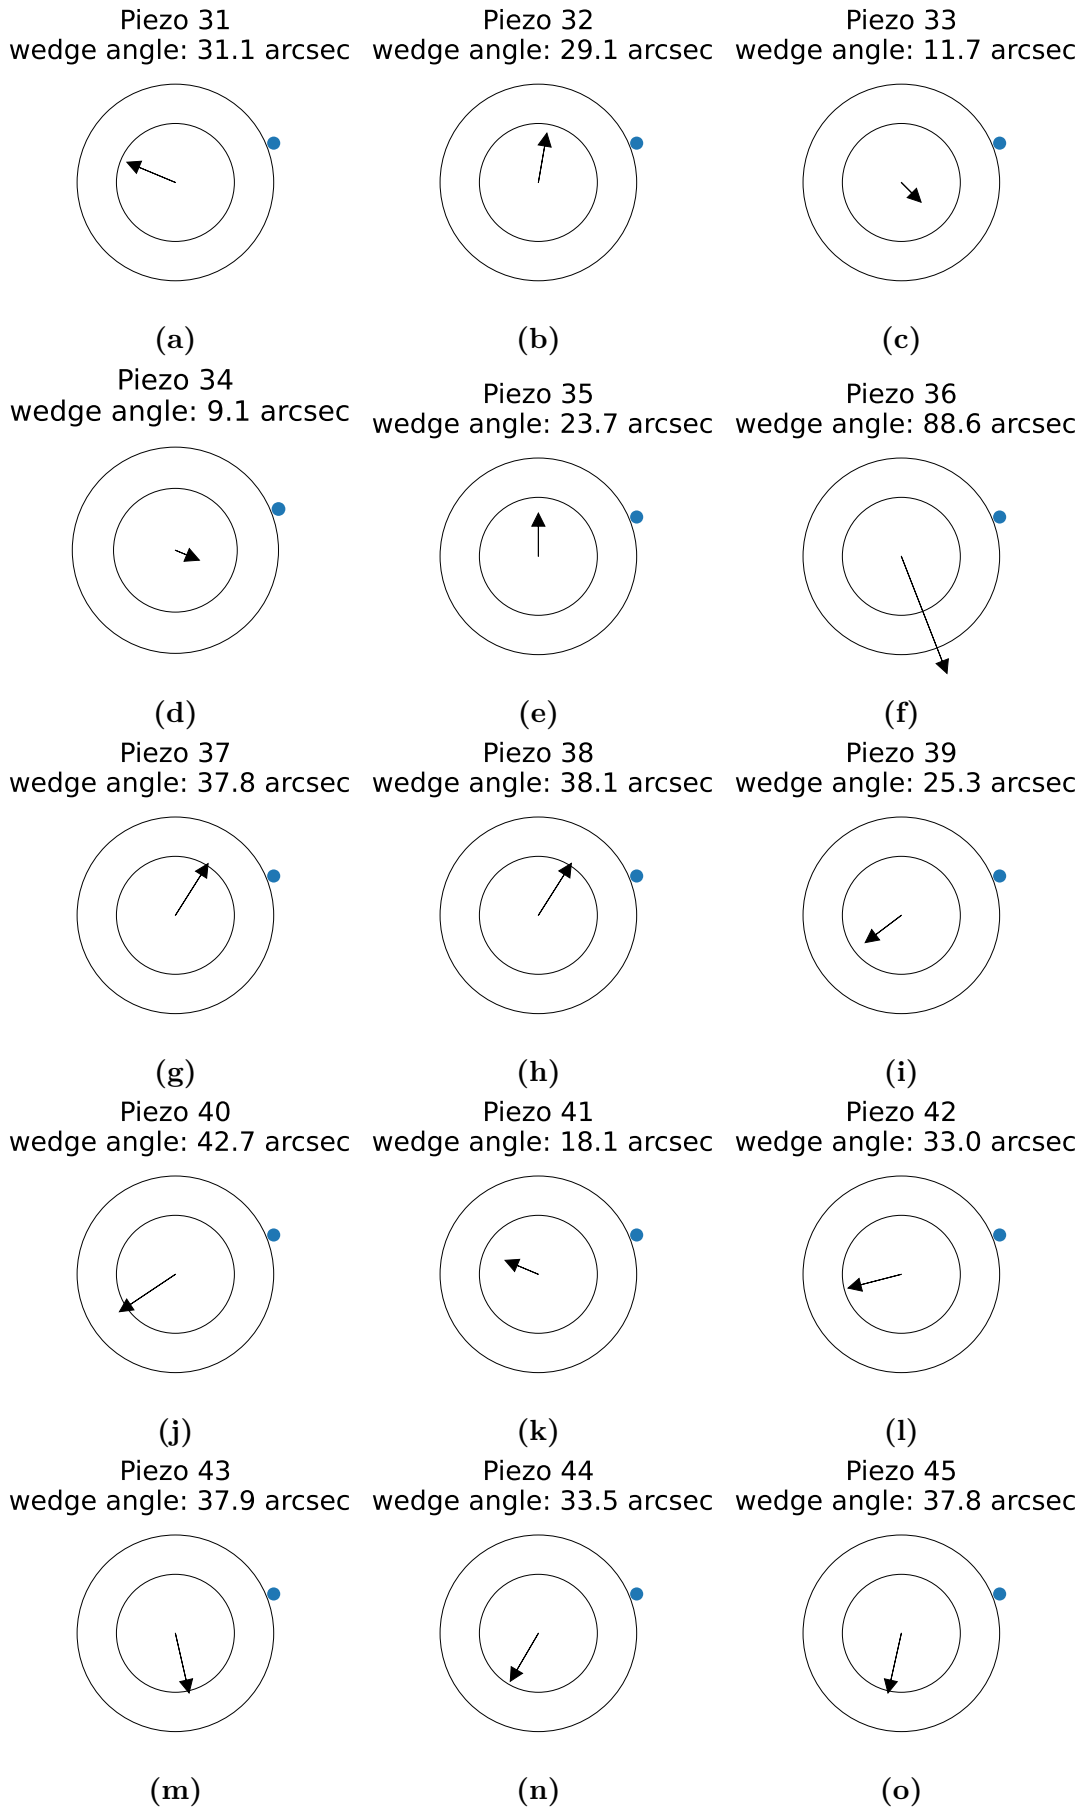


Figure 4.8: Diagrams indicating PZT wedge direction and the length of the arrows indicate the amount of wedging. The number in the bracket is the wedge angle

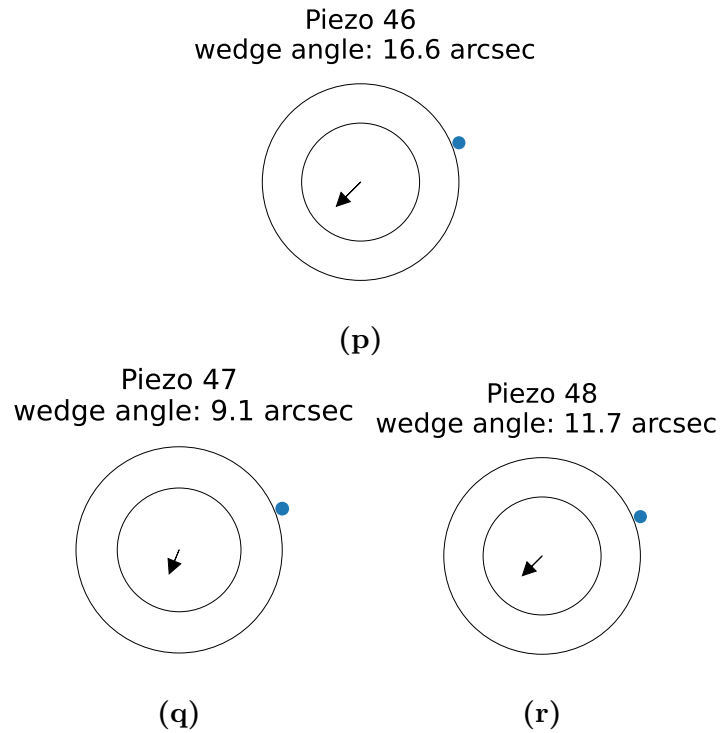


Figure 4.8: Diagrams indicating PZT wedge direction and the length of the arrows indicate the amount of wedging. The number in the bracket is the wedge angle

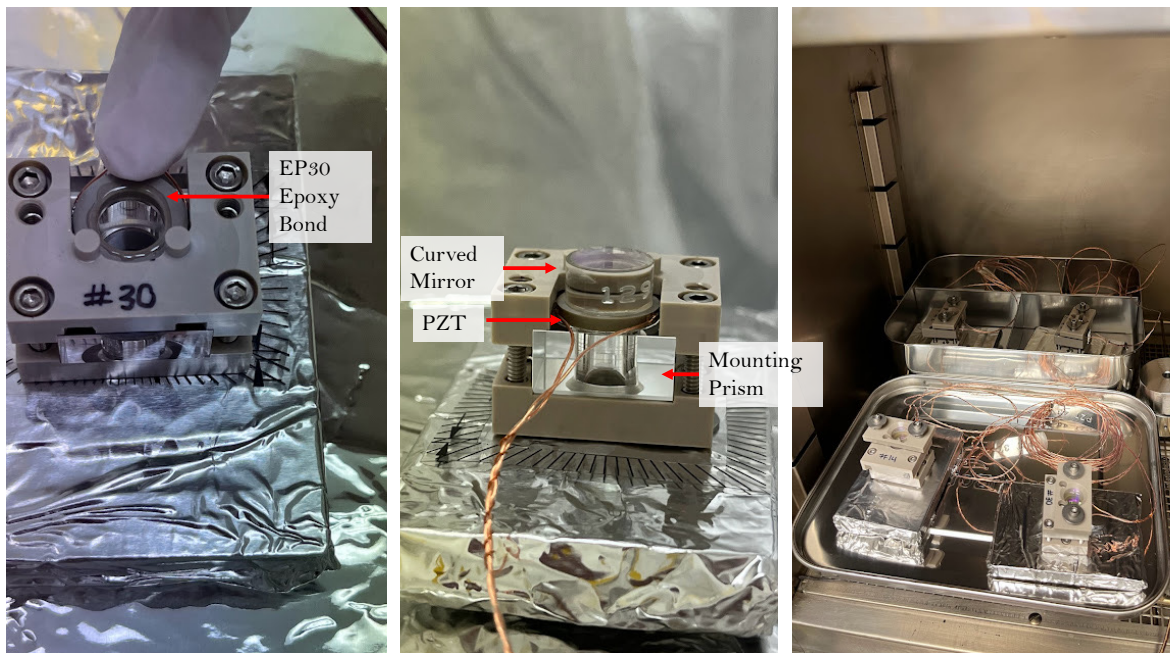


Figure 4.9: Subassembly preparation and bonding process for the A+ LIGO Output Mode Cleaner curved mirrors using a bonding fixture. **Left:** Subassembly bonding step showing EP30 epoxy applied to the PZT surface. **Middle:** Close-up of a fully stacked subassembly consisting of a curved mirror, a PZT actuator, and a mounting prism held in a precision bonding fixture. **Right:** Four completed subassemblies placed inside an oven for curing. This subassembly is later mounted on the optical breadboard along with the flat mirrors to form the full cavity.

along the vertical axis. This deviation should be less than 0.5 mm, or the deviation from perpendicularity should be lower than 40 arcsec [30]. Therefore, it is essential to align the orientations of the PZT sub assembly components in a manner that adheres to this constraint.

A python code was written by me to use the PZT wedge angle, mounting prism perpendicularity deviation and, curved mirror radius of curvature measurement using the beam profiler data along with the mirror scatter maps [29] to match sub-assembly components for the A+ OMC build [150]. We had access to only four subassembly bonding fixtures at a time. The python code goes through some iteration to choose for curved mirror-PZT-Mounting prism and flat prism high reflector combinations that have zero net deviation from perpendicularity. The idea is to prioritise the curved mirrors with least wedge angle, but also ensure that we pick the mirrors with less high scatter points. For the highlighted mirrors the risk factor includes not just the sensitivity to rotation but also the scatter point location. But these scatter points can be placed in the horizontal axis to help us avoid them by rotating the mirror with the help of micrometers.


Sr.No.	CM	Wedge in arcmin	Wedge dir in deg	TIS
1	30	1.161	105.842	1.4
2	14	1.504	9.329	1.6
3	29	1.154	214.892	1.9
4	22	1.506	162.341	2.1
5	21	1.065	126.481	2.4
6	17	1.251	160.430	2.4
7	18	1.047	257.557	2.7
8	12	1.785	269.077	2.0
9	6	1.729	18.443	2.2
10	16	1.608	129.497	2.3
11	2	1.903	311.015	2.9
12	4	1.159	141.296	2.8
13	25	0.218	304.796	4.0
14	23	1.252	25.164	4.0
14	20	1.298	257.958	4.6
16	5	1.082	112.812	5.0
17	11	0.098	37.842	5.1
18	19	0.810	160.040	5.5
19	26	0.351	205.072	8.4
20	7	1.144	300.549	2.2

Table 4.6: Wedge parameters and TIS for various CMs.

A total of 16 mounting prisms, 19 PZTs, and 20 curved mirrors were available for OMC construction. Visual inspection of the mirror scatter maps was used to screen and categorise the curved mirrors based on their scatter maps. Each mirror was assigned to one of three categories: “yes”, “maybe”, or “no”, depending on the density and location of scatter points in the TIS map. Mirrors in the “yes” category exhibited minimal scattering and were considered high quality. “Maybe” mirrors showed a moderate number of medium-level (green) scatter points, typically distributed away from the central region. Mirrors in the “no” category displayed high-intensity scatterers in critical areas and were excluded from use. Representative examples from each category are shown in Figure 4.10. This screening process is essential for constructing high-throughput (>99%) OMCs, where optical loss must be minimised through careful component selection.


Since we have only 16 mounting prisms this is the limiting number for the possible number of subassemblies. Also the goal of the A+ project is to build 6 OMC cavities and two curved mirror subassemblies per cavity implies 12 subassemblies in total. The code goes through all the possible combinations of mounting prisms and PZTs to find the best match for each of the curved mirrors in the table 4.6. For more details on matching refer to [151]. The idea is to direct wedging in the p- plane such that the resulting wedge in the s- plane is zero. Gluing instruction sheet for one of the combinations of mounting prism, PZT and curved mirror is shown in Figure 4.11.

For more details on subassembly bonding procedure from aLIGO efforts refer to [33] and further improvements for subsequent A+ OMC cavity build, refer to [111].



(a) TIS for CM 11. Here the red scatter on the mirror makes it relatively undesirable for cavity construction.

(b) TIS map for CM 16: Here medium scatterers are distributed across the mirror HR surface.



(c) This is a good mirror. CM 30 has very low scatter points and is an ideal mirror for cavity construction.

Figure 4.10: Scatter map for three of the A+ OMC curved mirrors. Measurements and plots by Dr. Liyuan Zhang, LIGO Lab, Caltech [183].

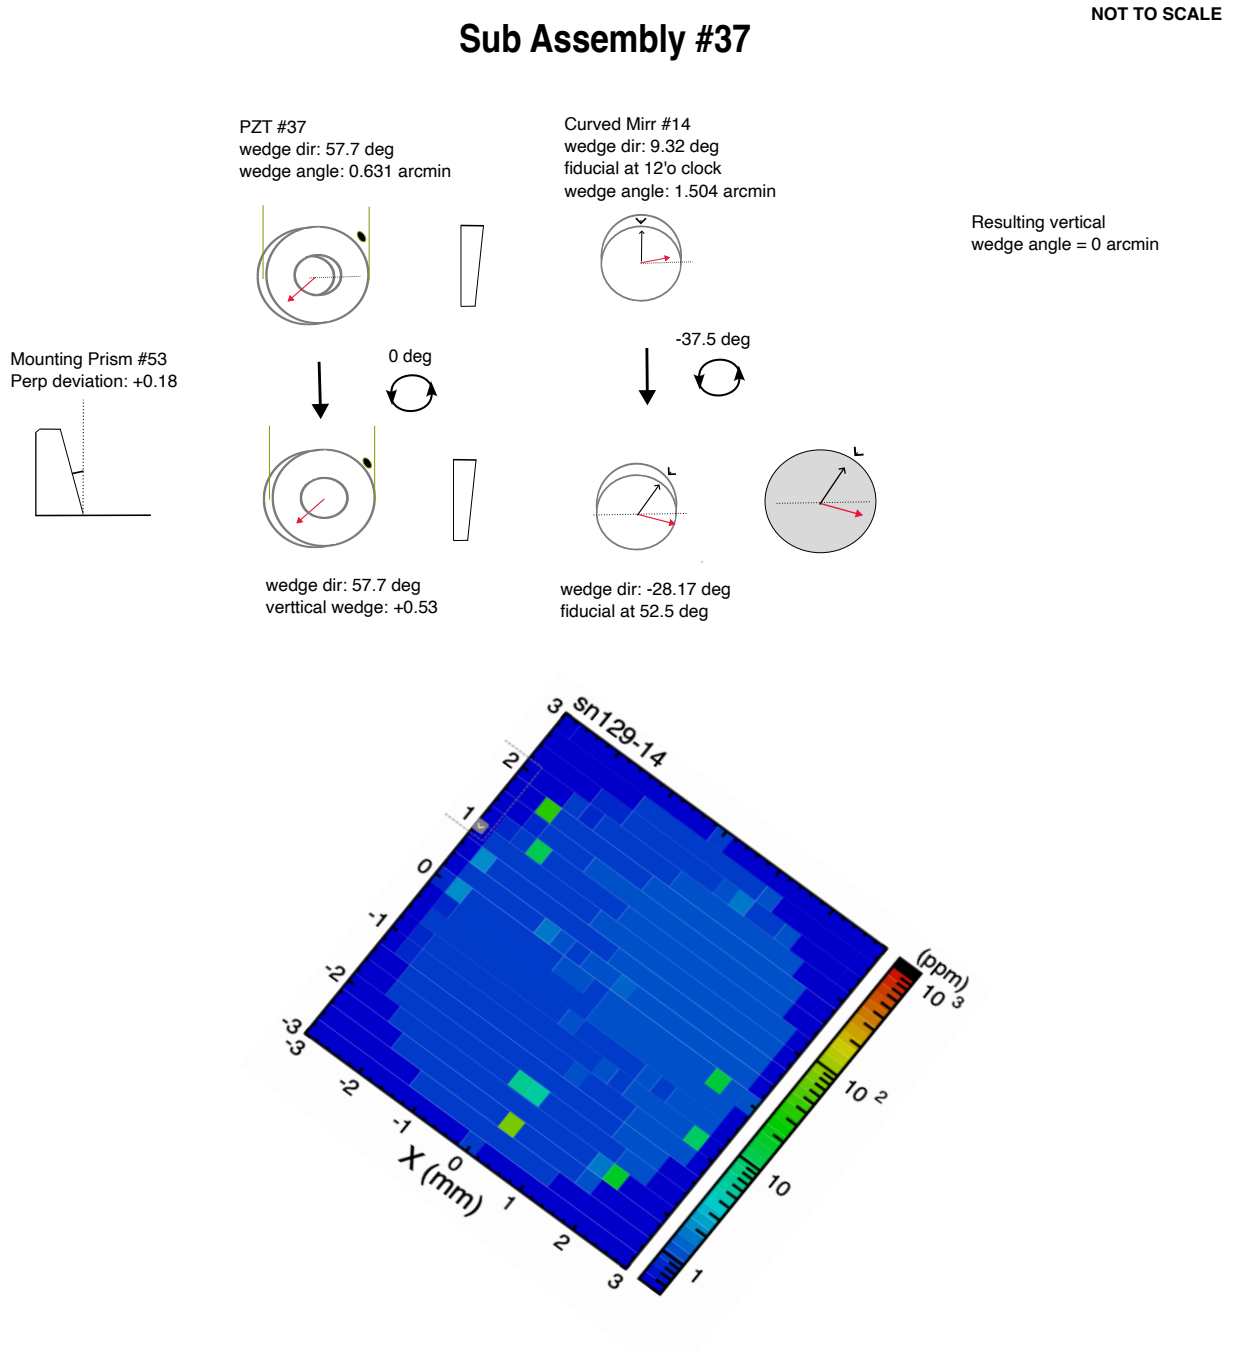


Figure 4.11: Example curved mirror subassembly glueing instruction sheet for A+ OMC curved mirror subassembly.

4.5 Cavity Assembly

The setup was built on an optical table in the West Bridge Labs at Caltech. The schematic diagram of the optical layout can be found in Figure 5.5. The optical cavity was stabilised using the Pound-Drever-Hall (PDH) locking technique. Phase modulation sidebands were imposed on the input laser field at a frequency of 32.7 MHz using a resonant electro-optic modulator (EOM). The light reflected from the cavity was detected by a broadband photodetector (Thorlabs PDA255, bandwidth ~ 50 MHz), and the resulting signal was demodulated at the modulation frequency to generate the PDH error signal. Feedback control was implemented using a Newport LB1005 servo controller to apply the appropriate filtering. The error signal was used to actuate the laser frequency via the fast piezoelectric transducer (PZT) on the laser.

The input beam to the cavity was delivered via a polarisation-maintaining single-mode fiber (Thorlabs P3-1064PM-FC-5), which ensured spatial mode filtering and polarisation stability. Each end of the fiber was equipped with a collimation lens (Thorlabs CFC-2X-C) mounted on a six-axis fiber positioner (Thorlabs K6X), allowing precise alignment for mode matching both into the fiber and from the fiber output to the input of the output mode cleaner (OMC) cavity. Accurate mode-matching of the input beam to the OMC cavity was essential for precise power budget measurements. The mode-matching telescope consisted of two plano-convex lenses with focal lengths of $f = 35$ mm and $f = 125$ mm. The telescope was assembled on a separate sub-breadboard allowing the mode-matching quality to be reliably reproduced by preserving the alignment of components on the sub-breadboard, including the fiber and fiber coupler. The beam was elevated to the height of the OMC cavity on the transport fixture using a periscope.

For PDH locking, the reflected beam from the cavity was detected by a Thorlabs PDA100CF photodetector. The transmitted light through the cavity was monitored using two Thorlabs PDA20CS photodiodes, which were used for alignment and transmission diagnostics.

The OMC was mounted on a transport fixture. The transport fixture was mounted rigidly on the optical table to ensure the cavity doesn't move during assembly and testing. The glass breadboards were cleaned thoroughly using high grade IPA red wipes after using a nitrogen gun to blow away contaminants (carefully away from rest of the optics on the bench). The breadboard was examined under halogen light for contaminants from the lens tissue as it can shed fibres that are hard to see under normal lighting 4.12.



Figure 4.12: Inspection and cleaning of the OMC glass breadboard using high-purity IPA wipes and halogen lighting to remove and detect particulate and fibrous contaminants prior to template bonding.

To centre the assembled template [27] on the glass breadboard, we used calipers to ensure precise alignment. The breadboard was then secured in place by tightening the surrounding brackets. It is important to apply minimal force during this step, as excessive pressure can cause the template to shift on the smooth glass surface. We found it most effective to begin by securing one pair of opposite sides before moving on to the remaining sides. For

added convenience and speed, we typically used two calipers and enlisted an extra pair of hands. The position mode-matching template was also adjusted such that the beam waist coincides with the centre of the template i.e. the cavity waist. The mirrors were placed in their respective template cutouts.




Figure 4.13: Cavity bonding template mounted and aligned on the OMC glass breadboard. The template is centred such that the cavity waist coincides with the beam waist, and mirrors are placed in their designated cutouts.

Each mirror and the top surface of the glass breadboard were carefully cleaned prior to assembly. The procedure involved blowing the surface with an ionized air gun, followed by wiping with lens tissue soaked in isopropyl alcohol. Particular care was taken to clean only the areas immediately surrounding the cavity mirror cutouts.

The flat mirrors (FM1 and FM2) were placed into their respective template openings and secured against the pads using an Allen key. The curved cavity mirrors (CM1 and CM2) were then positioned. This involved first adjusting the micrometer adapter plates [28] to their intended positions, centring the mounting screws within the available slot length to set the nominal cavity length. Both micrometers were extended approximately to the




Figure 4.14: Bonding template setup used for positioning the curved cavity mirrors (CM1 and CM2). The micrometer adapter plates are adjusted to set the nominal cavity length, and the mirrors are placed such that they simultaneously contact the micrometers and template pads. Visual and tactile checks are performed to ensure proper seating and alignment before final securing.

midpoint of their travel range. The curved mirrors were then placed into their template openings such that they simultaneously contacted the micrometers and template pad. Visual and tactile checks were used to ensure the mirrors were seated flush at the bottom face and free from interferences around the three contact points.

Once both curved mirrors were installed, their symmetry and alignment were verified. Final adjustments were made by repositioning the adapter plates as needed. The optics were then secured in place using Allen screws.

4.6 Cavity Alignment Procedure

The following steps outline the beam alignment process used to position the optical mode through the cavity mirrors with high precision. Here, the laser beam is walked using steering mirrors 1 and 2 (SM1 and SM2 respectively, SM2 is closer to the cavity) and enters the cavity through flat mirror 1 (FM1 - see Figure 4.3).

1. Initial Beam Centring on FM1 (Flat Mirror 1):

The beam was roughly centred on FM1 using the steering mirror SM1. The position was monitored using an IR sensor card. Vertical centring was defined as 15 mm above the optical table surface (or 8 mm below the top surface of FM1).

2. Initial Beam Centring on FM2 (Flat Mirror 2):

Similarly, the beam was roughly centred on FM2 using SM2, with vertical alignment again set to 15 mm above the table (or 8 mm below FM2's top surface), monitored via the IR sensor card.

3. Verification of Beam Path Between SM1 and SM2:

The beam spot was checked to ensure it propagated close to the midpoint between SM1 and SM2. If the spot was near the mirror edge or misaligned, the mirror positions were iteratively adjusted by:

- (a) Placing an IR card between SM2 and FM1 to visualise the beam.
- (b) Adjusting the relevant steering mirror to bring the beam within 1 mm of the mirror centre.
- (c) Confirming alignment to the input iris, FM1, and FM2 centres.
- (d) Fine-tuning the other steering mirror if necessary.

4. Beam Centring on CM1 (Curved Mirror 1):

The beam was directed toward CM1 by adjusting SM2 only. If alignment could not be achieved, it was verified that CM1 was correctly seated against its micrometer mounts and that both SM1 and SM2 were not misaligned. If issues persisted, the procedure reverted to step 1.

5. Beam Centring on CM2 (Curved Mirror 2):

Horizontal alignment was performed exclusively using adjustments to CM1. If vertical alignment was off, the beam was “walked” between CM1 and CM2 using coordinated adjustments of SM1 and SM2. CM1 was then gently positioned into contact with its micrometers and a three-point contact pad using minimal force to avoid stressing the optic.

6. Final Beam Centring on FM1 (Return Path Overlap):

Finally, the beam position was refined to ensure proper overlap with the original beam spot on FM1, confirming alignment consistency throughout the cavity.

Fine Alignment

Once a resonance flash is observed, a camera is placed at the transmission port to monitor the cavity output; initially, high-order transverse modes may be present. The beam is then walked towards the fundamental TEM_{00} mode via coordinated adjustments of the steering mirrors SM1 and SM2. After a clear TEM_{00} transmission is obtained, each cavity mirror is lightly touched at the top to verify it is well-seated. Finally, the OMC transmission and reflection photodiodes (TRANS PD and REFL PD) are aligned using steering mirrors so the beam is fully incident on the active area; lenses are used to focus the beam onto the photodiodes without saturation, enabling cavity lock acquisition once sufficient TEM_{00} power is available.

Once the cavity was locked, fine alignment was performed using pico-motorised steering mirrors.

4.7 Cavity Characterisation

The OMC cavity is characterised pre-bonding and post-bonding. Characterisation efforts involve the following:

1. Transverse Mode Spacing (TMS) and Free Spectral Range (FSR) measurement.
TMS and FSR are influenced by the geometry of the cavity,
2. PZT DC response,
3. Power Budget Measurement to evaluate the cavity throughput and loss.

Pre-bonding characterisation of the cavity is important as this helps us identify any issues with the cavity before bonding. UV bonding is irreversible so it's important to make sure that the cavity parameters meet the requirements. In Table 4.7 below we see all the important parameters as result of the measurements that we will see in detail in the upcoming subsections.

4.7.1 Cavity Transverse Mode Spacing and Free Spectral Range

Cavity Free Spectral Range (FSR) is the frequency spacing between consecutive TEM₀₀ fundamental resonance modes of the cavity. Transverse Mode Spacing (TMS) is the frequency mode separation between consecutive higher order modes. For more theoretical background on optical cavities and laser beams refer to [153][97].

The TMS and FSR of the OMC cavity was measured by scanning the sidebands with a locked carrier. An additional phase modulation was applied using a broadband EOM with the modulation frequency scanned by a network analyser functionality in MOKU. Some misalignment is introduced in pitch and yaw directions independently to measure TMS yaw and TMS pitch, which are different due to inherent astigmatism in the cavity. More details about this method can be found in [172]. A small clipping of the beam on

Index	Parameter	Value	Error (\pm)
[1]	Optical Throughput (%)	99.10	0.49
[2]	Optical Loss (%)	0.87	0.49
[3]	Free Spectral Range (MHz)	266.4614	0.0019
[4]	Roundtrip cavity length (m)	1.125867	0.000008
[5]	Transverse Mode Spacing, vertical (MHz)	58.7325	0.0010
[6]	TMS as fraction of FSR, vertical	0.220417	0.000004
[7]	Transverse Mode Spacing, horizontal (MHz)	58.66890	0.00100
[8]	TMS as fraction of FSR, horizontal	0.220178	0.000004
[9]	TMS to PZT1 voltage, linear coeff., vertical	-1.304E-05	5.7E-07
[10]	TMS to PZT1 voltage, y-intercept, vertical	0.220155	0.00063
[11]	TMS to PZT1 voltage, linear coeff., horizontal	-1.294E-05	4.1E-07
[12]	TMS to PZT1 voltage, y-intercept, horizontal	0.220518	0.000430
[13]	TMS to PZT2 voltage, linear coeff., vertical	-1.30E-05	9.4E-07
[14]	TMS to PZT2 voltage, y-intercept, vertical	0.220155	0.000630
[15]	TMS to PZT2 voltage, linear coeff., horizontal	-1.37E-05	3.7E-07
[16]	TMS to PZT2 voltage, y-intercept, horizontal	0.220518	0.000043
[17]	PZT1 Operation Range (V)	77	–
[18]	PZT2 Operation Range (V)	78	–
[19]	PZT1 Actuation Range (μm)	0.90845	–
[20]	PZT2 Actuation Range (μm)	0.99138	–
[21]	PZT1 Actuation Range (/FSR)	1.71	–
[22]	PZT2 Actuation Range (/FSR)	1.77	–
[23]	PZT1 DC Response (nm/V)	11.798	0.097
[24]	PZT2 DC Response (nm/V)	12.710	0.11

Table 4.7: A+ OMC cavity performance parameters summary

the transmission PD is also introduced. This is because lobes of Higher Order Modes have opposite polarity [97] and cancel each other when beat against the carrier, introducing a this clipping on the transmission PD favours detection of one polarity over the other. We then measured the transfer function between excitation of the EOM and transmission PD. As the modulation frequency is swept, resonances are observed whenever either sideband becomes coincident with a cavity higher-order mode (HOM). The resulting peaks occur at frequencies determined by the HOM order m , with the upper sideband producing peaks at

$$f_m^{(u)} \simeq m f_{\text{TMS}}, \quad (4.5)$$

and the lower sideband producing peaks at

$$f_m^{(l)} \simeq f_{\text{FSR}} - m f_{\text{TMS}}, \quad (4.6)$$

for $m = 1, 2, \dots$

In Figure 4.15, the transfer function of the transmission photodiode is shown for an EOM sideband frequency scan. The HOM resonances are labelled by mode order in the main panel, and the corresponding fits for each identified peak are shown in the subpanels below.

When the scanning sideband modulation frequency is at the FSR, a dip in the transfer function is seen. This happens because the cavity transfer function is symmetric for TEM₀₀ unlike for the HOM and both the positive and negative sidebands pass through the cavity at FSR frequencies.

Higher Order Mode Spectrum Analysis

The TMS values and FSR values measured go beyond just characterisation of cavity geometry. The ratio between TMS and FSR for the pitch and yaw degrees of freedom gives Gouy phase of the HOM:

$$\text{Gouy Phase of TEM}_{mn} = m \cdot \frac{\text{TMS}_x}{\text{FSR}} + n \cdot \frac{\text{TMS}_y}{\text{FSR}} \quad (4.7)$$

where $m, n = 0, 1, 2, 3, \dots$

This enables us to create an HOM plot for a particular voltage setting for PZT 1 and 2 i.e. for a certain cavity roundtrip length as shown in Figure 4.16. The goal here is to tune the cavity length before bonding the optics such that the vertical black line in the plot avoids HOM and their control sidebands. This ensures that in case a HOM is rung up in the laser beam, the cavity does not transmit it. As shown in the plot, we have successfully avoided all the lower HOMs by following a cavity length tuning procedure. In Figure 4.17, HOM bands for full PZT scan from 0 to 200 V for each PZTs are shown.

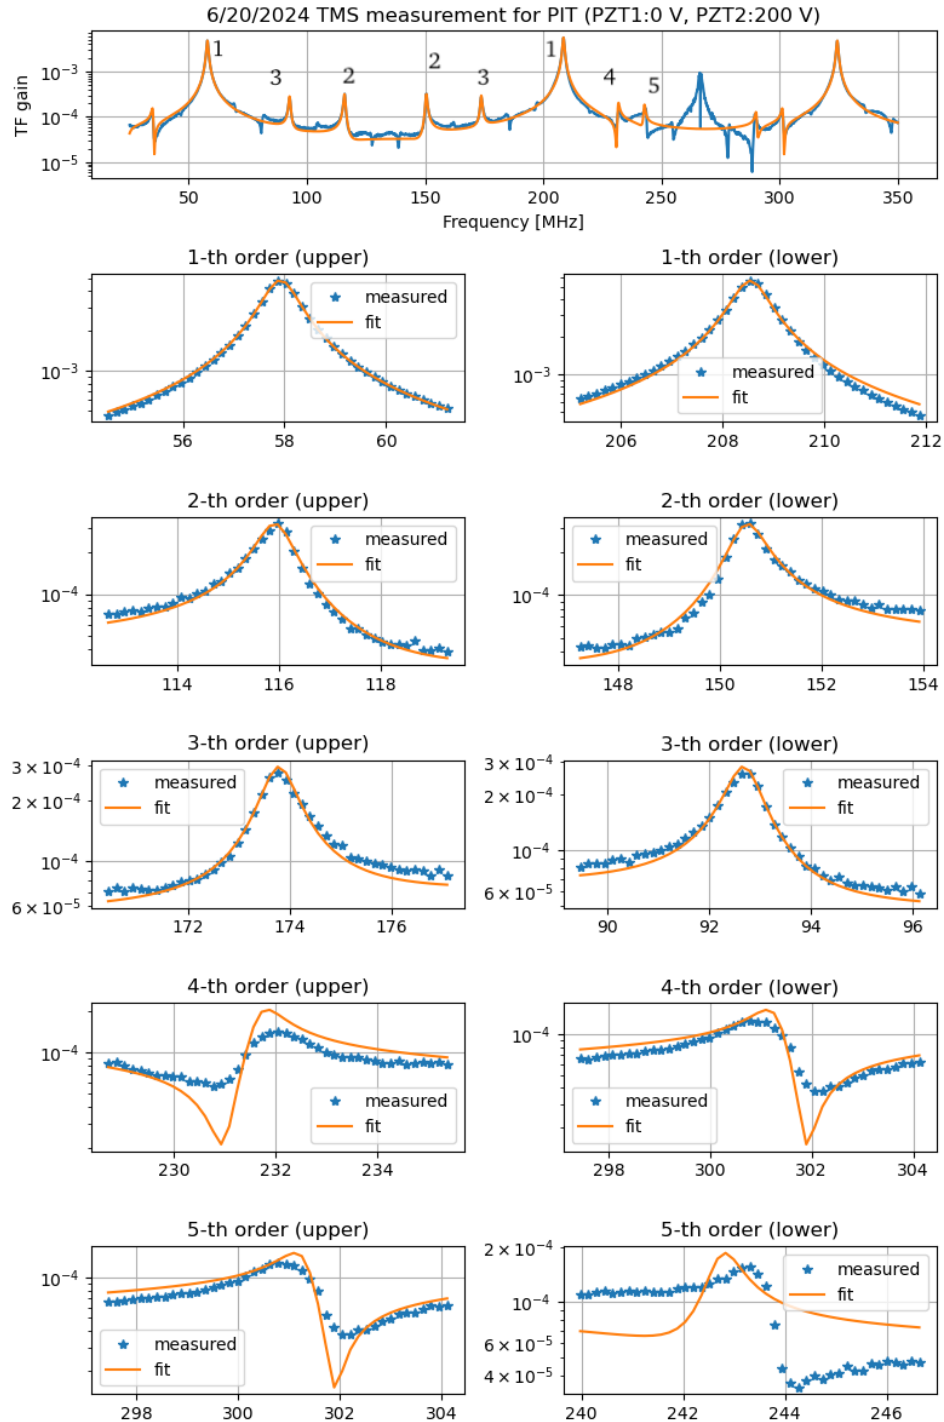


Figure 4.15: Fitting of transfer function between excitation of the EOM and the OMC transmission PD: As the EOM sideband is scanned, the various HOM sidebands beat with the carrier. When the carrier sideband gets closer to FSR, the transfer function sees a dip in frequency which is seen in the plot above between 250 MHz and 300 MHz. This scan is done with a pitch misalignment thus the HOMs in the plot are TEM_{x0} modes for $x = 1, 2, 3, \dots$. The OMC mirrors are bonded here (*Courtesy: M. Nakano, A+ OMC team*).




Figure 4.16: Higher Order Mode (HOM) distribution in the optical cavity, showing the Gouy phase (expressed in units of FSR) as a function of mode order. Each colour represents a different RF sideband frequency used for controlling the interferometer and auxiliary optical cavities in LIGO. The data corresponds to the configuration where PZT1 is held at 0 V and PZT2 at 50 V, resulting in measured mode spacings of $TMS_y/FSR = 0.21983$ for yaw degree of freedom and $TMS_p/FSR = 0.21950$ for pitch degree of freedom. The vertical black line at the centre marks the laser fundamental mode carrier resonance frequency. This measurement was taken immediately after bonding the optics on the cavity.





Figure 4.17: Transverse mode distribution as a function of PZT drive voltage, showing the frequency locations of higher-order modes (HOMs) and their sidebands relative to the fundamental mode. Each coloured band corresponds to a specific HOM and its frequency shift as the cavity length is varied through PZT actuation. The vertical spread of each band arises from cavity astigmatism, which causes the transverse mode spacing (TMS) to vary between the sagittal and tangential Gouy phase limits. The black dashed lines indicate the minimum required frequency separation from the TEM₀₀ carrier (two linewidths) for reliable OMC operation. The green dashed region marks the optimal PZT range where modal degeneracies are avoided while satisfying this separation criterion—labelled as the “useful PZT range” on each plot. The top panel corresponds to length tuning via PZT1, and the bottom panel to tuning via PZT2.

Cavity length tuning

In the pre-bonding phase, measurements of the transverse mode spacing (TMS) and free spectral range (FSR) help determine whether the cavity axis yields a higher-order mode (HOM) spectrum that meets OMC operational requirements—that is, no coincidental resonance of HOM carriers or sidebands with the fundamental mode, up to at least the 18th order. An example of an unsatisfactory configuration is shown in Figure 4.18. From Figure 4.16, we know that the optimal cavity length is approximately 112.5 cm. The micrometers shown in Figure 4.14 are used to rotate the curved mirrors, thereby tuning the cavity length by adjusting the optical axis.

The alignment procedure involves determining the required direction to steer the cavity axis toward the centre (within 1–2 mm), either by translating or pivoting the axis. There are three main ways to move the cavity axis between the curved mirrors (CMs):

- (a) Shift the axis toward FM1 and FM2 by rotating CM1 clockwise (CW) and CM2 counterclockwise (CCW).
- (b) Shift the axis away from FM1 and FM2 by rotating CM1 CCW and CM2 CW.
- (c) Yaw the cavity axis about the midpoint between CM1 and CM2.

The goal is to centre the beam spots on the curved mirrors, assuming those regions are free of scatterers. A typical rule of thumb is that a 1 mrad rotation corresponds to a 1 mm beam shift at the mirror, which equals a $10\ \mu\text{m}$ displacement on the micrometer. After tuning, the laser is aligned to the cavity to optimize the TEM₀₀ resonance.




Figure 4.18: Higher Order Mode (HOM) distribution in OMC cavity with roundtrip length of 111.7 cm: In this plot, we see the Gouy phase (in units of FSR) as a function of mode order. Each colour corresponds to a different RF sideband frequency used in the LIGO interferometer control scheme. In this configuration, PZT1 is held at 0 V and PZT2 at 50 V, resulting in measured mode spacings of $\text{TMS}_y/\text{FSR} = 0.2191$ and $\text{TMS}_p/\text{FSR} = 0.2186$. This HOM configuration is not suitable for OMC operation due to a coincidental resonance of the 9th-order mode sideband with the fundamental mode, which violates the requirement of avoiding higher-order mode overlap within at least two linewidths of the fundamental resonance. Such modal degeneracy can lead to excess power leakage through the OMC and reduced filtering performance.

4.8 Cavity Throughput Optimisation

In addition to satisfying the geometric constraints outlined above, the optical cavity must also meet the transmission throughput requirement. Specifically, the cavity should transmit approximately 98% of the incident power for a perfectly mode-matched TEM₀₀ carrier beam. However, in practice, point defects on the mirror surfaces can reduce the achievable throughput below this target if no further optimisation is applied. To mitigate this, careful alignment is necessary to position the beam spots away from scattering sites, thereby maximising transmission efficiency through the cavity.

Cavity throughput optimisation involves several key steps. The first step is the careful selection of mirrors from the available batch, ensuring that only those with minimal surface defects and low scattering loss are used. If localised scatterers are present, their positions are mapped and oriented such that they lie along the vertical axis and away from the central region of the mirror. This precaution reduces the likelihood of the beam interacting with high-scatter or high-absorption regions due to residual or accidental misalignments during alignment and operation. This process is covered in detail in section 4.4 (refer Figure 4.10). During the cavity alignment process, the beam spots on the mirrors are frequently inspected using CCD cameras operated at low exposure settings to reveal high-scatter regions. If a significant scatterer is identified—especially one visible to the naked eye under halogen illumination—it is typically addressed by carefully blowing inert gas using a gun on the affected mirror surface pointing away from the cavity. This procedure is complemented by power budget measurements at various points along the optical path, both inside and outside the cavity, to monitor losses and throughput.

If no obvious source of loss is visible on the mirror surfaces, a more systematic cleaning approach is employed. First Contact polymer is applied sequentially to each mirror surface, with cavity throughput measured after each cleaning cycle. This process allows the identification of the mirror contributing most significantly to the optical loss. Every time the cavity mirrors are touched it is crucial to perform FSR measurement to ensure the

cavity geometry is not altered. In the worst case event of presence of a high scatter on cavity axis that is not removable by any of the above techniques, then the cavity axis can be slightly tuned until the cavity loss is small. This should be complemented with FSR measurement to track changes in cavity geometry. Results from the cavity power budget measurements for the first A+ OMC cavity is shown in Figure 4.8.

Category	Parameter	Value	Error	Units
Power Measurements	BS TRANS	33.55	0.10	mW
	BS REFL	0.00	0.00	mW
	CM1 Transmission	23.9	0.10	μ W
	CM2 Transmission	23.8	0.10	μ W
	Cavity Injection	34.33	0.10	W
REFL PD	Locked	31.8	2.00	mV
	Unlocked	3.42	0.02	V
Offset	REFL	-6.01	0.01	mV
	Reference	3.90	0.01	mV
Cavity Parameter	Finesse	390	10	—
PB Result	Throughput	0.9918	0.0088	—
	Loss per mirror	14	18	ppm
	Mode matching	0.98898	0.00060	—
Transmission	FM loss	8000	210	ppm
	CM1 loss	5.72	0.15	ppm
	CM2 loss	5.72	0.15	ppm

Table 4.8: Summary of measured and derived cavity parameters from power budget analysis of the bonded A+ OMC cavity 101.

4.9 PZT Characterisation

The PZT DC response was measured by sweeping the PZT from 0 V to 120 V. We injected the PZT with a 3 Hz sinusoidal excitation signal and recorded the TRANS photodiode (PD) signal over 30 cycles. The recorded signal was then divided into 30 segments corresponding to increasing and decreasing PZT voltage, and averaged separately. The resulting TRANS PD signal is shown in Figure 4.19.

The top and bottom plots display the averaged PD signals for upward and downward sweeps, respectively. Using the TEM₀₀ and HOM peaks as frequency markers, the PZT displacement was estimated and is shown in the right figure.

The fitting was performed using the function:

$$\text{PZT}(v) = a_1 \cdot (v - v_{\min}) + \frac{\dot{v}}{|\dot{v}|} \cdot a_2 \cdot (v_{\max} - v) \cdot (v - v_{\min}) + \text{ofs}$$

where a_1 is the actuator efficiency, a_2 is the hysteresis parameter, v_{\min} and v_{\max} are the minimum and maximum values of the sinusoidal input, and $\frac{\dot{v}}{|\dot{v}|}$ indicates the direction of the voltage sweep. The first term in the above equation is the important term with a_1 giving the PZT response, the second term is the hysteresis term that assumes a parabolic shape with values equating to zero at $v = v_{\min}$ or $v = v_{\max}$ and ofs emphasises any zero offset. The fitting parameters are a_1 , a_2 , and ofs.

The fitted value of a_1 are,

- PZT1 DC response: 10.056 ± 0.082 nm/V
- PZT2 DC response: 10.54 ± 0.17 nm/V

and the hysteresis parameters a_2 are,

- PZT1 hysteresis parameter: 0.00237 ± 0.00012 /V
- PZT2 hysteresis parameter: 0.00218 ± 0.00018 /V

The AC response of the PZT actuator was characterised using transfer function measurements between the OMC PZT excitation signal and the laser PZT loop feedback signal.

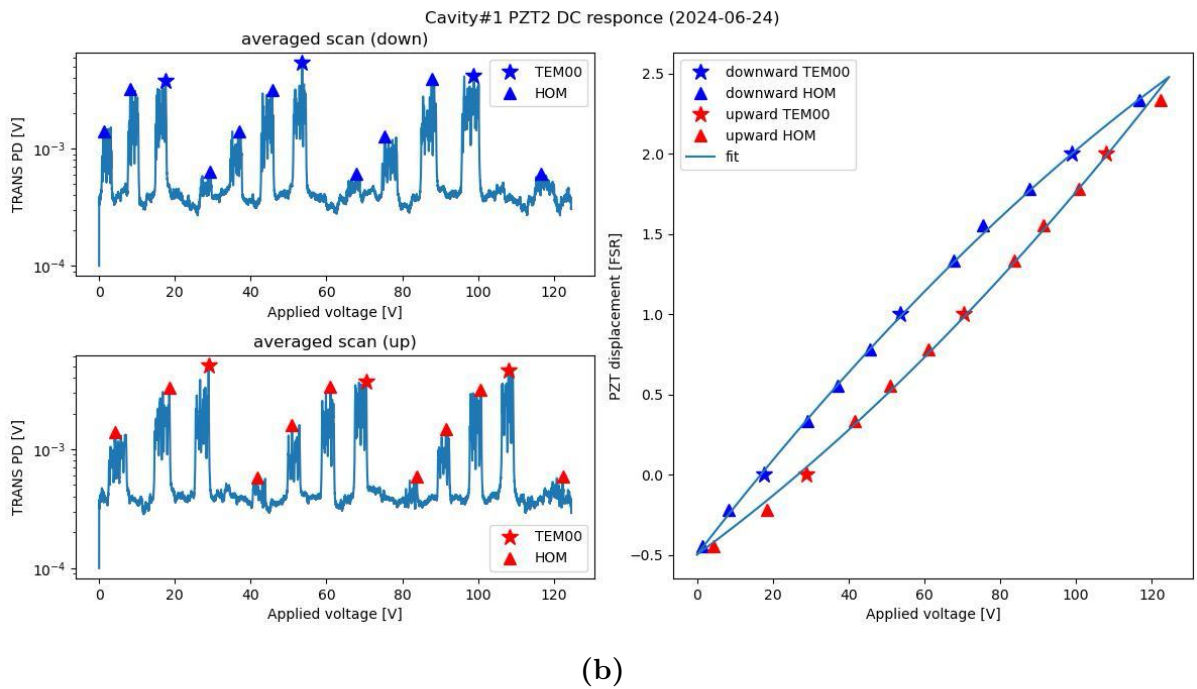
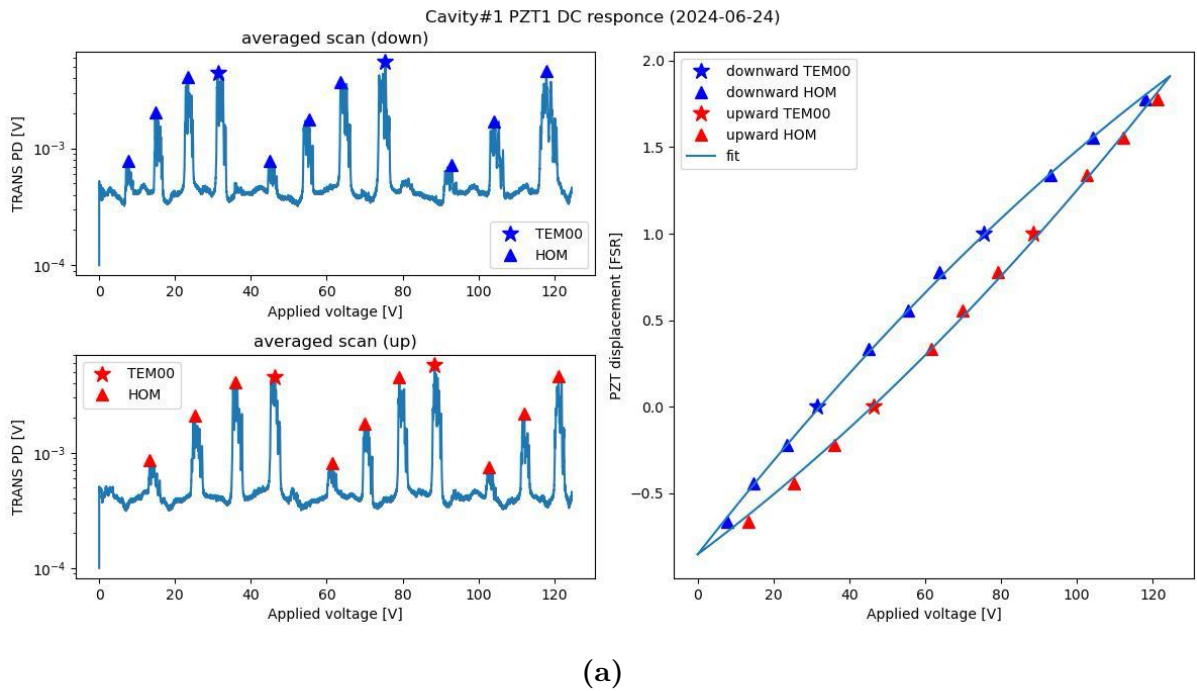


Figure 4.19: DC Response of PZT Actuators. (a) and (b) correspond to measurements for PZT 1 and PZT 2, respectively. The plots on the left show cavity transmission as a function of PZT drive voltage, scanning across several free spectral ranges (FSRs). Prominent peaks correspond to the fundamental TEM_{00} mode and higher-order modes (HOMs). Separate traces are shown for upward and downward voltage sweeps. Using the known cavity geometry to determine the FSR and transverse mode spacing (TMS), the plots on the right convert this data into PZT displacement versus applied voltage, characterising the DC response of each actuator (*Courtesy: M. Nakano, A+ OMC team*).

To begin, the OMC was locked using the laser PZT, and an excitation signal was injected into the OMC PZT. The servo output signal was then connected to a spectrum analyser to measure the transfer function from the excitation input to the servo output. This transfer function (TF1) can be expressed as

$$\text{TF1} = A_{\text{OMC}} \cdot H \cdot F / (1 + G), \quad (4.8)$$

where A_{OMC} is the actuator efficiency of the OMC PZT, H is the optical gain of the OMC photodetector, F is the servo transfer function, and G is the open-loop gain.

Next, the excitation point was moved to the error point of the loop (the alternate input to the servo), and the transfer function to the feedback signal was measured again—this time representing TF2. This measurement yields

$$\text{TF2} = F / (1 + G). \quad (4.9)$$

Therefore, by taking the ratio TF1/TF2, we obtain the AC actuator response as

$$\frac{\text{TF1}}{\text{TF2}} = A_{\text{OMC}} \cdot H. \quad (4.10)$$

Since H is known to be flat over the measurement bandwidth, this ratio directly provides the frequency-dependent AC response of the actuator. The DC response is calibrated independently and is available from previous measurements.

4.10 Cavity Bonding

Once all optical parameters were satisfactorily adjusted, the prism mirrors were bonded sequentially, starting with flat mirrors. Prior to bonding the curved mirrors CM1 and CM2, the transverse mode spacing (TMS) to free spectral range (FSR) ratio, $f_{\text{TMS}}/f_{\text{FSR}}$, was measured to ensure it remained within the specified tolerance. If deviations were observed, the cavity length was readjusted by repositioning CM1 and CM2.

After bonding CM1, the TMS/FSR ratio was re-evaluated. If the ratio remained out of tolerance, the cavity geometry was readjusted once more before proceeding. Finally, CM2 was bonded in place. Upon completion, the final values of both the TMS and FSR were confirmed and recorded for reference. The cavity beam spot positions are also closely monitored during every iteration of bonding as shown in figure 4.20.




Figure 4.20: Cavity beam spot position verification by imaging beam spots on the cavity mirrors. The cavity axis appears to be well centred by eye, with beam spots located within 1 mm of the mirror centres.




Figure 4.21: OMC 101 mirrors bonded.

4.11 Conclusion and Future Directions

This chapter was written at the end of my visit to Caltech in 2024. The cavity mirrors were bonded and summary of the cavity parameters is shown below:

Table 4.9: Cavity Summary (after bonding)

Parameter	Value
Serial #	101
FSR (MHz)	266.4614
Cavity Length (m)	1.1259
Throughput (%)	99.10%
PZT1 Actuation Range (V)	72.7
PZT1 Operating Voltage (V)	14 – 86.6
PZT2 Actuation Range (V)	70.5
PZT2 Operating Voltage (V)	13.6 – 84.1

During this visit the OMC team came up with the procedures and python scripts that would be used in the coming months to build the other five OMC cavities. The build summary of these cavities is tabulated below

Table 4.10: Measured OMC parameters (OMC1–OMC3).

	OMC1	OMC2	OMC3
Finesse	390	375	380
OMC unit throughput	0.991 ± 0.005	0.9798 ± 0.0054	0.983 ± 0.002
OMC loss per mirror (ppm)	16 ± 11	43.49 ± 11.44	32 ± 10

Table 4.11: Measured OMC parameters (OMC4–OMC6).

	OMC4	OMC5	OMC6
Finesse	380	387	378
OMC unit throughput	0.988 ± 0.007	0.986 ± 0.008	0.9836 ± 0.0049
OMC loss per mirror (ppm)	23.25 ± 14.25	27 ± 17	31 ± 10

Polarisation Balanced Homodyne Detection

5.1 Introduction

Balanced Homodyne Detection (BHD) was introduced in the previous chapter, it is a method in which a local oscillator (LO) operating at the same optical frequency as the signal is used to demodulate the signal from optical to lower frequencies [73]. In the A+ LIGO configuration, BHD is implemented using two output mode cleaner (OMC) cavities per interferometer to spatially filter the signal field and the LO field independently, as discussed in Section 4.1.1. This dual-OMC architecture is a significant increase in system complexity relative to DC readout, and it introduces new pathways for technical noise to couple into the readout. It also amplifies practical challenges in hardware delivery: as shown in the previous chapter, constructing OMC cavities is itself non-trivial. Achieving >98% throughput reproducibly across multiple units is demanding because the OMC is a monolithic, bonded assembly with limited post-bond alignment freedom, and the build process requires substantial time, cleanliness control, and coordinated person-effort.

A central challenge is that the noise rejection of BHD relies on symmetry between the two detection paths (Figure 4.2). If the two OMCs are not closely matched in their optical response—e.g., in throughput, mode filtering, higher-order-mode spectrum, or alignment fluctuations—then the two photodetector signals (PD1 and PD2) no longer represent equal-amplitude measurements of the same optical quadrature. Such asymmetries degrade the subtraction that provides common-mode rejection, and can reintroduce coupling of LO noise and other technical fluctuations, thereby reducing the quantum-noise advantage that motivates BHD in the first place.

This chapter explores a polarisation-based approach that can mitigate these complications by enabling both the LO and the signal to be spatially filtered by a single OMC. The concept, originally proposed by Arai [31], assigns orthogonal polarisations to the two fields—taking the interferometer signal to be p-polarised and the LO to be s-polarised—so that they can be co-propagated through the same cavity while remaining separable at detection. This approach was investigated before [100], here we revisit this concept and develop the mathematical framework further by investigating the drawbacks of the model (Section 5.4). Two polarisation BHD, although simpler in principle than usual BHD, gives rise to an additional noise term due to birefringence in OMC cavity: cavity length fluctuations couple to readout signal. We will try to understand mathematically how this can be mitigated (Section 5.4.3) and discuss challenges.

In the subsequent sections we understand how this can be implemented in LIGO and develop a table top experimntal method to verify this model.

5.2 Birefringence in Optical Cavities

Birefringence is an anisotropic optical property of materials that exhibit different refractive index that depends on the polarisation and propagation direction of light. The cavity mirrors in OMCs are not made of birefringent material and are composed of fused silica substrates with high-reflectivity dielectric coatings (plus anti-reflective coatings on the rear surfaces). In optical cavities it is the interaction of the optical fields with these coating layers that causes birefringence.

To achieve the high reflectivity required for mirrors in gravitational wave interferometers, multi-layer dielectric coatings are employed [134]. These coatings are composed of alternating layers of materials with high and low refractive indices, deposited onto a substrate. Each interface between layers partially reflects the incident light, and the layer thicknesses are typically designed to be one quarter of the optical wavelength, ensuring constructive interference of the reflected waves. Increasing the number of layers enhances the overall reflectivity, and a larger contrast in refractive indices between adjacent layers improves the reflection efficiency at each interface. Birefringence in optical coatings can occur due to thermal stress between the multi-layered optical coating during the fabrication process or sub-assembly bonding process or, non-normal angle of incidence [51]. In an OMC the angle of incidence on the curved mirrors is 4 degrees which contributes to some level of birefringence. These factors cause the difference in reflectivities for s- and p- polarisations thus introduce mode splitting.

The frequency dependent transfer function of a cavity, like the OMC, in transmission can approximately be written as [[50]]

$$G(f) = \frac{1}{1 + if/\gamma} = \frac{e^{i\phi}}{1 + (f/\gamma)^2} = |G|e^{i\phi} \quad (5.1)$$

The difference in the transfer functions for s- and p- polarisations due to coating birefringence leads to different cavity parameters for the cavity. Since we want to resonate both s- and p- polarisations simultaneously, the cavity is locked between the two resonance peaks.




Figure 5.1: OMC Circulating Power: This plot shows exaggerated birefringence effects in an optical cavity. $\delta\mathcal{F}/\mathcal{F}$ is 0.1 and $\delta f_0/\gamma$ is 0.5. The measured cavity birefringence for the A+ OMC cavity number 6 are 0.025 (0.0005) and 0.0038 (0.0012) respectively.

A quarter-wave plate (QWP) with a general birefringence parameter η can be described using the following Jones matrix representation [[157]]

$$\mathbf{J}_{\text{QWP}} = \begin{bmatrix} e^{i\frac{\eta}{2}} & 0 \\ 0 & e^{-i\frac{\eta}{2}} \end{bmatrix}, \quad (5.2)$$

the fast axis of the QWP is horizontal here. Similarly, the birefringence effect in the OMC cavity can be modelled as a waveplate with a birefringence of $\Delta\phi$ in the following way

$$\mathbf{J}_{\text{OMC}} = \begin{bmatrix} G_p & 0 \\ 0 & G_s e^{i\Delta\phi} \end{bmatrix}, \quad (5.3)$$

where G_p and G_s are p- and s- polarisation transfer function magnitudes respectively. The Jones matrix of a HWP with fast axis at an angle θ with respect to the horizontal can be represented as [75]

$$\mathbf{J}_{\text{HWP}}(\theta) = e^{-i\frac{\pi}{2}} \begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix}, \quad (5.4)$$

Since the phase difference between the s- and p-polarised fields in the cavity is inherently relative, one of the resonant fields was chosen as the zero-phase reference. This approach avoids the need to unwrap the cavity geometry and track the cumulative phase accrued over multiple round trips. In the future sections we will be using these Jones matrices to build a model for the experiment to understand birefringence noise cancellation.

5.3 Revisiting BHD equations

We discussed about the nominal plan for BHD scheme in the previous chapter in Section 4.1.1. Revisiting the concept, a strong local oscillator (LO) field $E_{\text{LO}}e^{i\delta} = \sqrt{P_{\text{LO}}}e^{i(\delta+\omega t)}$ is combined at a beam splitter with the interferometer output signal $E_{\text{IFO}} = \sqrt{P_{\text{IFO}}}e^{i\omega t}$, where P_{LO} and P_{IFO} represent the powers of the local oscillator and interferometer beams, respectively, and ζ denotes the relative phase between the two fields. Assuming the LO power dominates, $P_{\text{LO}} \gg P_{\text{IFO}}$, the detected power on each photodiode (see Figure 4.2) is given by [100][73]:

$$PD_1 = \frac{1}{2} \left[P_{\text{LO}} - 2\sqrt{P_{\text{LO}}P_{\text{IFO}}} \left(e^{i\delta} + e^{-i\delta} \right) \right], \quad (5.5)$$

$$PD_2 = \frac{1}{2} \left[P_{\text{LO}} + 2\sqrt{P_{\text{LO}}P_{\text{IFO}}} \left(e^{i\delta} + e^{-i\delta} \right) \right]. \quad (5.6)$$

Taking the difference between the two photodiode signals yields an output proportional to the interferometer signal

$$P = \frac{PD_2 - PD_1}{2} = \sqrt{P_{\text{LO}}P_{\text{IFO}}} \cos \delta. \quad (5.7)$$

5.4 Two-Polarisation BHD

Before introducing the polarisation-based implementation, it is useful to restate the core requirements of balanced homodyne detection (BHD) that must also be satisfied by the polarisation BHD scheme:

- **Homodyne-angle control:** the ability to control the relative optical phase between the local oscillator (LO) and the interferometer signal field, thereby selecting the measured quadrature.
- **Field mixing:** coherent mixing of the LO and signal fields at a 50/50 beamsplitter (or an equivalent optical element) prior to balanced detection.

In the two-polarisation BHD concept, the LO is derived from a pick-off in the same manner as in the nominal BHD architecture, but its polarisation is rotated so that it is orthogonal to the interferometer output field. The relative LO–signal phase (i.e. the homodyne angle) is controlled using the same approach as in nominal BHD, via a dedicated phase shifter in the LO path. The two fields are then mixed using polarisation optics: a half-wave plate (HWP) with its fast axis set to 22.5° rotates each field by 45° , producing equal projections of the orthogonally polarised LO and signal onto a common polarisation basis. This operation is equivalent to a 50/50 mixing step and enables subsequent balanced detection.

This mixing can be implemented in two configurations as discussed in the following sections.

5.4.1 BHD with HWP before the OMC




Figure 5.2: Schematic of the A+ LIGO interferometer output with polarisation homodyne readout chain: A local oscillator (LO) field is picked off from the input beam path and reflected off of a phase shifter mirror and a HWP to control its phase angle and power relative to the . This LO field is then recombined with the signal field E_{IFO} at a PBS. The resulting beams are filtered through a single OMC to remove higher-order spatial modes and control sidebands. Balanced homodyne detection (BHD) is performed using a PBS. The red and yellow paths represent the LO and signal fields, respectively, with polarisation states shown schematically at various points.

The Jones representation for the light fields past the combining PBS in Figure 5.2 is shown below

$$-\frac{1}{\sqrt{2}} \begin{bmatrix} \text{OMC} & \\ G_p & 0 \\ 0 & G_s e^{i\Delta\phi} \end{bmatrix} \begin{bmatrix} \text{HWP FA angled } 22.5^\circ \text{ with horizontal} & \\ 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \text{Input} \\ E_{\text{IFO}} \\ E_{\text{LO}} e^{i\delta} \end{bmatrix} = \begin{bmatrix} E_1 \\ E_2 \end{bmatrix}$$

Here the OMC matrix is used from equation 5.3 and E_{IFO} and E_{LO} are interferometer and local oscillator fields as shown in Figure 5.2.

$$E = -\frac{1}{\sqrt{2}} \left[G_p (E_{\text{IFO}} + E_{\text{LO}} e^{i\delta}) \right] \hat{p} - \frac{1}{\sqrt{2}} \left[G_s e^{i\Delta\phi} (E_{\text{IFO}} - E_{\text{LO}} e^{i\delta}) \right] \hat{s} \quad (5.8)$$

Here \hat{p} and \hat{s} are $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ matrices respectively. The s- and p- fields here resonate in the OMC cavity simultaneously. Homodyning (equation 5.7) this signal gives:

$$P_1 - P_2 = \frac{1}{2} (G_s^2 - G_p^2) (P_{\text{IFO}} + P_{\text{LO}}) - (G_p^2 + G_s^2) \cos \delta \sqrt{P_{\text{IFO}} P_{\text{LO}}} \quad (5.9)$$

If we assume $G_p \approx G_s$ (when the birefringence is very small) then the equation simplifies to:

$$P = \frac{P_1 - P_2}{2} = \sqrt{P_{\text{LO}} P_{\text{IFO}}} \cos \zeta. \quad (5.10)$$

which is the BHD signal we saw in equation 5.7 as expected. Now let's estimate the noise spectral density associated with the readout.

The derivative of the cavity transfer function can be estimated as

$$\frac{\partial |G_s|^2}{\partial x} \approx -\frac{\delta f_0}{\gamma^2} \frac{FSR}{\lambda} \quad (5.11)$$

Now assuming [100]

$$\frac{\partial |G_s|^2}{\partial x} = -\frac{\partial |G_p|^2}{\partial x} \quad (5.12)$$

for small cavity birefringence, the partial derivative of equation 5.9 with respect to x (cavity length) gives the noise spectral density

$$S_P^{1/2} = \frac{P_{\text{IFO}} + P_{\text{LO}}}{2} \frac{f_{\text{FSR}}}{\lambda} \frac{\delta f_0}{\gamma^2} S_x^{1/2} \quad (5.13)$$

where f_{FSR} is free spectral range of the cavity, γ is the cavity pole and δf_0 is difference resonance frequencies for s- and p- fields due to cavity birefringence. The assumption in equation 5.12 holds true if the laser locking point is half way between the s- and p-resonances and the cavity birefringence is small. From equation 5.13 we see that the readout noise ASD due to OMC cavity length fluctuations depends on the difference in the resonance frequencies of the two polarisations due to the cavity birefringence.

5.4.2 BHD with HWP after the OMC

In this section we will see how BHD can be performed using a HWP after the OMC to mix the LO and IFO beams as opposed to before the OMC.




Figure 5.3: After the OMC, a half-wave plate (HWP) is inserted to rotate the polarisation angle by 45° prior to balanced homodyne detection (BHD), which is performed using a PBS and a pair of photodiodes (P1 and P2). Red and yellow lines represent the LO and signal paths, respectively, with polarisation states indicated schematically.

The Jones representation for the light fields past the combining PBS in Figure 5.3 is shown below:

$$-\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{matrix} \text{HWP FA angled } 22.5^\circ \text{ with horizontal} \\ \text{OMC} \end{matrix} \begin{bmatrix} G_p & 0 \\ 0 & G_s e^{i\Delta\phi} \end{bmatrix} \begin{matrix} \text{Input} \\ \begin{bmatrix} E_{\text{IFO}} \\ E_{\text{LO}} e^{i\delta} \end{bmatrix} \end{matrix} = \begin{bmatrix} E_1 \\ E_2 \end{bmatrix}$$

Here the OMC matrix is used from equation 5.3 and E_{IFO} and E_{LO} are interferometer and local oscillator fields as shown in Figure 5.3.

Performing matrix multiplication gives:

$$E = -\frac{1}{\sqrt{2}} \left[G_p E_{IFO} + G_s E_{LO} e^{i(\Delta\phi + \delta)} \right] \hat{p} - \frac{1}{\sqrt{2}} \left[G_p E_{IFO} - G_s E_{LO} e^{i(\Delta\phi + \delta)} \right] \hat{s}$$

Here \hat{p} and \hat{s} are $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ matrices respectively. Homodyning gives:

$$\frac{P_1 - P_2}{2} = 2G_p G_s \sqrt{P_{LO} P_{IFO}} \cos(\delta + \Delta\phi) \quad (5.14)$$

This is of the form of homodyne readout signal in Equation 5.7. Now because

$$\frac{\partial \Delta\phi}{\partial x} = \frac{4\delta \mathcal{F}}{\lambda}, \quad (5.15)$$

taking partial derivative of equation 5.14 with respect to x (cavity length) and considering $\Delta\phi \sim 0$ gives noise spectral density in the homodyne readout as

$$S_p^{1/2} = \sqrt{P_{IFO} P_{LO}} \frac{2\delta \mathcal{F}}{\lambda} |\sin \delta| S_x^{1/2} \quad (5.16)$$

where δ is homodyne angle introduced by the phase shifter, $\delta \mathcal{F}$ is difference in finesse for s- and p- fields due to cavity birefringence. In equation 5.16 we see how the noise ASD for the BHD readout related to OMC cavity length fluctuations through the finesse difference for the two polarisations. This is again caused by the birefringence in the cavity.

From equations 5.16 and 5.13 we see that using HWPs before or after to perform mixing of the two beams gives rise to noise in the readout signal due to birefringence in the cavity. But this can be mitigated as discussed in the next subsection. Noticing that the mixing of the two fields due to rotation caused by the HWP (either before or after the OMC) can also be done using two HWPs, one before and one after such that we balance the rotation before and after the OMC. This enables cancellation of birefringence-induced noise while preserving net 45° rotation.

5.4.3 Simultaneous Half-Wave Plate Mixing - Cancelling Birefringence Noise

As we saw in the previous subsections, if we inject LO and IFO beams at orthogonal polarisations, we can perform BHD by mixing the two fields using a half wave plate whose fast axis is at 22.5° with respect to either polarisations. This enables a rotation of the polarisations by 45° . But this gives rise to a problem, this leads to coupling of length noise to the homodyne signal due to birefringence in the cavity. To perform BHD we want to be able to mix the two polarisations, this can be done using two half wave plates

$$\frac{1}{\sqrt{2}} \begin{matrix} \text{HWP}_2 \text{ FA at } 22.5^\circ - \theta \\ \begin{bmatrix} \cos 2\theta - \sin 2\theta & \cos 2\theta + \sin 2\theta \\ \sin 2\theta + \cos 2\theta & \sin 2\theta - \cos 2\theta \end{bmatrix} \end{matrix} \begin{matrix} \text{HWP}_1 \text{ FA at } \theta \\ \begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix} \end{matrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \quad (5.17)$$

The right hand side represents matrix of a rotation matrix, rotating the polarisation states by 45° . The matrix is also independent of θ as expected. This means we can use two HWPs, one before and one after the OMC and tune the angle θ such that the birefringence noise due to partial mixing using HWP_1 can be cancelled by that due to partial mixing using HWP_2 . The Jones representation of half wave plate mixing of the orthogonally polarised LO and IFO beams is given by:

$$\begin{array}{ccccccc}
& \text{HWP}_2 \text{ FA at } 22.5^\circ - \theta & & \text{OMC} & & \text{HWP}_1 \text{ FA at } \theta & & \text{Input} & & \text{Output} \\
\frac{1}{\sqrt{2}} \begin{bmatrix} \cos 2\theta - \sin 2\theta & \cos 2\theta + \sin 2\theta \\ \sin 2\theta + \cos 2\theta & \sin 2\theta - \cos 2\theta \end{bmatrix} & & & \begin{bmatrix} G_p & 0 \\ 0 & G_s e^{i\Delta\phi} \end{bmatrix} & & & & \begin{bmatrix} E_{\text{IFO}} \\ E_{\text{LO}} e^{i\delta} \end{bmatrix} & = & \begin{bmatrix} E_1 \\ E_2 \end{bmatrix} \\
& & & & & & & & & (5.18)
\end{array}$$

where $\Delta\phi$ is phase difference between s- and p- polarisations introduced by the OMC, δ is the homodyne angle.

Now we want to compute P_1 and P_2 and estimate $P_2 - P_1$ which is the BHD signal. We can then estimate $\frac{\partial P}{\partial x}$ to evaluate the effective noise contribution.

$$\begin{aligned}
E_1 = \frac{1}{\sqrt{2}} & \left[(\cos(2\theta) - \sin(2\theta))G_p \cos(2\theta) + (\cos(2\theta) + \sin(2\theta))G_s e^{i\Delta\phi} \sin(2\theta) \right] E_{\text{IFO}} \\
& + \frac{1}{\sqrt{2}} \left[(\cos(2\theta) - \sin(2\theta))G_p \sin(2\theta) - (\cos(2\theta) + \sin(2\theta))G_s e^{i\Delta\phi} \cos(2\theta) \right] E_{\text{LO}} e^{i\delta}.
\end{aligned}$$

This represents signal on the first photodiode.

$$|E_1|^2 = |A|^2 + |B|^2 + 2\text{Re}(A\bar{B})$$

where

$$A = \left[(\cos(2\theta) - \sin(2\theta))G_p \cos(2\theta) + (\cos(2\theta) + \sin(2\theta))G_s e^{i\Delta\phi} \sin(2\theta) \right] E_{\text{IFO}}$$

$$B = \left[(\cos(2\theta) - \sin(2\theta))G_p \sin(2\theta) - (\cos(2\theta) + \sin(2\theta))G_s e^{i\Delta\phi} \cos(2\theta) \right] E_{\text{LO}} e^{i\delta}.$$

Now evaluating these separately and substituting them in the above equation gives

$$|A|^2 = \left[(1 - \sin(4\theta))G_p^2 \cos^2(2\theta) + (1 + \sin(4\theta))G_s^2 \sin^2(2\theta) + \frac{\sin(8\theta)}{2}G_pG_s \cos(\Delta\phi) \right] P_{\text{IFO}}$$

$$|B|^2 = \left[(1 - \sin(4\theta))G_p^2 \sin^2(2\theta) + (1 + \sin(4\theta))G_s^2 \cos^2(2\theta) - \frac{\sin(8\theta)}{2}G_pG_s \cos(\Delta\phi) \right] P_{\text{LO}}$$

$$\begin{aligned} 2\text{Re}(A\bar{B}) &= \left\{ \left[(1 - \sin(4\theta))G_p^2 \frac{\sin(4\theta) \cos(\delta)}{2} \right] \right. \\ &\quad - \cos(4\theta)G_pG_s \cos^2(2\theta) \cos(\delta + \Delta\phi) \\ &\quad + \cos(4\theta) \sin^2(2\theta)G_pG_s \cos(\delta - \Delta\phi) \\ &\quad \left. - \left[(1 + \sin(4\theta))G_s^2 \frac{\sin(4\theta) \cos(\delta)}{2} \right] \right\} \\ &\quad \times 2\sqrt{P_{\text{IFO}}P_{\text{LO}}}. \end{aligned} \quad (5.19)$$

$$\begin{aligned} P_1 &= \frac{1}{2} \left[|A|^2 + |B|^2 + 2\text{Re}(A\bar{B}) \right] \\ &= \frac{1}{2} \left[\left((1 - \sin(4\theta))G_p^2 \cos^2(2\theta) + (1 + \sin(4\theta))G_s^2 \sin^2(2\theta) + \frac{\sin(8\theta)}{2}G_pG_s \cos(\Delta\phi) \right) P_{\text{IFO}} \right. \\ &\quad + \left((1 - \sin(4\theta))G_p^2 \sin^2(2\theta) + (1 + \sin(4\theta))G_s^2 \cos^2(2\theta) - \frac{\sin(8\theta)}{2}G_pG_s \cos(\Delta\phi) \right) P_{\text{LO}} \\ &\quad + \left\{ (1 - \sin(4\theta))G_p^2 \frac{\sin(4\theta) \cos(\delta)}{2} - \cos(4\theta)G_pG_s \cos^2(2\theta) \cos(\delta + \Delta\phi) \right. \\ &\quad \left. + \cos(4\theta)G_pG_s \sin^2(2\theta) \cos(\delta - \Delta\phi) - \frac{(1 + \sin(4\theta))G_s^2 \sin(4\theta) \cos(\delta)}{2} \right\} \cdot 2\sqrt{P_{\text{IFO}}P_{\text{LO}}} \left. \right] \end{aligned} \quad (5.20)$$

Similarly performing matrix multiplication for E_2 gives

$$E_2 = \left[(\cos(2\theta) + \sin(2\theta))G_p \cos(2\theta) + (\sin(2\theta) - \cos(2\theta))G_s \sin(2\theta) e^{i\Delta\phi} \right] E_{\text{IFO}}$$

$$+ \left[(\cos(2\theta) + \sin(2\theta))G_p \sin(2\theta) - (\sin(2\theta) - \cos(2\theta)) \cos(2\theta)G_s e^{i\Delta\phi} \right] E_{\text{LO}} e^{i\delta}.$$

This gives

$$\begin{aligned}
P_2 = \frac{1}{2} & \left[\left((1 + \sin(4\theta))G_p^2 \cos^2(2\theta) + (1 - \sin(4\theta))G_s^2 \sin^2(2\theta) \right. \right. \\
& \quad \left. \left. - \frac{\sin(8\theta)}{2} G_p G_s \cos(\Delta\phi) \right) P_{\text{IFO}} \right. \\
& + \left((1 + \sin(4\theta))G_p^2 \sin^2(2\theta) + (1 - \sin(4\theta))G_s^2 \cos^2(2\theta) \right. \\
& \quad \left. + \frac{\sin(8\theta)}{2} G_p G_s \cos(\Delta\phi) \right) P_{\text{LO}} \\
& + \left\{ \frac{(1 + \sin(4\theta))G_p^2 \sin(4\theta) \cos(\delta)}{2} + \cos(4\theta)G_p G_s \cos^2(2\theta) \cos(\delta + \Delta\phi) \right. \\
& \quad \left. - \cos(4\theta)G_p G_s \sin^2(2\theta) \cos(\delta - \Delta\phi) \right. \\
& \quad \left. - \frac{(1 - \sin(4\theta))G_s^2 \sin(4\theta) \cos(\delta)}{2} \right\} 2\sqrt{P_{\text{IFO}}P_{\text{LO}}} \Big] \tag{5.21}
\end{aligned}$$

Now when $|G_p|$ and $|G_s| \approx 1$ and $\Delta\phi$ that is the relative phase shift introduced by the cavity is zero, the equations for P_1 and P_2 become

$$P = 2\sqrt{P_{\text{LO}}P_{\text{IFO}}} \cos \delta \tag{5.22}$$

as expected.

We can then compute the length fluctuations in P_1 and P_2 by differentiating with respect to the cavity length and using equation 5.12, this gives

$$\begin{aligned}
\frac{\partial P_1}{\partial x} &= \left(\frac{P_{\text{LO}} - P_{\text{IFO}}}{2} \right) \left(\cos 4\theta \frac{\partial G_s^2}{\partial x} + \frac{\sin 8\theta}{2} G_p G_s \sin \Delta\phi \frac{\partial \Delta\phi}{\partial x} \right) \\
&\quad + \left(\frac{P_{\text{LO}} + P_{\text{IFO}}}{2} \right) \sin 4\theta \frac{\partial G_s^2}{\partial x} \\
&\quad + \frac{1}{2} \left\{ -\sin 4\theta \cos \delta \frac{\partial G_s^2}{\partial x} + \frac{\partial \Delta\phi}{\partial x} \cos 4\theta G_p G_s \right. \\
&\quad \times (\cos^2 2\theta \sin(\delta + \Delta\phi) + \sin^2 2\theta \sin(\delta - \Delta\phi)) \left. \right\} \\
&\quad \times 2\sqrt{P_{\text{IFO}}P_{\text{LO}}}
\end{aligned} \tag{5.23}$$

and

$$\begin{aligned}
\frac{\partial P_2}{\partial x} &= \left(\frac{P_{\text{LO}} - P_{\text{IFO}}}{2} \right) \left(\cos 4\theta \frac{\partial G_s^2}{\partial x} - \frac{\sin 8\theta}{2} G_p G_s \sin \Delta\phi \frac{\partial \Delta\phi}{\partial x} \right) \\
&\quad - \left(\frac{P_{\text{LO}} + P_{\text{IFO}}}{2} \right) \sin 4\theta \frac{\partial G_s^2}{\partial x} \\
&\quad + \frac{1}{2} \left\{ -\sin 4\theta \cos \delta \frac{\partial G_s^2}{\partial x} - \frac{\partial \Delta\phi}{\partial x} \cos 4\theta G_p G_s \right. \\
&\quad \times (\cos^2 2\theta \sin(\delta + \Delta\phi) + \sin^2 2\theta \sin(\delta - \Delta\phi)) \left. \right\} \\
&\quad \times 2\sqrt{P_{\text{IFO}}P_{\text{LO}}}
\end{aligned} \tag{5.24}$$

The above equations are generic and for any value of G_p, G_s and $\Delta\phi$. Now substituting $\Delta\phi = 0, G_p, G_s = 1$, the equations simplify to,

$$\begin{aligned}
2\frac{\partial P_1}{\partial x} &= \left[(P_{\text{LO}} - P_{\text{IFO}}) \cos 4\theta + (P_{\text{LO}} + P_{\text{IFO}}) \sin 4\theta - \sin 4\theta \cos \delta \sqrt{P_{\text{IFO}}P_{\text{LO}}} \right] \frac{\partial G_s^2}{\partial x} \\
&\quad + \left[\sqrt{P_{\text{IFO}}P_{\text{LO}}} \cos 4\theta \sin \delta \right] \frac{\partial \Delta\phi}{\partial x} \\
2\frac{\partial P_2}{\partial x} &= \left[(P_{\text{LO}} - P_{\text{IFO}}) \cos 4\theta - (P_{\text{LO}} + P_{\text{IFO}}) \sin 4\theta - \sin 4\theta \cos \delta \sqrt{P_{\text{IFO}}P_{\text{LO}}} \right] \frac{\partial G_s^2}{\partial x} \\
&\quad - \left[\sqrt{P_{\text{IFO}}P_{\text{LO}}} \cos 4\theta \sin \delta \right] \frac{\partial \Delta\phi}{\partial x}
\end{aligned}$$

Subtracting them gives the length fluctuations in the homodyne signal

$$\frac{\partial P_1 - P_2}{\partial x} = [(P_{\text{LO}} + P_{\text{IFO}}) \sin 4\theta] \frac{\partial G_s^2}{\partial x} + \left[\sqrt{P_{\text{IFO}} P_{\text{LO}}} \cos 4\theta \sin \delta \right] \frac{\partial \Delta\phi}{\partial x} \quad (5.25)$$

Now using equations 5.11 and 5.15 gives the total birefringence noise

$$\frac{\partial \frac{P_1 - P_2}{2}}{\partial x} = -(P_{\text{LO}} + P_{\text{IFO}}) \sin 4\theta \cdot \frac{\delta f_0 f_{\text{FSR}}}{\gamma^2 \lambda} + \sqrt{P_{\text{IFO}} P_{\text{LO}}} \cos 4\theta \sin \delta \cdot \frac{4\delta \mathcal{F}}{\lambda} \quad (5.26)$$

This equation is the total birefringence noise in the homodyne signal. This looks like the sum of equations 5.16 and 5.13 but projected into different quadratures.

Equating this to zero gives:

$$\theta_{\text{opt}} = -\frac{1}{4} \arctan \left(\sqrt{\frac{P_{\text{IFO}}}{P_{\text{LO}}}} \cdot \frac{4\delta \mathcal{F}}{f_0} \cdot \frac{\gamma^2}{\delta f_0} \cdot \sin \delta \right) \quad (5.27)$$

This implementation for LIGO is illustrated in Figure 5.4. Equation 5.27 shows that the optimum half-wave-plate balancing angles depend on both the cavity birefringence and the chosen homodyne angle. It is therefore necessary to characterise the cavity birefringence in order to determine the expected operating angles and to benchmark the experimental settings against the theoretical optimum.

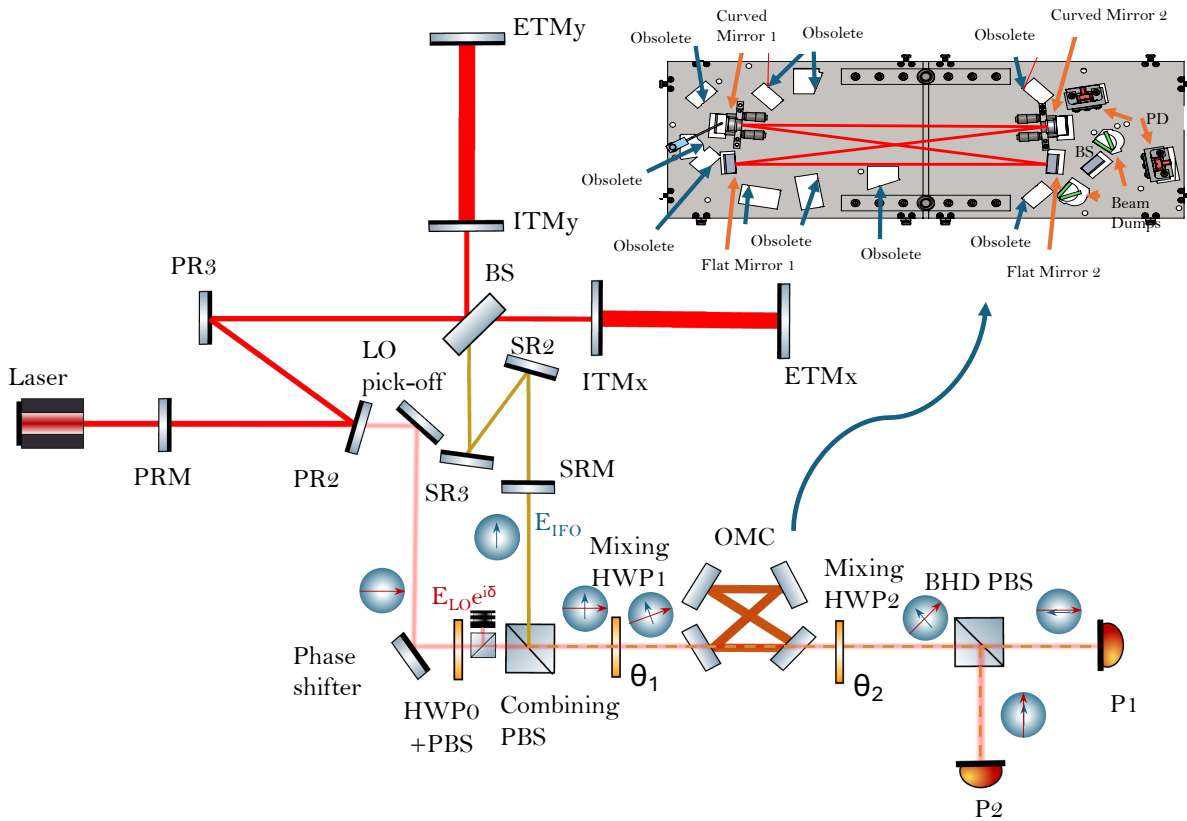


Figure 5.4: Simplified schematic of the A+ LIGO interferometer output with polarisation homodyne readout chain: The main interferometer consists of two orthogonal arm cavities formed by input (ITM) and end test masses (ETM), with the beamsplitter (BS) at the center. The power recycling cavity, formed by PRM, PR2, and PR3, enhances circulating laser power, while the signal recycling cavity, formed by SRM, SR2, and SR3, shapes the frequency response of the interferometer. A local oscillator (LO) field is picked off from the input beam and directed through a phase shifter mirror and a power-adjusting HWP-PBS combination before being combined with the signal field E_{IFO} at a polarizing beamsplitter (PBS). The combined beam is filtered by a single output mode cleaner (OMC), which suppresses higher-order spatial modes and unwanted sidebands. Before and after the OMC, half-wave plates (HWP) are inserted to rotate the polarisation angle by 45° , combined, prior to balanced homodyne detection (BHD), which is performed using a PBS and a pair of photodiodes (P1 and P2). Red and yellow lines represent the LO and signal paths, respectively, with polarisation states indicated schematically.

5.4.4 Investigating Discrepancy

In the plot in Figure 5.13 we observe discrepancy between the measured and the predicted optimum mixing angles. So we generalise the equations (5.23) and (5.24) further to investigate this. Equations (5.23) and (5.24) can be written (by disregarding the assumption in equation 5.12) more generally as

$$\begin{aligned}
2 \frac{\partial P_1}{\partial x} = & \left[\begin{aligned} & (1 - \sin 4\theta) \cos^2(2\theta) \frac{\partial G_p^2}{\partial x} \\ & + (1 + \sin 4\theta) \sin^2(2\theta) \frac{\partial G_s^2}{\partial x} \\ & + \frac{\sin 8\theta}{4} \left(\frac{\partial G_s^2}{\partial x} + \frac{\partial G_p^2}{\partial x} \right) \end{aligned} \right] P_{\text{IFO}} \\
& + \left[\begin{aligned} & (1 - \sin 4\theta) \sin^2(2\theta) \frac{\partial G_p^2}{\partial x} \\ & + (1 + \sin 4\theta) \cos^2(2\theta) \frac{\partial G_s^2}{\partial x} \\ & - \frac{\sin 8\theta}{4} \left(\frac{\partial G_p^2}{\partial x} + \frac{\partial G_s^2}{\partial x} \right) \end{aligned} \right] P_{\text{LO}} \\
& + 2 \left[\begin{aligned} & \frac{(1 - \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_p^2}{\partial x} \\ & - \frac{\cos^2(4\theta) \cos \delta}{2} \left(\frac{\partial G_p^2}{\partial x} + \frac{\partial G_s^2}{\partial x} \right) \\ & - \frac{(1 + \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_s^2}{\partial x} \\ & + \cos 4\theta \sin \delta \frac{\partial \Delta\phi}{\partial x} \end{aligned} \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}}.
\end{aligned} \tag{5.28}$$

$$\begin{aligned}
2 \frac{\partial P_2}{\partial x} = & \left[\begin{aligned} & (1 + \sin 4\theta) \cos^2(2\theta) \frac{\partial G_p^2}{\partial x} \\ & + (1 - \sin 4\theta) \sin^2(2\theta) \frac{\partial G_s^2}{\partial x} \\ & - \frac{\sin 8\theta}{4} \left(\frac{\partial G_s^2}{\partial x} + \frac{\partial G_p^2}{\partial x} \right) \end{aligned} \right] P_{\text{IFO}} \\
& + \left[\begin{aligned} & (1 + \sin 4\theta) \sin^2(2\theta) \frac{\partial G_p^2}{\partial x} \\ & + (1 - \sin 4\theta) \cos^2(2\theta) \frac{\partial G_s^2}{\partial x} \\ & + \frac{\sin 8\theta}{4} \left(\frac{\partial G_s^2}{\partial x} + \frac{\partial G_p^2}{\partial x} \right) \end{aligned} \right] P_{\text{LO}} \\
& + 2 \left[\begin{aligned} & \frac{(1 + \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_p^2}{\partial x} \\ & + \frac{\cos^2(4\theta) \cos \delta}{2} \left(\frac{\partial G_p^2}{\partial x} + \frac{\partial G_s^2}{\partial x} \right) \\ & - \frac{(1 - \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_s^2}{\partial x} \\ & - \cos 4\theta \sin \delta \frac{\partial \Delta\phi}{\partial x} \end{aligned} \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}}.
\end{aligned} \tag{5.29}$$

Factorising gives

$$\begin{aligned}
2 \frac{\partial P_1}{\partial x} = & \frac{1}{2} \left[\begin{array}{l} (1 - \sin 4\theta) \frac{\partial G_p^2}{\partial x} \\ + (1 + \sin 4\theta) \frac{\partial G_s^2}{\partial x} \end{array} \right] (P_{\text{IFO}} + P_{\text{LO}}) \\
& + \frac{1}{2} \left[\cos 4\theta \left(\frac{\partial G_p^2}{\partial x} - \frac{\partial G_s^2}{\partial x} \right) \right] (P_{\text{IFO}} - P_{\text{LO}}) \\
& + 2 \left[\begin{array}{l} \frac{(1 - \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_p^2}{\partial x} \\ - \frac{\cos^2(4\theta) \cos \delta}{2} \left(\frac{\partial G_p^2}{\partial x} + \frac{\partial G_s^2}{\partial x} \right) \\ - \frac{(1 + \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_s^2}{\partial x} \\ + \cos 4\theta \sin \delta \frac{\partial \Delta\phi}{\partial x} \end{array} \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}}. \quad (5.30)
\end{aligned}$$

$$\begin{aligned}
2 \frac{\partial P_2}{\partial x} = & \frac{1}{2} \left[\begin{array}{l} (1 + \sin 4\theta) \frac{\partial G_p^2}{\partial x} \\ + (1 - \sin 4\theta) \frac{\partial G_s^2}{\partial x} \end{array} \right] (P_{\text{IFO}} + P_{\text{LO}}) \\
& + \frac{1}{2} \left[\cos 4\theta \left(\frac{\partial G_p^2}{\partial x} - \frac{\partial G_s^2}{\partial x} \right) \right] (P_{\text{IFO}} - P_{\text{LO}}) \\
& + 2 \left[\begin{array}{l} \frac{(1 + \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_p^2}{\partial x} \\ + \frac{\cos^2(4\theta) \cos \delta}{2} \left(\frac{\partial G_p^2}{\partial x} + \frac{\partial G_s^2}{\partial x} \right) \\ - \frac{(1 - \sin 4\theta) \sin 4\theta}{2} \cos \delta \frac{\partial G_s^2}{\partial x} \\ - \cos 4\theta \sin \delta \frac{\partial \Delta\phi}{\partial x} \end{array} \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}}. \quad (5.31)
\end{aligned}$$

For small detuning $\delta f_0/\gamma \ll 1$ and small birefringence $|\beta| \ll 1$,

$$\frac{\partial |G_s|^2}{\partial x} \approx \frac{\delta f_0}{\gamma^2} \frac{\text{FSR}}{\lambda} \equiv A_0, \quad (5.32)$$

$$\frac{\partial |G_p|^2}{\partial x} \approx -\frac{2\delta f_0}{\gamma^2} (1 - \varepsilon)(1 + 2\beta) \frac{\text{FSR}}{\lambda} = -\kappa A_0, \quad \kappa \equiv 2(1 - \varepsilon)(1 + 2\beta). \quad (5.33)$$

Here ε indicated the lock point, $\varepsilon = 0$ if the lock point is on the -p resonance and $\varepsilon = 1$ for the lock point on -s resonance.

Substituting (5.32)–(5.33) into (5.30)–(5.31) gives

$$\begin{aligned}
2 \frac{\partial P_1}{\partial x} = & \frac{A_0}{2} \left\{ [(1 + \sin 4\theta) - \kappa(1 - \sin 4\theta)](P_{\text{IFO}} + P_{\text{LO}}) - \cos 4\theta (\kappa + 1)(P_{\text{IFO}} - P_{\text{LO}}) \right\} \\
& - A_0 \cos \delta \left[(1 - \sin 4\theta) \sin 4\theta \kappa + \cos^2(4\theta) (1 - \kappa) + (1 + \sin 4\theta) \sin 4\theta \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}} \\
& + 2 \cos 4\theta \sin \delta \frac{\partial \Delta \phi}{\partial x} \sqrt{P_{\text{LO}} P_{\text{IFO}}}.
\end{aligned} \tag{5.34}$$

$$\begin{aligned}
2 \frac{\partial P_2}{\partial x} = & \frac{A_0}{2} \left\{ [(1 - \sin 4\theta) - \kappa(1 + \sin 4\theta)](P_{\text{IFO}} + P_{\text{LO}}) - \cos 4\theta (\kappa + 1)(P_{\text{IFO}} - P_{\text{LO}}) \right\} \\
& + A_0 \cos \delta \left[\cos^2(4\theta) (1 - \kappa) - (1 + \sin 4\theta) \sin 4\theta \kappa - (1 - \sin 4\theta) \sin 4\theta \right] \sqrt{P_{\text{LO}} P_{\text{IFO}}} \\
& - 2 \cos 4\theta \sin \delta \frac{\partial \Delta \phi}{\partial x} \sqrt{P_{\text{LO}} P_{\text{IFO}}}.
\end{aligned} \tag{5.35}$$

$$\begin{aligned}
2 \left(\frac{\partial P_1}{\partial x} - \frac{\partial P_2}{\partial x} \right) = & A_0(1 + \kappa) \sin(4\theta) (P_{\text{IFO}} + P_{\text{LO}}) - 2A_0(1 - \kappa) \cos \delta \sqrt{P_{\text{LO}} P_{\text{IFO}}} \\
& + 4 \cos(4\theta) \sin \delta \frac{\partial \Delta \phi}{\partial x} \sqrt{P_{\text{LO}} P_{\text{IFO}}}.
\end{aligned} \tag{5.36}$$

$$A_0 \equiv \frac{\delta f_0}{\gamma^2} \frac{\text{FSR}}{\lambda}, \quad \kappa \equiv 2(1 - \varepsilon)(1 + 2\beta).$$

$$A_0(1 + \kappa)(P_{\text{IFO}} + P_{\text{LO}}) \sin(4\theta) + 4 \sin \delta \frac{\partial \Delta \phi}{\partial x} \sqrt{P_{\text{LO}} P_{\text{IFO}}} \cos(4\theta) = 2A_0(1 - \kappa) \cos \delta \sqrt{P_{\text{LO}} P_{\text{IFO}}}.$$

Defining

$$A \equiv A_0(1 + \kappa)(P_{\text{IFO}} + P_{\text{LO}}), \quad B \equiv 4 \sin \delta \frac{\partial \Delta \phi}{\partial x} \sqrt{P_{\text{LO}} P_{\text{IFO}}}, \quad D \equiv 2A_0(1 - \kappa) \cos \delta \sqrt{P_{\text{LO}} P_{\text{IFO}}}.$$

Then $A \sin(4\theta) + B \cos(4\theta) = D$. Let $R = \sqrt{A^2 + B^2}$ and $\phi = \text{atan2}(B, A)$. Since $A \sin x + B \cos x = R \sin(x + \phi)$, we obtain

$$\boxed{4\theta + \phi = \sin^{-1}\left(\frac{D}{R}\right) + 2\pi n \quad \text{or} \quad 4\theta + \phi = \pi - \sin^{-1}\left(\frac{D}{R}\right) + 2\pi n, \quad n \in \mathbb{Z}}$$

i.e.

$$\boxed{\theta = \frac{1}{4} \left[\sin^{-1}\left(\frac{D}{R}\right) - \phi \right] + \frac{\pi}{2} n \quad \text{or} \quad \theta = \frac{1}{4} \left[\pi - \sin^{-1}\left(\frac{D}{R}\right) - \phi \right] + \frac{\pi}{2} n}$$

(valid when $|D| \leq R$).

5.5 OMC6 Birefringence Measurement




Figure 5.5: Experimental setup to measure cavity birefringence: Schematic of the optical setup used for characterising the output mode cleaner (OMC) on the transport fixture. A 1064 nm Innolight laser (output power ~ 60 mW) is passed through a Faraday isolator to prevent back-reflections, and then phase modulated at 32 MHz using a resonant electro-optic modulator (EOM) for Pound-Drever-Hall (PDH) locking. The beam is mode-matched into the OMC using a telescope consisting of two plano-convex lenses ($f = 35$ mm and $f = 125$ mm). A PBS shown downstream the EOM cleans the polarisation entering the cavity. A BS directs a small fraction of light to a power reference photodiode for monitoring input power. The main beam is passes through a PBS to clean the polarisation. The input polarisation is adjusted at 45 degree angle using a combination of a half-wave plate (HWP_0) and a quarter-wave plate (QWP). After interacting with the OMC cavity, the reflected light is directed to a reflection photodetector and a CCD camera for PDH locking and alignment diagnostics, respectively. The transmitted beam is analyzed using a balanced homodyne detection (BHD) setup comprising a polarizing beam splitter (PBS) and two photodiodes (PD_1 and PD_2). The OMC cavity optics are mounted on a glass breadboard which is held in a dedicated transport fixture to maintain mechanical stability and optical alignment.

The experimental setup to measure the birefringence of the optical cavity is shown in Figure 5.5. Equal powers of s- and p- polarisation light fields were injected into the cavity and cavity was scanned across multiple free spectral ranges.

For the purposes of comparing the difference between the s- and p- polarisation cavity transfer functions, it is useful to define these transfer functions in terms of Lorentzians. Let's define cavity transfer function for p polarised beam by a Lorentzian of the form

$$L_p(f; f_0, \gamma, A) = \frac{A}{1 + \left(\frac{f-f_0}{\gamma}\right)^2} \quad (5.37)$$

where γ is cavity pole. If the frequencies are normalised to γ then we have

$$L_p(v; v_0, A) = \frac{A}{1 + v^2} \quad (5.38)$$

Then due to the birefringence effects described in the previous section, the cavity transfer function for s - polarised beam can be written as:

$$L_s(v; v_0 + \delta v_0, A + \delta A) = \frac{A + \delta A}{1 + (1 + \beta)^2(v + \delta v_0)^2} \quad (5.39)$$

where β is cavity finesse difference normalised to p- polarised cavity finesse, δv_0 corresponds to difference in the resonance frequencies and δA corresponds to difference in amplitude for s- and p- polarised beams. The difference between the s- and p- polarisation transfer functions can thus be written as

$$L_p(v; v_0, A) - L_s(v; v_0 + \delta v_0, A + \delta A) = \frac{A}{1 + v^2} - \frac{A + \delta A}{1 + (1 + \beta)^2(v + \delta v_0)^2} \quad (5.40)$$

And the cavity birefringence can be measured by fitting this to $PD_1 - PD_2$ (see Figure 5.7).




Figure 5.6: Normalised p-polarisation cavity resonances. Each trace shows the photodiode signal for a single resonance window after converting the fitted drive voltage to a normalised detuning $f = (V - V_0)/\gamma$. γ here was extracted by calculating the difference in voltage at half maximum of the resonance peak for every peak. The data used in the fitted using Lorentzians $L(f) = A/(1 + f^2)$. The dashed line at $f = 0$ marks the fitted resonance centre. The per-peak voltage-domain normalisation collapses distinct scans onto a common lineshape suitable for averaging and comparison.




Figure 5.7: Differential signal fit $L_p - L_s$. The measured differential photodiode signal versus the normalised detuning f . The scatter is fit with offset $+ \frac{A}{1 + f^2} - \frac{A + \delta A}{1 + [(1 + \beta)(f - \delta f_0)]^2}$, where δf_0 is the relative p-s resonance offset, β encodes the linewidth (finesse) asymmetry, and δA captures contrast (amplitude) mismatch. This parameterisation separates resonance spacing, finesse difference, and contrast into orthogonal fit knobs for subsequent statistical analysis. The data is averaged over 54 peaks.

The finesse of the cavity is 378. The difference in the finesse between s- and p- polarisations is $\Delta\mathcal{F} = 9.741$ and difference in resonance frequencies of the two polarisations is $\delta f_0 = 1.36$ kHz. The statistics of the cavity birefringence parameters are summarised with histograms in Figure 5.8.

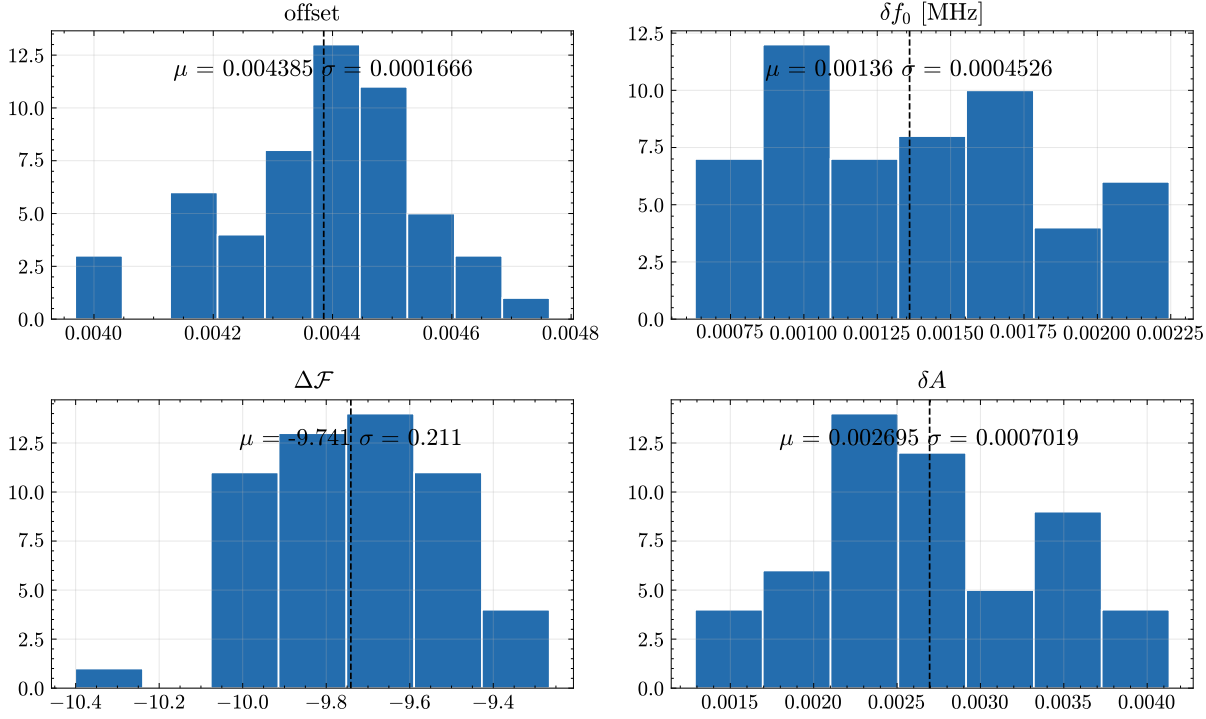


Figure 5.8: Distributions of birefringence and lineshape parameters from OMC cavity scans. For each resonance window the p-peak is fit in voltage (see Figure 5.6) to obtain (V_0, γ, A) , and the differential model $\text{offset} + \frac{A}{1+f^2} - \frac{A+\delta A}{1+[(1+\beta)(f-\delta f_0)]^2}$ (see Figure 5.7) yields (clockwise from top left) the DC *offset* of the differential channel, the polarisation resonance separation δf_0 (reported in MHz via the cavity calibration), the finesse difference $\Delta\mathcal{F} = \beta \mathcal{F}_0$ with $\mathcal{F}_0 = 378$, and the contrast mismatch δA . Each panel annotates the sample mean μ and standard deviation σ of the inlier set, as printed on the plots.

5.6 Measuring the Optimal Mixing Angles for OMC6

As shown in Figure 5.4, realising a tabletop polarisation BHD requires control of the relative phase (homodyne angle) between the local-oscillator (LO) and interferometer (IFO) fields. In our bench setup this control is implemented using a cascaded half-wave plate (HWP) and quarter-wave plate (QWP), as indicated. We first located the fast axes

and verified the performance of each waveplate, and aligned the polarizing beam splitters (PBSs) to a common plane (Appendix B). To account for any upstream phase imparted by preceding optics, the waveplates were then characterized sequentially in the order they appear in Figure 5.9 (Appendix B.1). This procedure ensures that any additional phase shifts introduced earlier in the train are incorporated into the final homodyne-angle calibration.

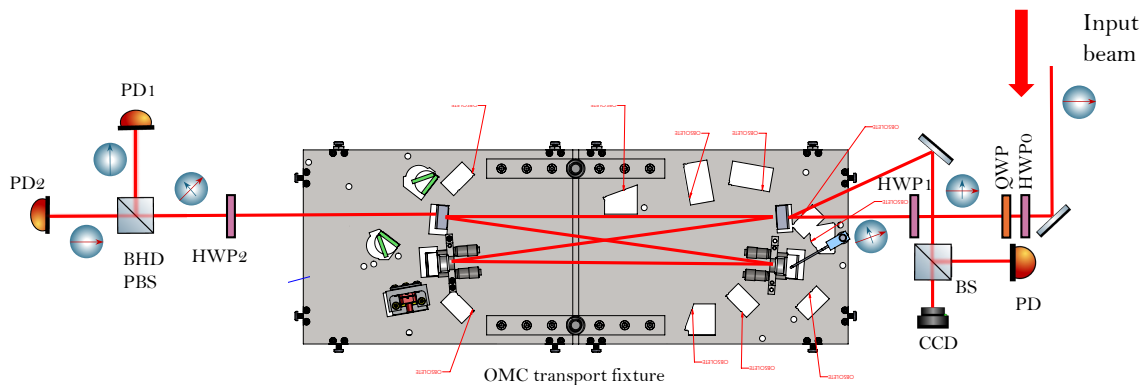


Figure 5.9: Tabletop setup for polarisation BHD The input laser setup is same as that used for birefringence measurement in Figure 5.5. Waveplate stack HWP₀–QWP provides control of the LO–IFO relative phase (homodyne angle) and also helps maintain a similar power ratio between the two orthogonal polarisations across measurements. Downstream, HWP₁ and HWP₂ before and after the OMC balance the rotation before and after the OMC enabling cancellation of birefringence-induced noise while preserving net 45 degrees rotation.

We first characterised the HWP₀–QWP pair to determine the mount settings that realise a target power ratio between the p - and s -polarised fields (Figure 5.10). This ratio was then fixed for the remainder of the experiment. Using those angle pairs, we estimated the relative optical phase (homodyne angle) δ for each setting (Figure 5.11) and measured the corresponding output powers P_1 and P_2 (Figure 5.12). The homodyne angles were obtained from the measured balanced-homodyne signal and the input powers in each quadrature via Eq. 5.7.

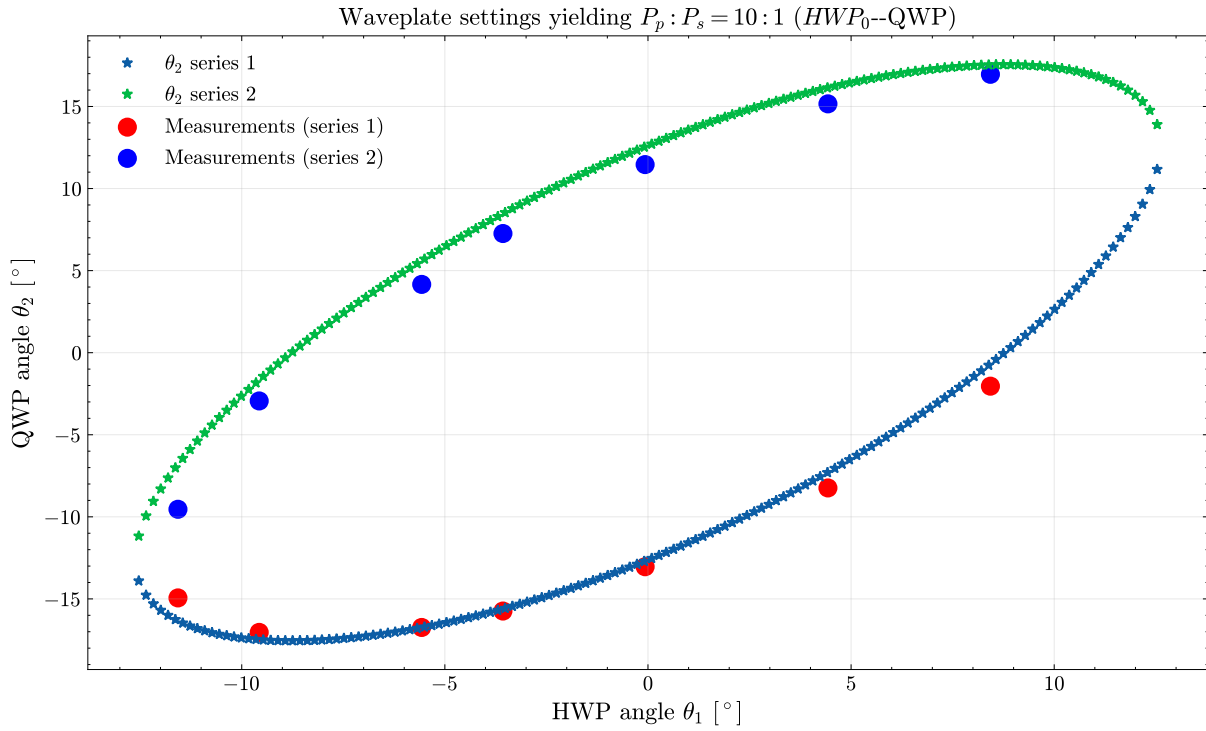


Figure 5.10: Characterisation of the HWP₀-QWP angle pairs that realize the input power ratio $P_p : P_s = 10 : 1$. Plotted are the measured QWP angles θ_2 versus HWP₀ angles θ_1 (two measurement series). Angles are reported relative to calibration offsets $\theta_1 = \text{HWP}_0 - \theta_{1,0}$ and $\theta_2 = \text{QWP} - \theta_{2,0}$.




Figure 5.11: Phase difference δ between the s - and p -polarized outputs evaluated at the waveplate settings from Figure 5.10. The phase is inferred from the balanced-homodyne relation $P = 2\sqrt{P_p P_s} \cos \delta$ with $P \equiv E_y - E_x$ and $(P_p, P_s) \equiv (P_{\text{LO}}, P_{\text{FO}})$. Two measurement series are shown.




Figure 5.12: Output powers P_1 and P_2 versus HWP0 angle θ_1 after the HWP–QWP–HWP stage set to 22.5° . Curves show the modeled intensities $|E_x|^2 (p)$ and $|E_y|^2 (s)$ for two series; markers show the corresponding measured channel fractions (data).

We excited the OMC piezoelectric transducer (PZT) to measure the noise transfer function from OMC length fluctuations to the homodyne signal. The input waveplates were configured to realise the desired homodyne angle. Using the HWP₁ calibration (see Appendix B.1), HWP₁ was set to its mixing angle. HWP₂ was then introduced and adjusted so that the FA did not introduce any further rotation of the polarisations. The homodyne signal was recorded and verified against the model. Finally, HWP₁ and HWP₂ were counter-rotated in opposite directions until the measured noise transfer function was minimised.

We determined the optimum mixing angle for a set of target homodyne angles by jointly rotating HWP₁ and HWP₂ and identifying the HWP₁ setting that minimised the measured transfer function. The results for multiple θ_{HD} values are shown in Figure 5.14. In Figure 5.13 we have plotted this as a function of homodyne angle.




Figure 5.13: Optimal polarisation angle as a function of homodyne angle. The solid curve shows the model prediction for Θ_{opt} from Eq. 5.27, evaluated using the measured birefringence parameters. The markers denote the measured HWP₁ angles that minimise the OMC length-to-homodyne noise transfer at the indicated homodyne angles.

As shown in Figure 5.13, the measured optimum fast-axis angles do not follow the theoretical prediction based on the independently measured birefringence parameters. The origin of this discrepancy is not yet understood. Nevertheless, each experimental point corresponds to the minimum of the measured transfer function from OMC length fluctuations to the homodyne signal, and thus represents the empirically optimal setting for that homodyne angle.

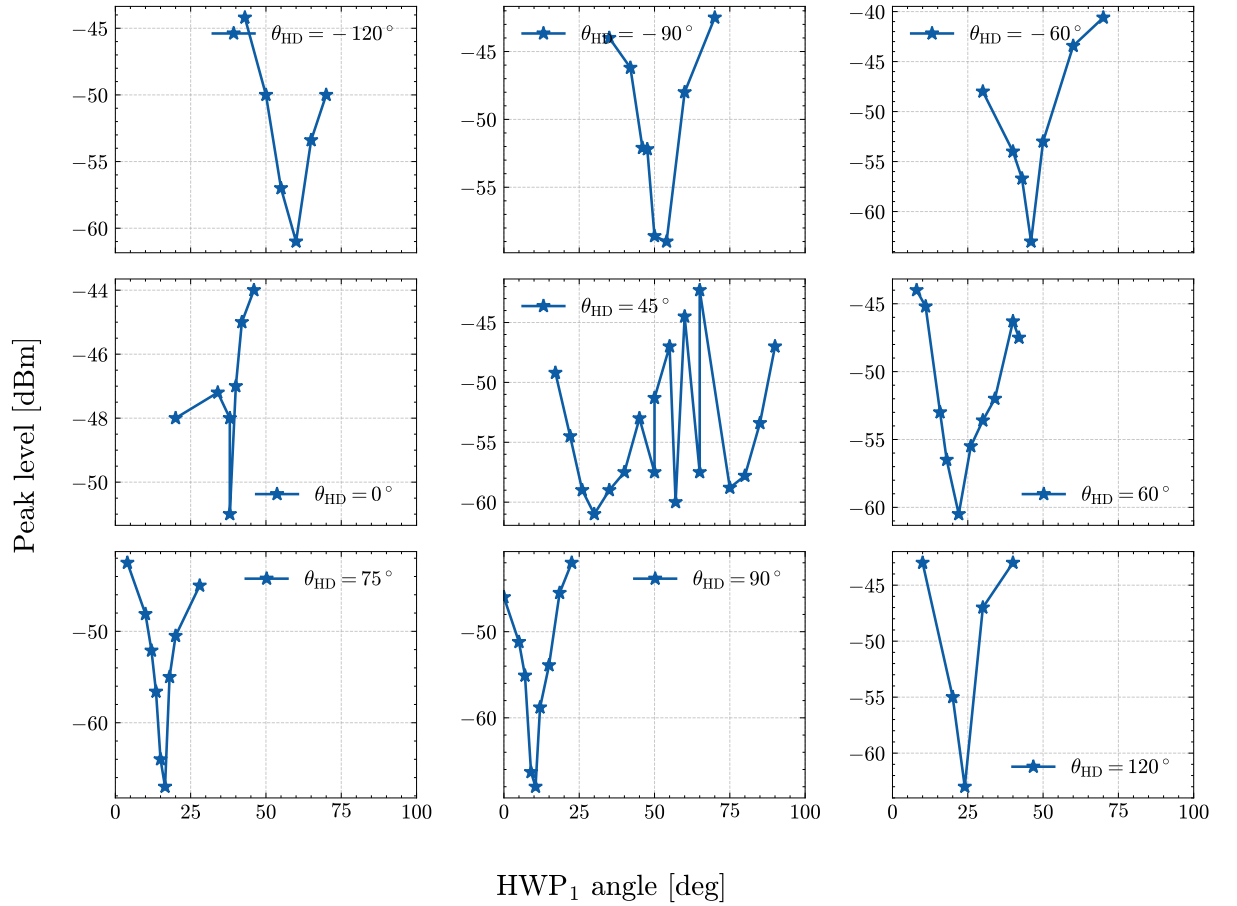


Figure 5.14: Determination of the optimum HWP₁ mixing angle for a range of homodyne angles. Each panel shows the measured peak level of the driven line (in dBm) versus the HWP₁ FA angle, while HWP₂ is counter-rotated so that the homodyne angle remains fixed at the value indicated in the panel (θ_{HD}). The minimum in each curve identifies the HWP₁ setting that minimises the OMC length-to-homodyne signal noise transfer for that θ_{HD} .

5.7 Discussion

This chapter developed and tested a polarisation-based balanced homodyne detection (BHD) scheme that uses a single output mode cleaner (OMC) to filter both the interferometer (IFO) and local-oscillator (LO) fields. By assigning orthogonal polarisations to the two beams and distributing the mixing between half-wave plates (HWPs) before and after the OMC, we derived a compact expression for the birefringence-induced coupling of OMC length noise to the homodyne readout and identified a tunable operating point that suppresses this coupling.

We first measured the OMC birefringence (differences in resonance frequency and finesse between s and p polarisations) and used these parameters to build a predictive model for the noise transfer function. We then implemented a bench-top experiment in which HWP_0 -QWP set the power ratio and homodyne angle, while HWP_1 and HWP_2 were counter-rotated to locate the minimum of the length-to-homodyne transfer function. The procedure yielded optimal mixing angles across homodyne phase, demonstrating empirical suppression of birefringence coupling. Although the trend is broadly consistent with the model, the measured optima deviate quantitatively from the prediction based on independently measured birefringence; the source of this discrepancy remains unresolved.

5.7.1 Future directions

1. **Model refinement:** Incorporate alignment jitter, finite extinction ratios, spatial retardance non-uniformity of waveplates, OMC depolarisation, and coating inhomogeneity into the noise model; propagate uncertainties from birefringence fits through to θ_{opt} .
2. **Improved calibration:** Re-calibrate HWP/QWP fast-axis angles and retardances over the beam footprint; characterise temperature dependence and wavelength sensitivity.
3. **Scaling to observatory conditions:** Repeat the demonstration with a suspended OMC and interferometer-like seismic and thermal environments; evaluate long-term noise, thermal lensing, and alignment cross-couplings.
4. **Comparison with nominal BHD:** Quantitatively compare polarisation BHD against the dual-OMC scheme in terms of quantum noise, classical noise couplings, complexity, and maintainability.

In summary, polarisation BHD provides a practical path to reduce optical complexity while preserving the essential benefits of BHD. With improved modelling and calibration, and with validation on a suspended platform, the technique could become a competitive readout option for future gravitational-wave detectors.

Chapter 6

Conclusion

This thesis has addressed a set of closely connected experimental and modelling problems motivated by the next steps in gravitational-wave detector technology: the move toward cryogenic operation for reduced thermal noise, and the increasing complexity of optical readout chains required for quantum-noise-limited sensitivity. Across these themes, the work presented here contributes (i) modelling and design studies for a cryogenic, multi-material suspension architecture for the Glasgow Cryogenic Interferometer Facility (GCIF), (ii) development and cryogenic investigation of a low-heat-load local displacement sensor for suspension damping, (iii) Output Mode Cleaner development for the Advanced LIGO A+ upgrade, including both high-throughput output mode cleaner (OMC) cavities, (iv) and a simplified balanced-homodyne readout concept based on a single OMC.

6.1 Summary of contributions

6.1.1 Common-platform suspension modelling for GCIF

A central goal of the GCIF is the direct measurement of coating thermal noise using a short, cryogenic Fabry–Pérot cavity formed by suspended silicon test masses. Achieving the required displacement sensitivity demands stringent control of seismic coupling while preserving practical damping and lock-acquisition capability. In Chapter 2, a Mathematica state-space model was developed to study a *common-platform* (split-pendulum) suspension architecture, in which two suspension chains share a rigid upper mass. The principal motivation of this design is the potential for common-mode rejection: disturbances transmitted through the shared top stage can be mechanically correlated between the input and end chains, thereby reducing the differential motion that appears as cavity length noise.

The modelling results show that the split geometry reorganises the resonance structure into clusters with systematic mixed-DOF content, in contrast to the more separable mode families typical of conventional single-chain triple suspensions. This has two practical consequences. First, sensing and damping must be designed with mixed families and occasional cross-coupling pathways in mind. Second, the parameter space of the common platform (mass properties, wire attachment geometry, and chain-to-chain symmetry) can be used deliberately to manage degeneracies, improve observability, and mitigate troublesome couplings.

Several targeted studies were carried out to inform design choices: (i) shifting the top-mass centre of mass toward the geometric reference plane improves overall mode visibility and reduces longitudinal-to-pitch coupling, motivating compensating-mass strategies in the presence of recoil structures; (ii) increasing the final-stage length generally improves average mode visibility while leaving cross-coupling transfer functions broadly unchanged

within the explored set, supporting the preference for longer monolithic lower stages subject to mechanical constraints; and (iii) increasing the common-platform mass reduces certain cross-couplings by raising the platform inertia, but simultaneously reduces the platform visibility of several resonance families, highlighting a trade-off between cross-coupling robustness and local-control observability.

Finally, realistic chain-to-chain asymmetries were introduced to reflect tolerances and uncertainties in monolithic attachment geometry at the optic stage. The modelling indicates that modest asymmetry does not strongly alter the overall cross-coupling transfer functions, while beneficially lifting degeneracies (mode splitting). These results support the practical feasibility of the common-platform approach for GCIF and provide a modelling basis for the next iteration of mechanical design and controls planning.

6.1.2 Cryogenic shadow sensor for local damping

Reliable local sensing and damping remain essential for commissioning and stable operation, particularly in cryogenic systems where actuator authority and alignment margins can be constrained. Chapter 3 presented the development of a cryogenic-compatible shadow sensor intended for local control of the GCIF suspensions. A systematic characterisation campaign was performed on infrared LEDs, quadrant photodiodes, and flag geometries down to cryogenic temperatures, supported by a detailed performance and noise model.

A key outcome is that the prototype sensor meets the combined requirements of sensitivity, dynamic range, and heat load for cryogenic suspension use. In particular, the normalised QPD readout scheme was demonstrated to suppress LED intensity noise to near shot-noise levels, and cryogenic operation was validated in testing at 5 K with stable performance consistent with the model predictions. The modelling and measurements jointly clarify

the design trade space between responsivity (slope at the operating point), dynamic range (linear range about that point), and heat load (set largely by LED drive power and heat extraction pathways). This provides a rational route to selecting geometry and operating point for the GCIF implementation rather than relying on purely empirical optimisation.

The chapter also identified the most important practical limitations to address in the next hardware iteration: LED package robustness under thermal cycling, improved optical collection/collimation strategies compatible with vacuum and cryogenic conditions, and closed-loop demonstrations on suspended masses to directly validate displacement-noise performance under feedback.

6.1.3 A+ LIGO OMC cavities: design, assembly, and performance

OMCs are critical components of the interferometer readout chain: they reject higher-order spatial modes and RF sidebands while transmitting the carrier field with minimal loss. In Chapter 4, the design, fabrication, and characterisation of the A+ LIGO monolithic bow-tie OMC cavities were presented, including the screening and metrology of coated optics, the bonding and assembly workflow, and the measurement methods used to optimise throughput and confirm cavity geometry.

A major practical outcome of this work is the demonstrated optical performance of the A+ OMC builds. For the first bonded cavity, an optical throughput of 99.10% was achieved, with measured cavity parameters (FSR, round-trip length, transverse mode spacing) meeting the design requirements. The refined procedures and scripts developed during the build of the first cavity were then used to support the production of the subsequent cav-

ities, enabling consistent measurement and acceptance of key metrics such as finesse, unit throughput, and loss per mirror. Taken together, these results represent a meaningful improvement in OMC performance and reproducibility relative to earlier builds, and directly support squeezed-light operation in the upgraded detectors.

6.1.4 Polarisation-based balanced homodyne detection with a single OMC

Balanced homodyne detection (BHD) is a cornerstone of the A+ quantum-noise upgrade but introduces significant optical complexity in its nominal implementation, including the need for two OMC cavities and the associated risk of differential imperfections degrading the “balanced” subtraction. In Chapter 5, we developed and tested a polarisation-based BHD scheme in which the interferometer signal and local oscillator are carried on orthogonal polarisations and filtered by a *single* OMC. The work established the theoretical framework for this configuration, explicitly treating birefringence-induced differences between the s- and p-polarised cavity responses, and derived a compact expression describing the coupling of OMC length noise into the homodyne readout.

The central conceptual result is that the coupling can be tuned and, in principle, suppressed by distributing the required polarisation mixing before and after the OMC using appropriately chosen waveplate angles. A bench-top experiment demonstrated this principle operationally: for a range of homodyne angles, an empirically optimal mixing setting was found that minimised the measured length-to-homodyne transfer function, thereby confirming practical suppression of birefringence coupling. While the observed trend was broadly consistent with the model, the measured optimal angles deviated quantitatively from predictions based on independently measured birefringence parameters, indicating that additional effects (e.g. waveplate non-idealities, alignment jitter, finite extinction ratios, or spatial retardance variations) must be incorporated before a full predictive model can be claimed.

Overall, the results position polarisation BHD as a credible path to reduce optical complexity while preserving the essential benefits of homodyne readout, provided that modelling, calibration, and suspended-platform validation are pursued.

6.2 Outlook and recommended future work

The studies in this thesis naturally motivate several next steps, grouped by subsystem.

6.2.1 GCIF suspension and controls integration

The common-platform modelling should now be coupled more tightly to the mechanical design iteration: candidate top-mass geometries (including recoil structures, compensating masses, and realistic attachment tolerances) can be evaluated against the combined criteria of (i) low longitudinal-to-angular cross-coupling, (ii) sufficient mode visibility at available sensor/actuator stages, and (iii) robust common-mode rejection in the projected length-noise budget. As hardware prototypes mature, model-to-measurement validation should become a priority: transfer functions and resonance frequencies measured on prototype assemblies will provide high-leverage constraints on stiffness and damping assumptions in the model.

6.2.2 Cryogenic shadow sensor further work and miniaturisation

For the shadow sensor, development of a mechanical package that withstands thermal cycles and a closed-loop demonstration on a suspended mass to quantify in-situ displacement noise, stability margins, and the interplay between sensor noise and actuator authority. Finally, the modelling indicates a clear path to increased sensitivity (higher optical power and responsivity), but this must be pursued alongside a rigorous heat-load and heat-extraction analysis consistent with GCIF thermal constraints.

6.2.3 A+ OMC deployment and extension

For the A+ OMCs, continued emphasis on reproducible metrology, bonding workflows, and loss budgeting will remain essential as cavities are installed, transported, and operated in observatory conditions. The methods developed here can also inform future OMC designs for third-generation detectors and also existing detectors looking to improve their readout.

6.2.4 Polarisation BHD refinement and scaling

The polarisation BHD concept warrants further work in four directions: (i) model refinement to include realistic non-idealities and to propagate parameter uncertainties through to predicted optimal settings; (ii) improved calibration of waveplate retardance and fast-axis angles over the relevant beam footprint and over environmental drifts; (iii) demonstration in a more interferometer-like environment, ideally including a suspended OMC and long-duration stability measurements; and (iv) quantitative comparison against the nominal dual-OMC BHD chain, including both classical noise couplings and practical maintainability. The aim is also to understand the existing discrepancy between the experimental results and the model prediction.

6.3 Closing remarks

The experimental direction of gravitational-wave detection continues to be shaped by the need to suppress thermal noise and to realise low-loss, low-noise optical readout at ever higher levels of quantum enhancement. The work presented in this thesis advances both aims: it helps establish a credible cryogenic suspension architecture and local-sensing solution for GCIF, and it contributes directly to A+ readout hardware through improved OMC cavity performance and a demonstrated route to simplifying homodyne readout using polarisation multiplexing. Together, these results strengthen the technical foundations for near-term upgrades and provide practical methods and insights that are relevant to the design choices facing third-generation detectors.

Appendices

A PCB Images

A.1 LED Driver




Figure 6: Photograph of the LED driver printed circuit board. The circuit diagram for this is found in Figure 3.1

A.2 QPD Amplifier




Figure 7: Photograph of the QPD amplifier printed circuit board.

B Wave Plate Calibration Curves

Below we plot the calibration curves for the waveplates. They are calibrated independently.




Figure 8: Calibration of HWP0, the power ratio half wave plate. The normalised photodiode signals in the s - and p -polarisation channels are shown as a function of the HWP0 rotation angle. The data are fitted with the expected sinusoidal responses, $f_s(\theta) = A_s \sin^2(2(\theta - \theta_{0,s})) + \text{off}_s$ and $f_p(\theta) = A_p \cos^2(2(\theta - \theta_{0,p})) + \text{off}_p$, confirming the orthogonal mixing of the two polarisations. Vertical dashed and dotted lines indicate the fitted extrema (peaks/dips) within the scan range, which define the effective fast-axis reference angles used for subsequent polarisation-BHD alignment and balancing.




Figure 9: Calibration of the quarter-wave plate (QWP), the power ratio quarter wave plate. The photodiode signal after the PBS is shown as a function of QWP rotation angle. The data are fitted with the expected periodic dependence of the transmitted power for an initially linearly polarised field passing through a QWP and analysed by a PBS, $f(\theta) = 2A \sin^2(\theta - \theta_0) \cos^2(\theta - \theta_0) + \text{off}$. The vertical dashed and dotted lines indicate the fitted extrema (dip and peak) within the scanned range, which define the effective QWP fast-axis reference angle used for setting the desired polarisation state in the polarisation-BHD experiment.




Figure 10: Calibration of the first mixing half-wave plate (HWP1)




Figure 11: Calibration of the second mixing half-wave plate (HWP2)

B.1 Calibration of Waveplates In-Situ




Figure 12: Calibration of the first mixing half-wave plate (HWP1) in the presence of all the other waveplates upstream i.e. HWP0 and QWP, and other optical components (see Figure 5.9)

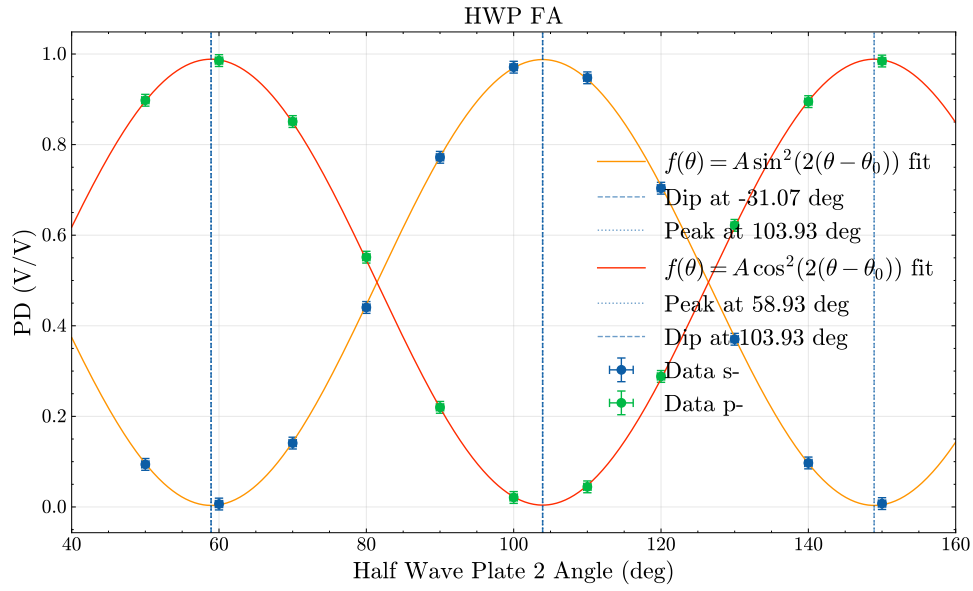


Figure 13: Calibration of the first mixing half-wave plate (HWP1) in the presence of all the other waveplates upstream i.e. HWP0, QWP and HWP1, and other optical components (see Figure 5.9)

C Suspension Model Parameters for State Space Export

Table 6: Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: masses and inertias. Here, mass 1 denotes the upper/common platform, mass 2 the intermediate masses, and mass 3 the optics. Subscripts *I* and *E* refer to the input and end chains respectively.

Parameter	Description
Masses	
m_1	Mass of the upper/common platform (mass 1).
m_{2I}, m_{2E}	Masses of the intermediate stages (mass 2) in the input and end chains.
m_{3I}, m_{3E}	Masses of the optics (mass 3) in the input and end chains.
Moments/products of inertia	
I_{1x}, I_{1y}, I_{1z}	Principal moments of inertia of the upper/common platform (mass 1) about its body x, y, z axes.
$I_{1xy}, I_{1yz}, I_{1zx}$	Products of inertia of the upper/common platform (mass 1).
$I_{2Ix}, I_{2Iy}, I_{2Iz}$	Principal moments of inertia of the input-chain intermediate mass (mass 2I).
$I_{2Ixy}, I_{2Iyz}, I_{2Izx}$	Products of inertia of the input-chain intermediate mass (mass 2I).
$I_{2Ex}, I_{2Ey}, I_{2Ez}$	Principal moments of inertia of the end-chain intermediate mass (mass 2E).
$I_{2Exy}, I_{2Eyz}, I_{2Ezx}$	Products of inertia of the end-chain intermediate mass (mass 2E).
$I_{3Ix}, I_{3Iy}, I_{3Iz}$	Principal moments of inertia of the input-chain optic (mass 3I).
$I_{3Ixy}, I_{3Iyz}, I_{3Izx}$	Products of inertia of the input-chain optic (mass 3I).
$I_{3Ex}, I_{3Ey}, I_{3Ez}$	Principal moments of inertia of the end-chain optic (mass 3E).
$I_{3Exy}, I_{3Eyz}, I_{3Ezx}$	Products of inertia of the end-chain optic (mass 3E).

Table 7: Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: attachment-point heights and wire separations. Subscripts I and E refer to the input and end chains, respectively.

Parameter	Description
Attachment-point heights (along z) and blade-to-mass offsets	
d_{top}	Distance down from the upper blades (on the structure) to the upper-wire attachment points.
d_0	Distance up from the upper-mass FRP to the upper-wire attachment points.
d_{1I}, d_{1E}	Distance down from intermediate blades (on the upper mass) to attachment points for intermediate wires (input/end chain).
d_{2I}, d_{2E}	Distance up from the intermediate-mass FRP to attachment points for intermediate wires (input/end chain).
d_{3I}, d_{3E}	Distance down from blades (on the intermediate mass) to attachment points for lower wires (input/end chain).
d_{4I}, d_{4E}	Distance up from the lower-mass FRP to attachment points for lower wires/ribbons/rods (input/end chain).
Wire separations	
n_0	y -direction half-separation at the <i>top</i> end of the upper wires.
n_1	y -direction half-separation at the <i>bottom</i> end of the upper wires.
n_{2I}, n_{2E}	y -direction half-separation at the <i>top</i> end of the intermediate wires (input/end chain).
n_{3I}, n_{3E}	y -direction half-separation at the <i>bottom</i> end of the intermediate wires (input/end chain).
n_{4I}, n_{4E}	y -direction half-separation at the <i>top</i> end of the lower wires (input/end chain).
n_{5I}, n_{5E}	y -direction half-separation at the <i>bottom</i> end of the lower wires (input/end chain).
s_u	One-sided x -direction separation of the upper wires.
s_{iI}, s_{iE}	x -direction half-separation of intermediate wires (input/end chain).
s_{lI}, s_{lE}	x -direction half-separation of lower wires (input/end chain).

Table 8: Key model parameters used to construct the split-pendulum (common-platform) Mathematica suspension model: stage lengths and elastic properties. Subscripts I and E refer to the input and end chains, respectively.

Parameter	Description
Stage lengths	
l_1	Upper wire length: structure blades \rightarrow upper/common platform.
l_{2I}, l_{2E}	Intermediate wire lengths: upper/common platform \rightarrow intermediate mass (input/end chain).
l_{3I}, l_{3E}	Lower wire/ribbon lengths: intermediate mass \rightarrow optic (input/end chain).
Elastic constants (vertical blades and longitudinal wire stretch)	
k_{buz}	Net vertical elasticity of the upper blades (on structure).
k_{bizI}, k_{bizE}	Net vertical elasticity of the intermediate blades (on the upper mass), input/end chain.
k_{w1}	Net longitudinal elasticity of upper wires (structure \rightarrow upper mass).
k_{w2I}, k_{w2E}	Net longitudinal elasticity of intermediate wires (upper mass \rightarrow intermediate mass), input/end chain.
k_{w3I}, k_{w3E}	Net longitudinal elasticity of lower wires/ribbons (intermediate mass \rightarrow optic), input/end chain.
Material stiffness and bending	
Y_1	Young's modulus of the upper wires.
Y_{2I}, Y_{2E}	Young's modulus of intermediate wires (input/end chain).
Y_{3I}, Y_{3E}	Young's modulus of lower wires/ribbons (input/end chain).
M_{11}	Second moment of area of the upper wires in direction "1".
M_{21I}, M_{21E}	Second moment of area of the intermediate wires in direction "1" (input/end chain).
M_{31I}, M_{31E}	Second moment of area of the lower wires/ribbons/rods in direction "1" (input/end chain).

D Suspension Modes for Common Platform Suspensions

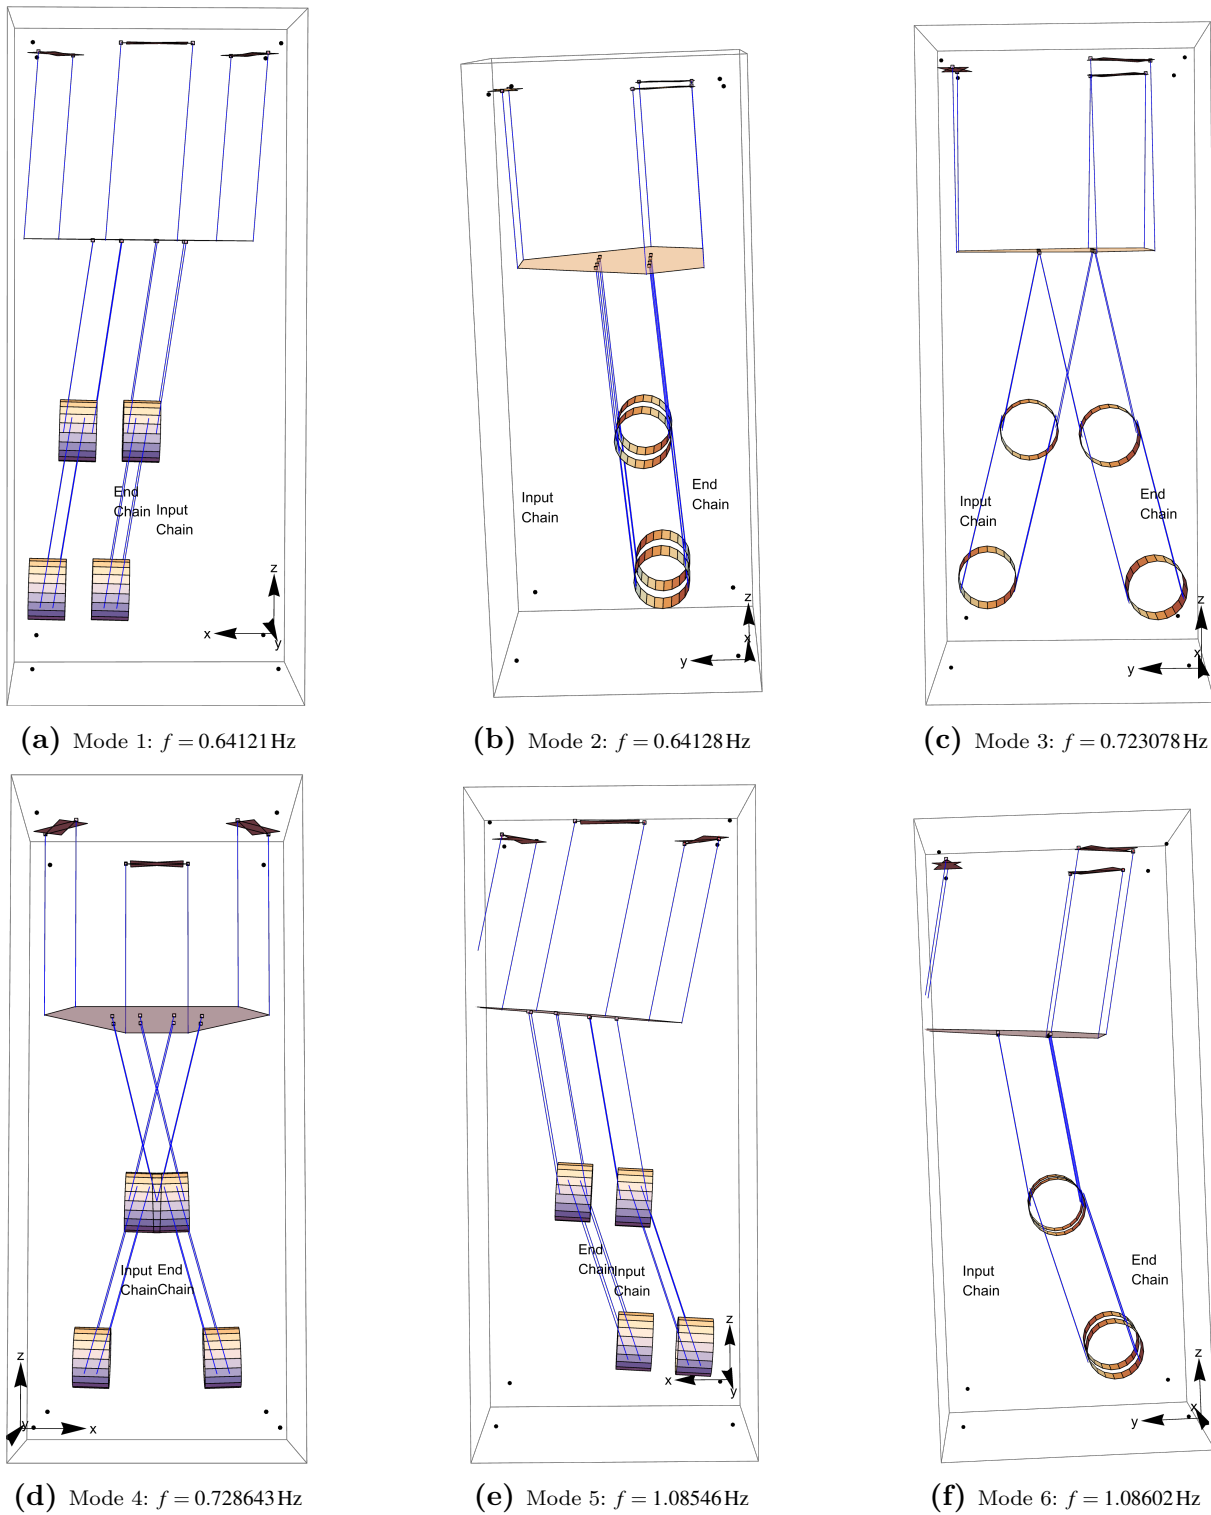


Figure 14: Common platform suspension modes (1–6).

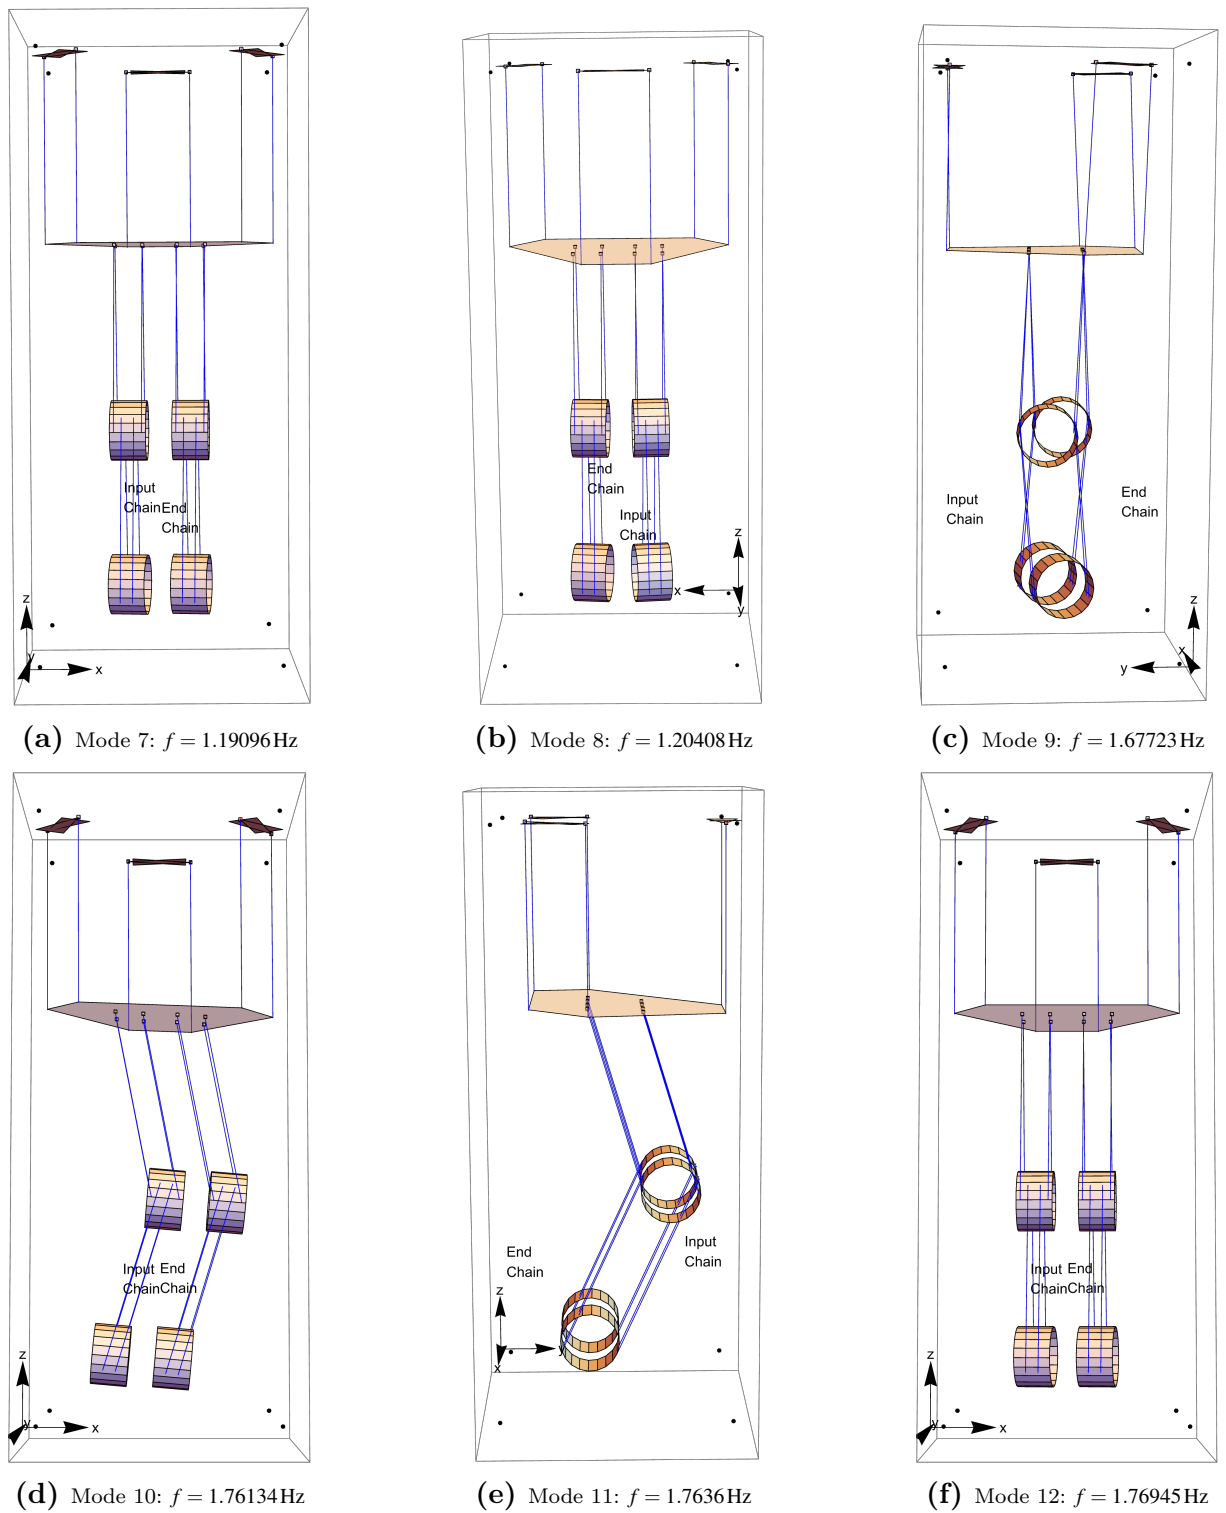


Figure 15: Common platform suspension modes (7–12).

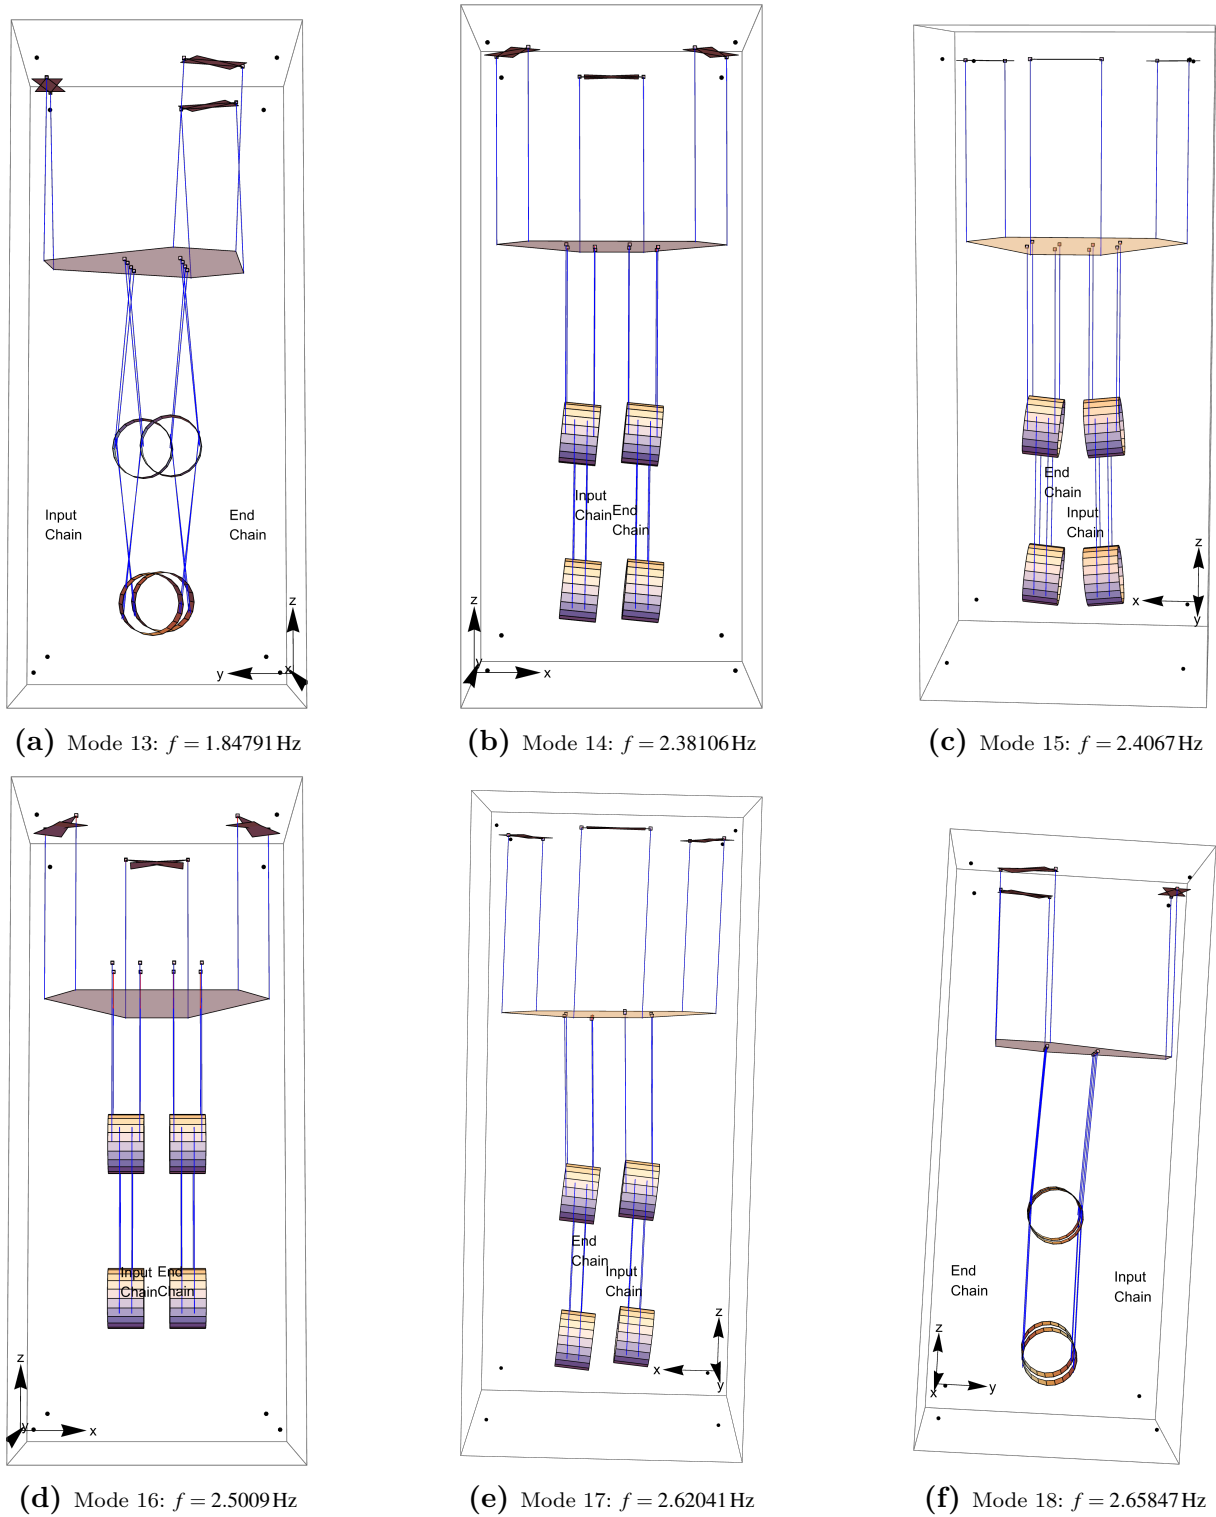


Figure 16: Common platform suspension modes (13–18).

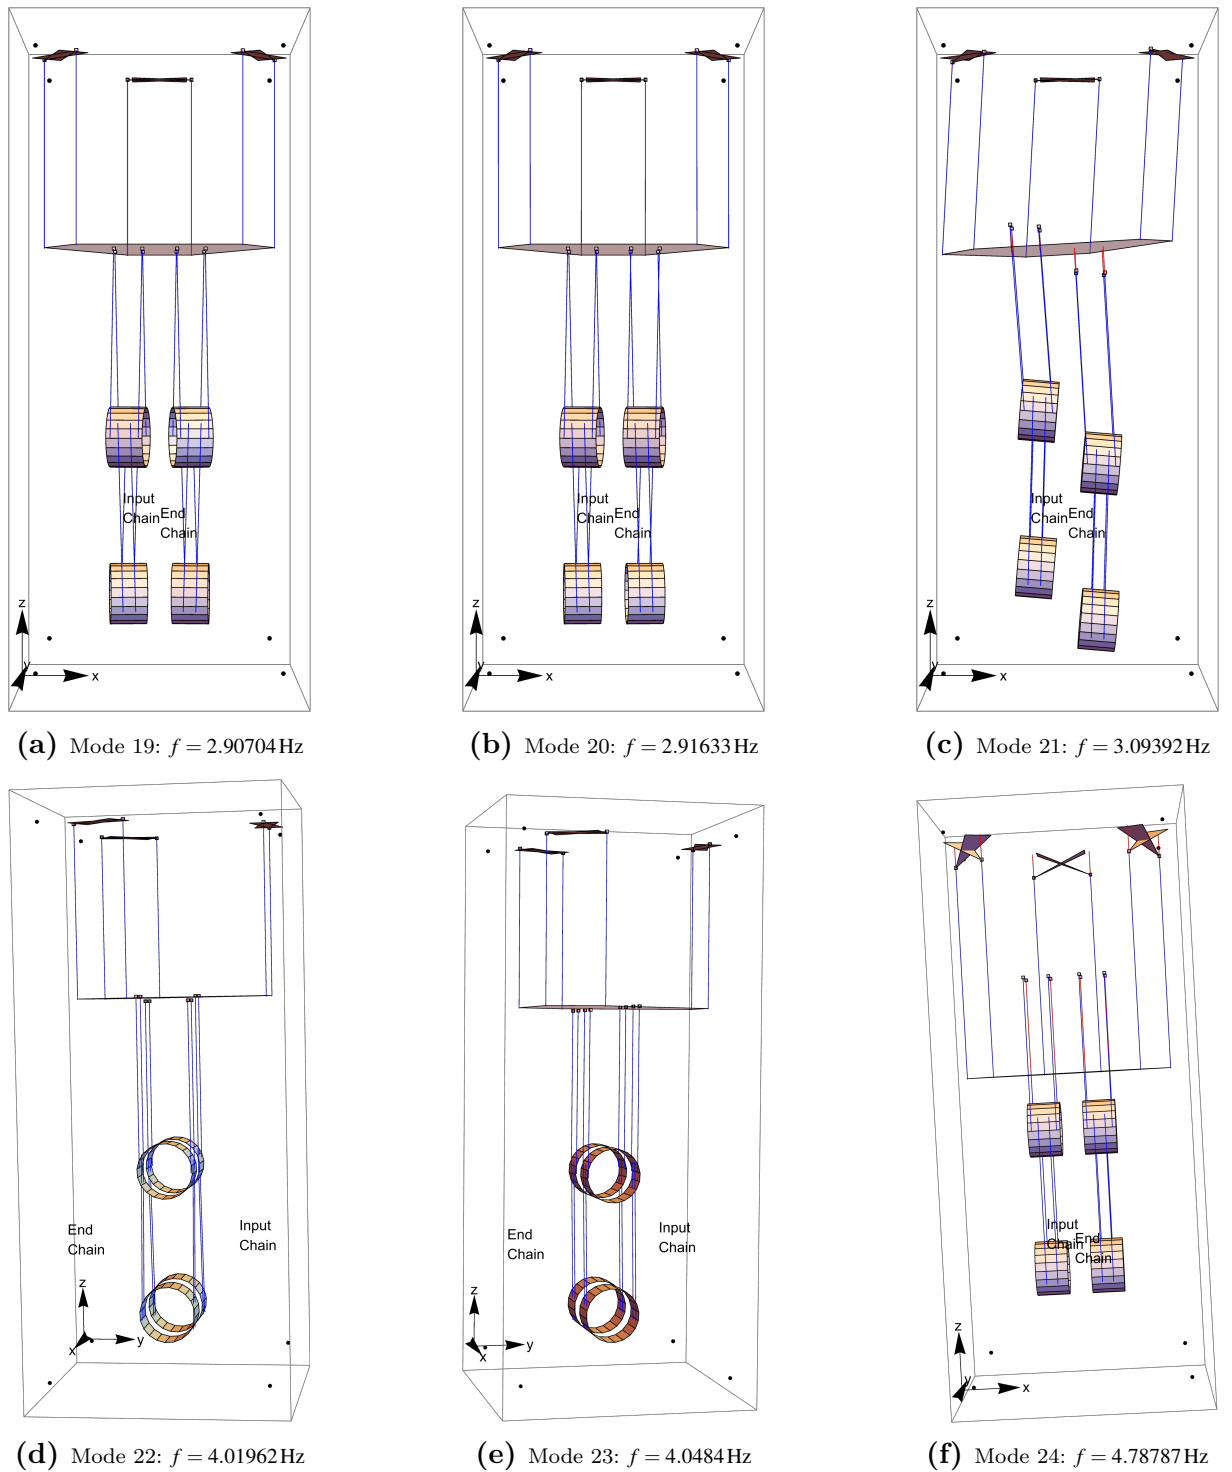





Figure 17: Common platform suspension modes (19–24).




(a) Mode 25: $f = 66.1018\text{Hz}$




(b) Mode 26: $f = 66.1018\text{Hz}$




(c) Mode 27: $f = 194.578\text{Hz}$



(d) Mode 28: $f = 194.578\text{Hz}$



(e) Mode 29: $f = 258.154\text{Hz}$



(f) Mode 30: $f = 258.154\text{Hz}$

Figure 18: Common platform suspension modes (25–30).

Bibliography

- [1] J. Aasi, B. P. Abbott, R. Abbott et al. ‘Advanced LIGO’. In: *Classical and Quantum Gravity* 32.7 [2015], p. 074001. doi: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001).
- [2] A. G. Abac et al. ‘GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run’. In: [2025]. arXiv:2508.18082. arXiv: [2508.18082](https://arxiv.org/abs/2508.18082) [gr-qc].
- [3] A. G. Abac et al. ‘Search for Continuous Gravitational Waves from Known Pulsars in the First Part of the Fourth LIGO–Virgo–KAGRA Observing Run’. In: *The Astrophysical Journal* 983 [2025], p. 99. doi: [10.3847/1538-4357/adb3a0](https://doi.org/10.3847/1538-4357/adb3a0). arXiv: [2501.01495](https://arxiv.org/abs/2501.01495) [gr-qc]. url: <https://doi.org/10.3847/1538-4357/adb3a0>.
- [4] B. P. Abbott et al. ‘GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence’. In: *Physical Review Letters* 116 [2016], p. 241103. doi: [10.1103/PhysRevLett.116.241103](https://doi.org/10.1103/PhysRevLett.116.241103). arXiv: [1606.04855](https://arxiv.org/abs/1606.04855) [gr-qc]. url: <https://link.aps.org/doi/10.1103/PhysRevLett.116.241103>.
- [5] B. P. Abbott et al. ‘GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral’. In: *Physical Review Letters* 119 [2017], p. 161101. doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [gr-qc].
- [6] B. P. Abbott et al. ‘GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs’. In: *Physical Review X* 9 [2019], p. 031040. doi: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040).

- [7] B. P. Abbott et al. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Physical Review Letters* 116.6 [2016], p. 061102. doi: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). url: <https://doi.org/10.1103/PhysRevLett.116.061102>.
- [8] B. P. Abbott et al. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Phys Rev Lett* 116.6 [Feb. 2016]. Edition: 2016/02/27, p. 061102. issn: 1079-7114 (Electronic) 0031-9007 (Linking). doi: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). url: <https://www.ncbi.nlm.nih.gov/pubmed/26918975>.
- [9] B. P. et al Abbott. ‘GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral’. In: *Phys. Rev. Lett.* 119 [16 2017], p. 161101. doi: [10.1103/PhysRevLett.119.161101](https://link.aps.org/doi/10.1103/PhysRevLett.119.161101). url: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
- [10] B. P. et al Abbott. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Phys. Rev. Lett.* 116 [6 2016], p. 061102. doi: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). url: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [11] R. Abbott et al. ‘All-sky search for short gravitational-wave bursts in the third Advanced LIGO and Advanced Virgo run’. In: *Physical Review D* 104.12 [2021], p. 122004. doi: [10.1103/PhysRevD.104.122004](https://doi.org/10.1103/PhysRevD.104.122004). arXiv: [2107.03701 \[gr-qc\]](https://arxiv.org/abs/2107.03701).
- [12] R. Abbott et al. ‘GW190814: Gravitational Waves from the Coalescence of a $23 M_{\odot}$ Black Hole with a $2.6 M_{\odot}$ Compact Object’. In: *The Astrophysical Journal Letters* 896.2 [2020], p. L44. doi: [10.3847/2041-8213/ab960f](https://doi.org/10.3847/2041-8213/ab960f). arXiv: [2006.12611 \[astro-ph.HE\]](https://arxiv.org/abs/2006.12611).
- [13] R. Abbott et al. ‘GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run’. In: *Physical Review X* 11 [2021], p. 021053. doi: [10.1103/PhysRevX.11.021053](https://doi.org/10.1103/PhysRevX.11.021053).
- [14] R. Abbott et al. ‘GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run’. In: *Physical Review X* 13 [2023], p. 041039. doi: [10.1103/PhysRevX.13.041039](https://doi.org/10.1103/PhysRevX.13.041039). arXiv: [2111.03606 \[gr-qc\]](https://arxiv.org/abs/2111.03606). url: <https://link.aps.org/doi/10.1103/PhysRevX.13.041039>.

- [15] R. Abbott et al. ‘GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run’. In: *Physical Review X* 13 [2023], p. 041039. doi: [10.1103/PhysRevX.13.041039](https://doi.org/10.1103/PhysRevX.13.041039).
- [16] R. Abbott et al. ‘Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences’. In: *The Astrophysical Journal Letters* 915.1 [2021], p. L5. doi: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e). arXiv: [2106.15163](https://arxiv.org/abs/2106.15163) [astro-ph.HE]. url: <https://arxiv.org/abs/2106.15163>.
- [17] R. Abbott et al. ‘Search for anisotropic gravitational-wave backgrounds using data from Advanced LIGO and Advanced Virgo’s first three observing runs’. In: *Physical Review D* 104.2 [2021], p. 022005. doi: [10.1103/PhysRevD.104.022005](https://doi.org/10.1103/PhysRevD.104.022005). arXiv: [2103.08520](https://arxiv.org/abs/2103.08520) [gr-qc]. url: <https://arxiv.org/abs/2103.08520>.
- [18] R. Abbott et al. ‘Search for Gravitational-wave Transients Associated with Magnetar Bursts in Advanced LIGO and Advanced Virgo Data from the Third Observing Run’. In: *The Astrophysical Journal* 966 [2024], p. 137. doi: [10.3847/1538-4357/ad27d3](https://doi.org/10.3847/1538-4357/ad27d3). arXiv: [2210.10931](https://arxiv.org/abs/2210.10931) [astro-ph.HE].
- [19] R. Abbott et al. ‘Upper Limits on the Isotropic Gravitational-Wave Background from Advanced LIGO’s and Advanced Virgo’s Third Observing Run’. In: *Physical Review D* 104.2 [2021], p. 022004. doi: [10.1103/PhysRevD.104.022004](https://doi.org/10.1103/PhysRevD.104.022004). arXiv: [2101.12130](https://arxiv.org/abs/2101.12130) [gr-qc]. url: <https://arxiv.org/abs/2101.12130>.
- [20] A. et al Abramovici. ‘LIGO: The Laser Interferometer Gravitational-Wave Observatory’. In: *Science* 256.5055 [1992], pp. 325–333. doi: [10.1126/science.256.5055.325](https://doi.org/10.1126/science.256.5055.325).
- [21] Sadao Adachi. *Properties of Semiconductor Alloys: Group-IV, III-V and II-VI Semiconductors*. John Wiley & Sons, 2009.
- [22] R. X. Adhikari et al. ‘A cryogenic silicon interferometer for gravitational-wave detection’. In: *Classical and Quantum Gravity* 37.16 [2020], p. 165003. doi: [10.1088/1361-6382/ab9143](https://doi.org/10.1088/1361-6382/ab9143).
- [23] Gabriella Agazie et al. ‘The NANOGrav 15-year Data Set: Evidence for a Gravitational-Wave Background’. In: *The Astrophysical Journal Letters* [2023]. doi: [10.3847/2041-8213/acdac6](https://doi.org/10.3847/2041-8213/acdac6). arXiv: [2306.16213](https://arxiv.org/abs/2306.16213) [astro-ph.HE]. url: <https://arxiv.org/abs/2306.16213>.

- [24] T. Akutsu, M. Ando, K. Arai et al. ‘Overview of KAGRA: Detector design and construction history’. In: *Progress of Theoretical and Experimental Physics* 2021.5 [2021], 05A101. doi: [10.1093/ptep/ptaa125](https://doi.org/10.1093/ptep/ptaa125). arXiv: [2005.05574](https://arxiv.org/abs/2005.05574).
- [25] Pau Amaro-Seoane, Heather Audley, Stanislav Babak et al. ‘Laser Interferometer Space Antenna’. In: *arXiv e-prints* [2017], arXiv:1702.00786. arXiv: [1702.00786](https://arxiv.org/abs/1702.00786) [[astro-ph](https://arxiv.org/archive/astro-ph).IM].
- [26] Y. Aoyama et al. ‘Cryogenic suspension systems for the KAGRA gravitational-wave detector’. In: *Classical and Quantum Gravity* 35.2 [2018], p. 025006. doi: [10.1088/1361-6382/aa960f](https://doi.org/10.1088/1361-6382/aa960f).
- [27] Stephen Appert. *A+ OMC Alignment and Bonding Assembly*. <https://dcc.ligo.org/LIGO-D2300300>. LIGO Document D2300300-v1. 2023.
- [28] Stephen Appert. *A+ OMC Alignment and Bonding Assembly, Piezo Actuator Bracket*. <https://dcc.ligo.org/LIGO-D2300299>. LIGO Document D2300299-v1. 2023.
- [29] Stephen Appert and Liyuan Zhang. *A+ OMC, Component Testing, Total Integrated Scatter (TIS) and Transmittance Measurements*. Tech. rep. LIGO-E2300029-v1. Measurements completed on the Caltech RTS setup, 2023 Jan – Feb. LIGO Laboratory, 2023. url: <https://dcc.ligo.org/LIGO-E2300029/public>.
- [30] Koji Arai. *Output Mode Cleaner Design*. 2013.
- [31] Koji Arai. *Some thoughts on polarization BHD*. Oct. 2018.
- [32] Koji Arai, William Korth and Jeffrey Lewis. *aLIGO Output Mode Cleaner: Optical Testing and Results*. Technical Report LIGO-T1500060-v3. Updated August 2023. Available at <https://dcc.ligo.org/LIGO-T1500060-v3>. LIGO Project, 2015.
- [33] Koji Arai, Jeffrey Lewis and Peter Fritschel. *Advanced LIGO Output Mode Cleaner Glass Breadboard Assembly and Installation Procedure*. Engineering Document E1300201-v1. <https://dcc.ligo.org/LIGO-E1300201>. LIGO Document Control Center, 2013.
- [34] Koji Arai et al. *Output Mode Cleaner (OMC) Design*. Technical Report LIGO-T1000276-v5. Updated January 2013. Available at <https://dcc.ligo.org/LIGO-T1000276-v5>. LIGO Project, 2010.

- [35] M. Armano and LISA Pathfinder Collaboration. ‘Capacitive sensing of test mass motion with nanometer precision over millimeter-wide sensing gaps for space-borne gravitational reference sensors’. en. In: *Physical Review D* 96.6 [Sept. 2017], p. 062004. issn: 2470-0010, 2470-0029. doi: [10.1103/PhysRevD.96.062004](https://doi.org/10.1103/PhysRevD.96.062004). url: <https://link.aps.org/doi/10.1103/PhysRevD.96.062004> [visited on 16/06/2022].
- [36] M. Armano et al. ‘Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μHz ’. In: *Physical Review Letters* 120 [2018], p. 061101. doi: [10.1103/PhysRevLett.120.061101](https://doi.org/10.1103/PhysRevLett.120.061101).
- [37] M. Armano et al. ‘Sub-Femto-g Free Fall for Space-Based Gravitational Wave Observatories: LISA Pathfinder Results’. In: *Physical Review Letters* 116 [2016], p. 231101. doi: [10.1103/PhysRevLett.116.231101](https://doi.org/10.1103/PhysRevLett.116.231101).
- [38] Y. Aso et al. ‘Interferometer design of the KAGRA gravitational wave detector’. In: *Phys. Rev. D* 88.043007 [2013]. Describes TM/IM/MN with corresponding suspended recoil masses (RM chain). doi: [10.1103/PhysRevD.88.043007](https://doi.org/10.1103/PhysRevD.88.043007). url: <https://link.aps.org/doi/10.1103/PhysRevD.88.043007>.
- [39] S. M. Aston et al. ‘Update on quadruple suspension design for Advanced LIGO’. In: *Classical and Quantum Gravity* 29.23 [Oct. 2012]. Publisher: IOP Publishing, p. 235004. doi: [10.1088/0264-9381/29/23/235004](https://doi.org/10.1088/0264-9381/29/23/235004). url: <https://doi.org/10.1088/0264-9381/29/23/235004>.
- [40] Stuart Aston. ‘BOSEM Design Document & Test Report’. In: *LIGO DCC* [2014].
- [41] Stuart Mark Aston. ‘Optical Read-out Techniques for the Control of Test-masses in Gravitational Wave Observatories’. Volume: PhD. PhD thesis. University of Birmingham, 2011.
- [42] Lisa Barsotti. *LIGO Evolutions*. Tech. rep. LIGO-G2300963. Talk at ET Symposium (Cagliari), 11 May 2023. LIGO Document Control Center, 2023. url: https://indico.ego-gw.it/event/562/contributions/4895/attachments/2822/4954/LIGO_ETSymposium_G2300963.pdf.
- [43] Lisa Barsotti, Jan Harms and Roman Schnabel. ‘Squeezed vacuum states of light for gravitational wave detectors’. In: *Reports on Progress in Physics* 82.1 [2019], p. 016905. doi: [10.1088/1361-6633/aab906](https://doi.org/10.1088/1361-6633/aab906).

- [44] Lisa Barsotti et al. *The A+ design curve*. Tech. rep. LIGO-T1800042-v4. LIGO Document Control Center, 2018. url: <https://dcc.ligo.org/public/0149/T1800042/004/T1800042-v4.pdf>.
- [45] Mark Barton. *Models of the Advanced LIGO Suspensions in Mathematica™*. LIGO Technical Report T020205-02D. Available from the LIGO Document Control Center. LIGO Laboratory, 2006.
- [46] Mark Barton et al. *A+ OMC Assembly Procedure*. Engineering Document LIGO-E2100109-v2. Contents revised 05 Aug 2021; metadata revised 02 Sep 2022. LIGO Scientific Collaboration, Mar. 2021. url: <https://dcc.ligo.org/LIGO-E2100109-v2> [visited on 20/02/2026].
- [47] Mark A. Barton. *Advanced LIGO Suspension Models and State-Space Tools*. LIGO Technical Report T1400034. Available from the LIGO Document Control Center. LIGO Laboratory, 2014.
- [48] Luc Blanchet. ‘Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries’. In: *Living Reviews in Relativity* 9.4 [2006]. doi: [10.12942/lrr-2006-4](https://doi.org/10.12942/lrr-2006-4).
- [49] D. et al. Bohomolov. ‘Low- and High-Frequency Noise in LEDs’. In: *Physica Status Solidi* [2024].
- [50] Charlotte Bond et al. ‘Interferometer techniques for gravitational-wave detection’. In: *Living Reviews in Relativity* 19.1 [Feb. 2017], p. 3. issn: 1433-8351. doi: [10.1007/s41114-016-0002-8](https://doi.org/10.1007/s41114-016-0002-8). url: <https://doi.org/10.1007/s41114-016-0002-8>.
- [51] F. Brandi et al. ‘Measurement of the phase anisotropy of very high reflectivity interferential mirrors’. In: *Applied Physics B* 65.3 [Sept. 1997], pp. 351–355. issn: 1432-0649. doi: [10.1007/s003400050283](https://doi.org/10.1007/s003400050283). url: <https://doi.org/10.1007/s003400050283>.
- [52] Alessandra Buonanno and Yanbei Chen. ‘Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors’. In: *Physical Review D* 64.4 [2001], p. 042006. doi: [10.1103/PhysRevD.64.042006](https://doi.org/10.1103/PhysRevD.64.042006).

- [53] Craig Cahillane and Georgia Mansell. ‘Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four’. In: *Galaxies* 10.1 [2022]. issn: 2075-4434. doi: [10.3390/galaxies10010036](https://doi.org/10.3390/galaxies10010036). url: <https://www.mdpi.com/2075-4434/10/1/36>.
- [54] E Capote, L Dartez and D Davis. ‘Technical noise, data quality, and calibration requirements for next-generation gravitational-wave science’. In: *Classical and Quantum Gravity* 41.18 [2024], p. 185001. doi: [10.1088/1361-6382/ad694d](https://doi.org/10.1088/1361-6382/ad694d). url: <https://doi.org/10.1088/1361-6382/ad694d>.
- [55] E. Capote et al. ‘Advanced LIGO detector performance in the fourth observing run’. In: *Phys. Rev. D* 111 [6 2025], p. 062002. doi: [10.1103/PhysRevD.111.062002](https://doi.org/10.1103/PhysRevD.111.062002). url: <https://link.aps.org/doi/10.1103/PhysRevD.111.062002>.
- [56] L. Carbone et al. ‘Sensors and actuators for the Advanced LIGO mirror suspensions’. In: *Classical and Quantum Gravity* 29.11 [2012]. Section: 115005. issn: 0264-9381 1361-6382. doi: [10.1088/0264-9381/29/11/115005](https://doi.org/10.1088/0264-9381/29/11/115005).
- [57] J. et al Carron. ‘Cryogenic characterization of LEDs for space application’. In: *Proc. SPIE 10378, International Conference on Space Optics*. 2017. url: <https://ui.adsabs.harvard.edu/abs/2017SPIE10378E..05C/abstract>.
- [58] C. M. Caves. ‘Quantum-Mechanical Noise in an Interferometer’. In: *Physical Review D* 23.8 [1981], pp. 1693–1708. doi: [10.1103/PhysRevD.23.1693](https://doi.org/10.1103/PhysRevD.23.1693).
- [59] Whitney Clavin. *India Approves Construction of Its Own LIGO*. LIGO Laboratory, California Institute of Technology. 17th Apr. 2023. url: <https://www.ligo.caltech.edu/news/ligo20230417> [visited on 16/02/2026].
- [60] William Kingdon Clifford. ‘On the Space-Theory of Matter’. In: *Proceedings of the Cambridge Philosophical Society* 2 [1876], pp. 157–158.
- [61] S. J. Cooper et al. ‘Sensors and actuators for the Advanced LIGO A+ upgrade’. In: *Review of Scientific Instruments* 94.1 [2023], p. 014502. doi: [10.1063/5.0117605](https://doi.org/10.1063/5.0117605).
- [62] Alan Cumming. ‘Aspects of mirrors and suspensions for advanced gravitational wave detectors’. PhD thesis. 2008.
- [63] Stefan L. Danilishin and Farid Ya. Khalili. ‘Quantum Measurement Theory in Gravitational-Wave Detectors’. In: *Living Reviews in Relativity* 15.1 [2012], p. 5. doi: [10.12942/lrr-2012-5](https://doi.org/10.12942/lrr-2012-5).

- [64] Alastair Effler et al. *A Roadmap for the LIGO Observatories in the Era of Cosmic Exploration*. White paper submitted to NSF. 2023. url: <https://nsf-gov-resources.nsf.gov/files/ligo-observatories-white-paper-submission.pdf>.
- [65] A. Einstein and N. Rosen. ‘On gravitational waves’. In: *Journal of the Franklin Institute* 223.1 [1937], pp. 43–54. issn: 0016-0032. doi: [https://doi.org/10.1016/S0016-0032\(37\)90583-0](https://doi.org/10.1016/S0016-0032(37)90583-0). url: <https://www.sciencedirect.com/science/article/pii/S0016003237905830>.
- [66] Albert Einstein. ‘Naherungsweise Integration der Feldgleichungen der Gravitation’. In: *Sitzungsberichte der Koniglich Preussischen Akademie der Wissenschaften (Berlin)* [1916], pp. 688–696.
- [67] European Space Agency. *LISA factsheet*. Accessed 2026-01-31. 2024. url: https://www.esa.int/Science_Exploration/Space_Science/LISA/LISA_factsheet.
- [68] Matthew Evans et al. *A Horizon Study for Cosmic Explorer: Science, Observatories, and Community*. Tech. rep. CE–P2100003–v7. Cosmic Explorer, 2021. doi: [10.48550/arXiv.2109.09882](https://arxiv.org/abs/10.48550/arXiv.2109.09882). arXiv: [2109.09882](https://arxiv.org/abs/2109.09882) [astro-ph.IM]. url: <https://dcc.cosmicexplorer.org/public/0163/P2100003/007/ce-horizon-study.pdf> [visited on 16/02/2026].
- [69] Matthew Evans et al. ‘Observation of Parametric Instability in Advanced LIGO’. In: *Physical Review Letters* 114.16 [2015], p. 161102. doi: [10.1103/PhysRevLett.114.161102](https://doi.org/10.1103/PhysRevLett.114.161102).
- [70] J. Franc et al. ‘Mirror thermal noise in laser interferometer gravitational wave detectors operating at room and cryogenic temperature’. In: *Journal of Physics: Conference Series* 118 [2008], p. 012010. doi: [10.1088/1742-6596/118/1/012010](https://doi.org/10.1088/1742-6596/118/1/012010).
- [71] Tobin T Fricke et al. ‘DC readout experiment in Enhanced LIGO’. In: *Classical and Quantum Gravity* 29.6 [2012], p. 065005. doi: [10.1088/0264-9381/29/6/065005](https://doi.org/10.1088/0264-9381/29/6/065005). url: <https://dx.doi.org/10.1088/0264-9381/29/6/065005>.
- [72] P. Fritschel and the LIGO II Seismic Isolation Team. *LIGO II Seismic Isolation Subsystem (SEI) Design Requirements Review*. LIGO Technical Note LIGO-G010017-00-D. LIGO Laboratory, 2001. url: <https://dcc.ligo.org/LIGO-G010017/public>.

- [73] Peter Fritschel, Matthew Evans and Valery Frolov. ‘Balanced homodyne readout for quantum limited gravitational wave detectors’. en. In: *Optics Express* 22.4 [Feb. 2014], p. 4224. issn: 1094-4087. doi: [10.1364/OE.22.004224](https://doi.org/10.1364/OE.22.004224). url: <https://opg.optica.org/oe/abstract.cfm?uri=oe-22-4-4224> [visited on 14/05/2023].
- [74] D. Ganapathy et al. ‘Broadband Quantum Enhancement of the LIGO Detectors with Frequency-Dependent Squeezing’. In: *Physical Review X* 13 [2023], p. 041021. doi: [10.1103/PhysRevX.13.041021](https://doi.org/10.1103/PhysRevX.13.041021).
- [75] Anthony Gerrard and James M. Burch. *Introduction to Matrix Methods in Optics*. English. A Wiley-Interscience publication. Bibliography: p. 338–341. London, New York: Wiley-Interscience, 1975, pp. xiii + 355. isbn: 0471296856.
- [76] S. et al. Glemža. ‘Low-frequency noise of near-UV LEDs’. In: [2019].
- [77] K. Goda et al. ‘A quantum-enhanced prototype gravitational-wave detector’. In: *Nature Physics* 4.6 [2008], pp. 472–476. doi: [10.1038/nphys920](https://doi.org/10.1038/nphys920).
- [78] J. W. Goodman. *Statistical Optics*. Wiley, 2015.
- [79] N. A. Gordon. ‘Characterisation and control of coupled optical springs for future gravitational wave detectors’. PhD thesis. University of Glasgow, 2015.
- [80] S. Goßler et al. ‘The AEI 10 m prototype interferometer’. In: *Classical and Quantum Gravity* 27.8 [2010], p. 084023. doi: [10.1088/0264-9381/27/8/084023](https://doi.org/10.1088/0264-9381/27/8/084023).
- [81] Victoria Graham. ‘Development of optical suspensions and laser technology for a cryogenic interferometer prototype’. College of Science and Engineering, School of Physics and Astronomy. PhD thesis. Glasgow, UK: University of Glasgow, 2025. url: <https://theses.gla.ac.uk/84994/>.
- [82] S. Gras et al. ‘Audio-band coating thermal noise measurement for Advanced LIGO with a multimode optical resonator’. In: *Physical Review D* 95.2 [2017], p. 022001. doi: [10.1103/PhysRevD.95.022001](https://doi.org/10.1103/PhysRevD.95.022001). arXiv: [1609.05595](https://arxiv.org/abs/1609.05595).
- [83] *GWINC Documentation (pygwinc)*. LIGO Laboratory. 2021. url: <https://gwinc.docs.ligo.org/pygwinc> [visited on 16/02/2026].
- [84] Evan D. Hall, Matthew Evans et al. ‘Cosmic Explorer: A Next-Generation Ground-Based Gravitational-Wave Observatory’. In: *Galaxies* 10.4 [2022], p. 90. doi: [10.3390/galaxies10040090](https://doi.org/10.3390/galaxies10040090).
- [85] Hamamatsu Photonics. *Quadrant Photodiodes Technical Notes*. Available at: <https://www.hamamatsu.com>. 2021.

- [86] Jan Harms. ‘Terrestrial Gravity Fluctuations’. In: *Living Reviews in Relativity* 18 [2015], p. 3. doi: [10.1007/lrr-2015-3](https://doi.org/10.1007/lrr-2015-3). arXiv: [1507.05850](https://arxiv.org/abs/1507.05850) [gr-qc].
- [87] Jan Harms et al. ‘Squeezed-input, optical-spring, signal-recycled gravitational-wave detectors’. In: *Phys. Rev. D* 68 [4 2003], p. 042001. doi: [10.1103/PhysRevD.68.042001](https://doi.org/10.1103/PhysRevD.68.042001). url: <https://link.aps.org/doi/10.1103/PhysRevD.68.042001>.
- [88] G. M. Harry. ‘Thermal noise from optical coatings in gravitational wave detectors’. In: *Applied Optics* 45.7 [2006], pp. 1569–1574. doi: [10.1364/AO.45.001569](https://doi.org/10.1364/AO.45.001569).
- [89] G. M. Harry, A. M. Gretarsson, P. R. Saulson et al. ‘Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings’. In: *Classical and Quantum Gravity* 19.5 [2002], pp. 897–917. doi: [10.1088/0264-9381/19/5/305](https://doi.org/10.1088/0264-9381/19/5/305).
- [90] M. A. Herman and H. Sitter. *Molecular Beam Epitaxy: Fundamentals and Current Status*. 2nd. Springer, 1996.
- [91] Stefan Hild, Mark Abernathy, Francesco Acernese et al. ‘Sensitivity studies for third-generation gravitational wave observatories’. In: *Classical and Quantum Gravity* 28.9 [2011], p. 094013. doi: [10.1088/0264-9381/28/9/094013](https://doi.org/10.1088/0264-9381/28/9/094013). arXiv: [1012.0908](https://arxiv.org/abs/1012.0908) [gr-qc].
- [92] M. E. Husman et al. ‘Modeling of multistage pendulums: Triple pendulum suspension for GEO 600’. In: *Review of Scientific Instruments* 71.6 [2000], pp. 2546–2551. doi: [10.1063/1.1150646](https://doi.org/10.1063/1.1150646).
- [93] S. H. Huttner et al. ‘Techniques in the optimization of length sensing and control systems for a three-mirror coupled cavity’. In: *Classical and Quantum Gravity* 25.23 [2008], p. 235003. doi: [10.1088/0264-9381/25/23/235003](https://doi.org/10.1088/0264-9381/25/23/235003).
- [94] Russell Jones et al. *A+ SUS OMC Breadboard Assembly*. <https://dcc.ligo.org/LIGO-D2000172>. Version v4. LIGO Document D2000172-v4, updated 22 Aug 2023. Available at: <https://dcc.ligo.org/LIGO-D2000172>. Aug. 2023.
- [95] KAGRA Collaboration. ‘KAGRA: 2.5 generation interferometric gravitational wave detector’. In: *Nature Astronomy* 3.1 [2019], pp. 35–40. doi: [10.1038/s41550-018-0658-y](https://doi.org/10.1038/s41550-018-0658-y).
- [96] J. Kim and S. Kim. ‘High-resolution position detection using a quadrant photodiode’. In: *Optics Letters* 29 [2004], pp. 1875–1877.

- [97] H. Kogelnik and T. Li. ‘Laser Beams and Resonators’. In: *Appl. Opt.* 5.10 [1966], pp. 1550–1567. doi: [10.1364/AO.5.001550](https://doi.org/10.1364/AO.5.001550). url: <https://opg.optica.org/ao/abstract.cfm?URI=ao-5-10-1550>.
- [98] Xhesika Korovesi et al. ‘Cryogenic payloads for the Einstein Telescope: Baseline design with heat extraction, suspension thermal noise modeling, and sensitivity analyses’. In: *Physical Review D* 108.12 [2023], p. 123009. doi: [10.1103/PhysRevD.108.123009](https://doi.org/10.1103/PhysRevD.108.123009). arXiv: [2305.01419](https://arxiv.org/abs/2305.01419).
- [99] Rahul Kumar. ‘Aspects of suspension design for the development of advanced gravitational wave detectors’. Volume: PhD. PhD thesis. University of Glasgow, 2013.
- [100] Kevin Aaron Kuns. ‘Future Networks of Gravitational Wave Detectors: Quantum Noise and Space Detectors’. ProQuest ID: *Kuns_ucsb0035D14255*, *MerrittID* : *ark* : /13030/*m5wm6cdv*. Ph.D. Dissertation. University of California, Santa Barbara, 2019. url: <https://escholarship.org/uc/item/72v62v6n>.
- [101] LIGO Laboratory. *Design Requirement Document of the A+ filter cavity and relay optics*. Tech. rep. LIGO-T1800447. LIGO Document Control Center, 2020. url: https://dcc.ligo.org/public/0156/T1800447/007/FC_DRD_T1800447-v7.pdf.
- [102] LIGO Laboratory. *LIGO Voyager Upgrade: Design Concept*. LIGO Document Server. Version dated 25 Sep 2024. 2024. url: <https://docs.ligo.org/voyager/voyagerwhitepaper/main.pdf>.
- [103] LIGO Scientific Collaboration. *Unofficial sensitivity curves (ASD) for aLIGO, KAGRA, Virgo, Voyager, Cosmic Explorer, and Einstein Telescope*. LIGO Document Control Center. 2020. url: <https://dcc.ligo.org/ligo-t1500293/public> [visited on 16/02/2026].
- [104] LIGO Scientific Collaboration and Virgo Collaboration. ‘The Advanced LIGO Detectors in the Era of First Discoveries’. In: *Physical Review Letters* 116 [2016], p. 131103. doi: [10.1103/PhysRevLett.116.131103](https://doi.org/10.1103/PhysRevLett.116.131103).

- [105] LIGO Scientific Collaboration, Virgo Collaboration and KAGRA Collaboration. ‘All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data’. In: *Physical Review D* 106.10 [2022], p. 102008. doi: [10.1103/PhysRevD.106.102008](https://doi.org/10.1103/PhysRevD.106.102008). arXiv: [2201.00697 \[gr-qc\]](https://arxiv.org/abs/2201.00697). url: <https://doi.org/10.1103/PhysRevD.106.102008>.
- [106] LSC Post-O5 Study Group. *Report of the LSC Post-O5 Study Group*. Tech. rep. LIGO-T2200287-v3. LIGO Document Control Center, 2023. url: <https://dcc.ligo.org/LIGO-T2200287/public>.
- [107] LVK Collaboration. *The LVK Operations White Paper (2025 edition)*. Tech. rep. LIGO-T2400388. LIGO Document Control Center, 2025. url: https://dcc.ligo.org/public/0197/T2400388/002/20250501_opswp2024.pdf.
- [108] J. Macarthur. ‘Towards surpassing the standard quantum limit using optical springs’. PhD thesis. University of Glasgow, 2014.
- [109] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford: Oxford University Press, 2008. isbn: 9780198570745.
- [110] Michele Maggiore et al. ‘Science case for the Einstein telescope’. In: *Journal of Cosmology and Astroparticle Physics* 2020.03 [2020], p. 050. doi: [10.1088/1475-7516/2020/03/050](https://doi.org/10.1088/1475-7516/2020/03/050).
- [111] Camille Makarem et al. *A+ OMC Assembly Procedure (Final)*. Technical Note LIGO-T2500267-v1. LIGO Document Control Center; contents revised 26 Jan 2026. LIGO Scientific Collaboration, July 2025. url: <https://dcc.ligo.org/LIGO-T2500267-v1> [visited on 20/02/2026].
- [112] Fabien Matichard, Brian Lantz, Richard Mittleman et al. ‘Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance’. In: *Classical and Quantum Gravity* 32.18 [2015], p. 185003. doi: [10.1088/0264-9381/32/18/185003](https://doi.org/10.1088/0264-9381/32/18/185003). arXiv: [1502.06300 \[gr-qc\]](https://arxiv.org/abs/1502.06300).
- [113] Matthew Edward Husman. ‘SUSPENSION AND CONTROL FOR INTERFEROMETRIC GRAVITATIONAL WAVE DETECTORS’. PhD thesis. Glasgow, 2000.
- [114] L. McCuller et al. ‘Frequency-Dependent Squeezing for Advanced LIGO’. In: *Physical Review Letters* 124 [2020], p. 171102. doi: [10.1103/PhysRevLett.124.171102](https://doi.org/10.1103/PhysRevLett.124.171102).

- [115] Kirk McKenzie et al. ‘Technical limitations to homodyne detection at audio frequencies’. In: *Appl. Opt.* 46.17 [2007], pp. 3389–3395. doi: [10.1364/AO.46.003389](https://doi.org/10.1364/AO.46.003389). url: <https://opg.optica.org/ao/abstract.cfm?URI=ao-46-17-3389>.
- [116] B. J. Meers. ‘Recycling in laser-interferometric gravitational-wave detectors’. In: *Physical Review D* 38.8 [1988], pp. 2317–2326. doi: [10.1103/PhysRevD.38.2317](https://doi.org/10.1103/PhysRevD.38.2317).
- [117] John Miller et al. ‘Prospects for doubling the range of Advanced LIGO’. In: *Physical Review D* 91.6 [Mar. 2015], p. 062005. doi: [10.1103/PhysRevD.91.062005](https://doi.org/10.1103/PhysRevD.91.062005). url: <https://doi.org/10.1103/PhysRevD.91.062005>.
- [118] Vladimir Mitin. ‘Generation–Recombination Noise in Semiconductors’. In: Overview of trap kinetics and Lorentzian spectra/corner frequencies. 2010. url: https://www-ee.eng.buffalo.edu/faculty/mitin/old/Papers/VMitin_doc_1-20-10.pdf.
- [119] L. Naticchioni and on behalf of the Virgo Collaboration. ‘The payloads of Advanced Virgo: current status and upgrades’. In: *Journal of Physics: Conference Series* 957.1 [2018], p. 012002. doi: [10.1088/1742-6596/957/1/012002](https://doi.org/10.1088/1742-6596/957/1/012002). url: <https://doi.org/10.1088/1742-6596/957/1/012002>.
- [120] Ardiana Nela et al. *Crystalline Silicon Fibres in Cryogenic Suspension Systems for Next-Generation Gravitational Wave Detectors*. LIGO Document P2500545-v2. Conference proceedings for work presented at the GR24-Amaldi16 meeting. LIGO Scientific Collaboration, Aug. 2025. url: <https://dcc.ligo.org/LIGO-P2500545-v2>.
- [121] E. Oelker et al. ‘Audio-Band Frequency-Dependent Squeezing for Gravitational-Wave Detectors’. In: *Physical Review Letters* 116 [2016], p. 041102. doi: [10.1103/PhysRevLett.116.041102](https://doi.org/10.1103/PhysRevLett.116.041102).
- [122] OSI Optoelectronics. *FCI-InGaAs-Q3000 Datasheet*. Available at: <https://osioptoelectronics.com>. 2022.
- [123] V. Palenskis. ‘Low-frequency noise in high-power LEDs: a review’. In: [2021].
- [124] Vilius Palenskis, Jonas Matukas and Sandra Pralgauskaitė. ‘Review of Low-Frequency Noise Properties of High-Power White LEDs during Long-Term Aging’. In: *Materials* 15.1 [2021], p. 13. url: <https://www.mdpi.com/1996-1944/15/1/13>.

- [125] Erica R. et al Pantel. *Effects of Cryogenic Temperatures on LEDs and Optical Fiber*. NASA Technical Reports Server (NTRS), Summer Student Research Presentations. Accessed 2025-10-09. 2005. url: <https://ntrs.nasa.gov/citations/20060010268>.
- [126] Jun Hyuk et al Park. ‘Variation of the external quantum efficiency with temperature in commercial LEDs’. In: *Journal of Applied Physics* 119.2 [2016], p. 023101. url: <https://pubs.aip.org/aip/jap/article/119/2/023101/143171/Variation-of-the-external-quantum-efficiency-with>.
- [127] Guido Pizzella. *Resonant Detectors for Gravitational Waves*. Tech. rep. LNF-95/057 (PT). Presented at TAUP’95 (Toledo, Spain), 17–21 Sep 1995. Laboratori Nazionali di Frascati, 1995. url: <https://www.osti.gov/etdeweb/servlets/purl/434963>.
- [128] M. V. Plissi et al. ‘GEO 600 triple pendulum suspension system: Seismic isolation and control’. In: *Review of Scientific Instruments* 71.6 [2000]. Section: 2539, pp. 2539–2545. issn: 0034-6748 1089-7623. doi: [10.1063/1.1150645](https://doi.org/10.1063/1.1150645).
- [129] Henri Poincaré. ‘Sur la dynamique de l’électron’. In: *Comptes Rendus Hebdomadaires des Séances de l’Académie des Sciences* 140 [1905], pp. 1504–1508.
- [130] Subhra Priyadarshini. ‘India to begin construction of gravitational wave project’. In: *Nature India* [17th Apr. 2025]. News Feature. doi: [10.1038/d44151-025-00061-x](https://doi.org/10.1038/d44151-025-00061-x). url: <https://www.nature.com/articles/d44151-025-00061-x> [visited on 14/10/2024].
- [131] David L. Pulfrey. *Photonic Devices*. 2nd. Cambridge University Press, 2021.
- [132] M. Punturo et al. ‘The Einstein Telescope: A third-generation gravitational wave observatory’. In: *Classical and Quantum Gravity* 27.19 [2010], p. 194002. doi: [10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002).
- [133] *pygwinc: Gravitational Wave Interferometer Noise Calculator*. LIGO Laboratory. 2017. url: <https://git.ligo.org/gwinc/pygwinc> [visited on 16/02/2026].
- [134] Stuart Reid and Iain W. Martin. ‘Development of Mirror Coatings for Gravitational Wave Detectors’. In: *Coatings* 6.4 [2016]. issn: 2079-6412. doi: [10.3390/coatings6040061](https://doi.org/10.3390/coatings6040061). url: <https://www.mdpi.com/2079-6412/6/4/61>.

- [135] David Reitze et al. *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*. 2019. doi: [10.48550/arXiv.1907.04833](https://doi.org/10.48550/arXiv.1907.04833). arXiv: [1907.04833](https://arxiv.org/abs/1907.04833) [astro-ph.IM].
- [136] Keith Riles. ‘Searches for continuous-wave gravitational radiation’. In: *Living Reviews in Relativity* 26.1 [2023], p. 3. doi: [10.1007/s41114-023-00044-3](https://doi.org/10.1007/s41114-023-00044-3). url: <https://doi.org/10.1007/s41114-023-00044-3>.
- [137] D. I. Robertson et al. ‘The Glasgow 10 m prototype laser interferometric gravitational wave detector’. In: *Review of Scientific Instruments* 66.9 [1995], pp. 4447–4452. doi: [10.1063/1.1145339](https://doi.org/10.1063/1.1145339).
- [138] N. A. Robertson et al. *Advanced LIGO Suspension System Conceptual Design*. LIGO Technical Note LIGO-T010103. LIGO Laboratory, 2001. url: <https://dcc.ligo.org/LIGO-T010103/public>.
- [139] N. A. Robertson et al. ‘Quadruple suspension design for Advanced LIGO’. In: *Classical and Quantum Gravity* 19.15 [2002], pp. 4043–4058. doi: [10.1088/0264-9381/19/15/311](https://doi.org/10.1088/0264-9381/19/15/311).
- [140] Antoni Rogalski. ‘Infrared detectors: status and trends’. In: *Progress in Quantum Electronics* 27.2-3 [2003], pp. 59–210.
- [141] S. et al. Romyantsev. ‘Low-frequency noise and long-term stability of noncoherent light sources’. In: [2004].
- [142] M. Saleem et al. ‘The science case for LIGO–India’. In: *Classical and Quantum Gravity* 39.2 [20th Jan. 2022], p. 025004. doi: [10.1088/1361-6382/ac3b99](https://doi.org/10.1088/1361-6382/ac3b99).
- [143] Bahaa E. A. Saleh and Malvin Carl Teich. *Fundamentals of Photonics*. Wiley, 1991.
- [144] B.E.A. Saleh and M.C. Teich. *Fundamentals of Photonics*. 2nd. Wiley, 2007.
- [145] B. S. Sathyaprakash and Bernard F. Schutz. ‘Physics, Astrophysics and Cosmology with Gravitational Waves’. In: *Living Reviews in Relativity* 12.2 [2009]. doi: [10.12942/lrr-2009-2](https://doi.org/10.12942/lrr-2009-2). url: <https://link.springer.com/article/10.12942/lrr-2009-2>.
- [146] P. R. Saulson. ‘Thermal noise in mechanical experiments’. In: *Physical Review D* 42.8 [1990], pp. 2437–2445. doi: [10.1103/PhysRevD.42.2437](https://doi.org/10.1103/PhysRevD.42.2437).
- [147] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, 1994.
- [148] E. Fred Schubert. *Light-Emitting Diodes*. 2nd. Cambridge University Press, 2006.

- [149] Bernard F. Schutz. ‘Gravitational Radiation’. In: [2000]. arXiv:gr-qc/0003069. arXiv: [gr-qc/0003069](https://arxiv.org/abs/gr-qc/0003069) [gr-qc].
- [150] Thejas Seetharamu. *A+ OMC Subassembly Matching*. Tech. rep. LIGO-T2400229-v1. Notebook for matching PZT wedge angle, prism deviation, and mirror curvature radius for A+ OMC subassemblies, using 2023 batch data. LIGO Laboratory, 2024. url: <https://dcc.ligo.org/LIGO-T2400229/public>.
- [151] Thejas Seetharamu. *A+ OMC Subassembly Matching*. Technical Note T2400229-v1. <https://dcc.ligo.org/LIGO-T2400229>. LIGO Document Control Center, 2024.
- [152] Brett N. Shapiro. ‘Modal control with state estimation for Advanced LIGO quadruple suspensions’. Master’s thesis. MA thesis. Cambridge, MA, USA: Massachusetts Institute of Technology, 2008.
- [153] A.E. Siegman. *Lasers*. G - Reference, Information and Interdisciplinary Subjects Series. University Science Books, 1986. isbn: 9780935702118. url: <https://books.google.co.uk/books?id=1BZVwUZLTkAC>.
- [154] K. D. Skeldon et al. ‘Performance of the Glasgow 10 m prototype gravitational wave detector operating at 1064 nm’. In: *Physics Letters A* 271.1–2 [2000], pp. 7–12. doi: [10.1016/S0375-9601\(00\)00511-9](https://doi.org/10.1016/S0375-9601(00)00511-9).
- [155] K. Somiya. ‘Detector configuration of KAGRA: The Japanese cryogenic gravitational-wave detector’. In: *Classical and Quantum Gravity* 29.12 [2012], p. 124007. doi: [10.1088/0264-9381/29/12/124007](https://doi.org/10.1088/0264-9381/29/12/124007).
- [156] Clive C Speake and Peter L Andrews. ‘Capacitive sensing for drag-free satellites’. In: *Classical and Quantum Gravity* 14.6 [1997], p. 1557. doi: [10.1088/0264-9381/14/6/027](https://doi.org/10.1088/0264-9381/14/6/027). url: <https://doi.org/10.1088/0264-9381/14/6/027>.
- [157] Andrew P Spencer. ‘Advanced Techniques in Laser Interferometry for Current and Future Gravitational Wave Detectors’. en. PhD thesis. University Of Glasgow, 2020. url: <https://theses.gla.ac.uk/id/eprint/81266>.
- [158] M S Stefszky et al. ‘Balanced homodyne detection of optical quantum states at audio-band frequencies and below’. In: *Classical and Quantum Gravity* 29.14 [2012], p. 145015. doi: [10.1088/0264-9381/29/14/145015](https://doi.org/10.1088/0264-9381/29/14/145015). url: <https://dx.doi.org/10.1088/0264-9381/29/14/145015>.

- [159] Sebastian Steinlechner et al. ‘Local-oscillator noise coupling in balanced homodyne readout for advanced gravitational wave detectors’. In: *Phys. Rev. D* 92 [7 2015], p. 072009. doi: [10.1103/PhysRevD.92.072009](https://doi.org/10.1103/PhysRevD.92.072009). url: <https://link.aps.org/doi/10.1103/PhysRevD.92.072009>.
- [160] K. A. Strain and B. J. Meers. ‘Experimental demonstration of dual recycling for interferometric gravitational-wave detectors’. In: *Physical Review Letters* 66.11 [1991], pp. 1391–1394. doi: [10.1103/PhysRevLett.66.1391](https://doi.org/10.1103/PhysRevLett.66.1391).
- [161] K. A. Strain and B. N. Shapiro. ‘Damping and local control of mirror suspensions for laser interferometric gravitational wave detectors’. In: *Rev Sci Instrum* 83.4 [Apr. 2012]. Edition: 2012/05/09, p. 044501. issn: 1089-7623 (Electronic) 0034-6748 (Linking). doi: [10.1063/1.4704459](https://doi.org/10.1063/1.4704459). url: <https://www.ncbi.nlm.nih.gov/pubmed/22559557>.
- [162] Ken Strain and Russell Jones. *OMC FDR Update*. Technical Report LIGO-T2100078-v2. Updated August 2021. Available at <https://dcc.ligo.org/LIGO-T2100078-v2>. LIGO Project, 2021.
- [163] Ben G. Streetman and Sanjay Banerjee. *Solid State Electronic Devices*. 7th. Pearson, 2015.
- [164] Marek J. Szczepańczyk et al. ‘Optically targeted search for gravitational waves emitted by core-collapse supernovae during the third observing run of Advanced LIGO and Advanced Virgo’. In: *Physical Review D* 110.4 [2024], p. 042007. doi: [10.1103/PhysRevD.110.042007](https://doi.org/10.1103/PhysRevD.110.042007). arXiv: [2305.16146](https://arxiv.org/abs/2305.16146) [astro-ph.HE].
- [165] S.M. Sze and K.K. Ng. *Physics of Semiconductor Devices*. 3rd. Wiley-Interscience, 2006.
- [166] Hareem Tariq et al. ‘The linear variable differential transformer (LVDT) position sensor for gravitational wave interferometer low-frequency controls’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 489.1–3 [Aug. 2002], pp. 570–576. doi: [10.1016/S0168-9002\(02\)00802-1](https://doi.org/10.1016/S0168-9002(02)00802-1). url: [https://doi.org/10.1016/S0168-9002\(02\)00802-1](https://doi.org/10.1016/S0168-9002(02)00802-1).
- [167] Thorlabs, Inc. *LED122E Product Documentation*. Accessed: Jul. 2025. 2025. url: <https://www.thorlabs.com>.

- [168] Karl Toland et al. ‘Design and development of a prototype single crystalline silicon cryogenic suspension for 3rd generation gravitational wave detectors’. In: *Classical and Quantum Gravity* 42.17 [2025], p. 175006. doi: [10.1088/1361-6382/adf8a9](https://doi.org/10.1088/1361-6382/adf8a9). url: <https://doi.org/10.1088/1361-6382/adf8a9>.
- [169] C. Torrie. ‘DEVELOPMENT OF SUSPENSIONS FOR THE GEO 600 GRAVITATIONAL WAVE DETECTOR’. PhD thesis. Nov. 1999. url: <https://theses.gla.ac.uk/76201/>.
- [170] M. Tse et al. ‘Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy’. In: *Physical Review Letters* 123 [2019], p. 231107. doi: [10.1103/PhysRevLett.123.231107](https://doi.org/10.1103/PhysRevLett.123.231107).
- [171] Amit Singh Ubhi. *Cryogenic optical shadow sensors for future gravitational wave detectors*. May 2022.
- [172] N. Uehara and K. Ueda. ‘Accurate measurement of the radius of curvature of a concave mirror and the power dependence in a high-finesse Fabry–Perot interferometer’. en. In: *Applied Optics* 34.25 [Sept. 1995], p. 5611. issn: 0003-6935, 1539-4522. doi: [10.1364/AO.34.005611](https://doi.org/10.1364/AO.34.005611). url: <https://opg.optica.org/abstract.cfm?URI=ao-34-25-5611> [visited on 16/02/2023].
- [173] Takafumi Ushiba. ‘Personal communication (email)’. Email to Thejas Seetharamu, 16 February 2026. 16th Feb. 2026.
- [174] Takafumi Ushiba et al. ‘Cryogenic suspension design for a kilometer-scale gravitational-wave detector’. In: *Classical and Quantum Gravity* 38.8 [Mar. 2021], p. 085013. doi: [10.1088/1361-6382/abe9f3](https://doi.org/10.1088/1361-6382/abe9f3). url: <https://iopscience.iop.org/article/10.1088/1361-6382/abe9f3>.
- [175] Y. P. Varshni. ‘Temperature dependence of the energy gap in semiconductors’. In: *Physica* 34.1 [1967], pp. 149–154.
- [176] Anna-Maria A. van Veggel. ‘Quasi-monolithic mirror suspensions in ground-based gravitational-wave detectors: an overview and look to the future’. In: *Philosophical Transactions of the Royal Society A* 376.2120 [2018], p. 20170281. doi: [10.1098/rsta.2017.0281](https://doi.org/10.1098/rsta.2017.0281).
- [177] Virgo Collaboration. ‘Advanced Virgo: a second-generation interferometric gravitational wave detector’. In: *Classical and Quantum Gravity* 32.2 [2015], p. 024001. doi: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001).

- [178] R. L. Ward et al. ‘DC readout experiment at the Caltech 40 m prototype interferometer’. In: *Classical and Quantum Gravity* 25.11 [2008], p. 114030. doi: [10.1088/0264-9381/25/11/114030](https://doi.org/10.1088/0264-9381/25/11/114030).
- [179] Nils Leander Weickhardt, Artem Basalaeu and Oliver Gerberding. ‘Reducing suspension control noise with interferometric sensors – an experimental concept’. In: *arXiv preprint* [2025]. arXiv: [2507.18369](https://arxiv.org/abs/2507.18369) [gr-qc]. url: <https://arxiv.org/abs/2507.18369>.
- [180] Rainer Weiss. ‘LIGO and the Discovery of Gravitational Waves I’. In: *Reviews of Modern Physics* 90 [2018], p. 040501. doi: [10.1103/RevModPhys.90.040501](https://doi.org/10.1103/RevModPhys.90.040501).
- [181] B. Willke et al. ‘The GEO 600 gravitational wave detector’. In: *Classical and Quantum Gravity* 19.7 [2002], pp. 1377–1387. doi: [10.1088/0264-9381/19/7/321](https://doi.org/10.1088/0264-9381/19/7/321).
- [182] J. L. Wright. ‘Optical springs to create macroscopic optical traps and negative inertia for gravitational wave detectors’. PhD thesis. University of Glasgow, 2022.
- [183] Liyuan Zhang. *A+ OMC, Component Testing, Total Integrated Scatter (TIS) and Transmittance Measurements*. LIGO Document LIGO-E2300029-v1. LIGO Laboratory, Feb. 2023. url: <https://dcc.ligo.org/LIGO-E2300029-v1>.
- [184] C. Zhao et al. ‘Gingin High Optical Power Test Facility’. In: *Journal of Physics: Conference Series*. Vol. 32. Proceedings of the Sixth Edoardo Amaldi Conference on Gravitational Waves. 2006, pp. 368–373. doi: [10.1088/1742-6596/32/1/056](https://doi.org/10.1088/1742-6596/32/1/056).
- [185] M. Selim Ünlü and Manijeh Razeghi. ‘Cryogenic operation of InP/InGaAsP light-emitting devices’. In: *IEEE Journal of Quantum Electronics* 27.3 [1991], pp. 688–696.



Seetharamu, Thejas (2026) *Advancement of cryogenic suspensions and optical cavities for gravitational wave detection*. PhD thesis.

<https://theses.gla.ac.uk/85942/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk>

research-enlighten@glasgow.ac.uk