

Applications of Precision Interferometry in Quantum Measurement and Gravitational Wave Detection

Author:
Jiří Smetana

Supervisor:
Dr. Denis Martynov

A thesis submitted to the University of Birmingham for the degree of
Doctor of Philosophy

Institute for Gravitational Wave Astronomy
School of Physics & Astronomy
College of Engineering and Physical Sciences
University of Birmingham
January 26, 2023

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Declaration of Authorship

I, Jiří SMETANA, declare that this thesis titled, "Applications of Precision Interferometry in Quantum Measurement and Gravitational Wave Detection" and the work presented in it are my own. I confirm that:

- This work was done wholly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: 

Date: January 26, 2023

UNIVERSITY OF BIRMINGHAM

Abstract

College of Engineering and Physical Sciences
School of Physics & Astronomy

Doctor of Philosophy

Applications of Precision Interferometry in Quantum Measurement and Gravitational Wave Detection

by Jiří SMETANA

Interferometry sits at the forefront of precision measurement. This versatile technique is used to detect the smallest signals in the universe, including microscopic fluctuations of quantum phenomena, and gravitational waves. In this thesis, I present the development of (i) new interferometer configurations, (ii) new interferometric sensors, and (iii) quantum-limited systems of fundamental physics importance.

Concerning contribution (i), I present work on two new interferometer configurations for gravitational-wave detection. (a) Theoretical analysis of an optomechanical filter cavity that acts as a quantum amplifier. The purpose of the amplifier is to enhance the quantum-limited sensitivity of future detectors. The technique uses a dual frequency- and time-domain approach to verify the effect. (b) A high-frequency design that is capable of measuring the signal from a post-merger neutron star collapse. My contribution concerns the modelling of classical noise sources in the region of the spectrum where the conventional noise models lose their validity.

The device comprising contribution (ii) is a compact, interferometric displacement sensor. The position sensors and inertial sensors on contemporary detectors are limited by readout noise, which couples into the gravitational-wave readout channel. The sensor presented utilises the sharp response of a Michelson interferometer, with the additional use of deep frequency modulation. I show that the device improves upon existing displacement sensors by a factor of 300.

Contribution (iii) concerns the measurement of quantum noise over a broad band in a macroscopic system. Future gravitational-wave detectors will be able to operate at a quantum-limited level of sensitivity at the level of the so-called standard quantum limit. However, this limit has not yet been reached, despite the opportunities to study matters of fundamental physics importance that such a system would bring. I present a high-finesse, suspended optical cavity theoretically capable of reaching this level of operation. I analyse the performance of the main aspects of the experiment and present initial results from the operation of the cavity.

Acknowledgements

I would like to thank my supervisor, Dr Denis Martynov, for his leadership, valuable feedback, willingness to explain (and re-explain), and motivating enthusiasm to push projects forward.

I appreciate the assistance of Prof. Alberto Vecchio, Prof. Andreas Freise, Dr Haixing Miao, Dr Clive Speake, and Dr Conor Mow-Lowry, who have advised me on practical matters in the lab and given me guidance throughout my studies.

I would also like to extend my thanks to my colleagues, Aaron, Alex, Artemiy, Chiara, Joe, Leo, Riccardo, Sam, Teng, and Tianliang for their help in the lab and for creating an environment in which it was a pleasure to work. A particular mention must be made of Amit Ubhi, whose countless trips to the whiteboard have solved many a conundrum.

Thank you to David Stops for devoting many hours that have been irrecoverably lost to dealing with my relentless stream of IT requests and to Jo Cox for being a shield against the torrent of bureaucracy.

Finally, I would like to acknowledge the support of my friends and family, particularly during the challenges of the pandemic. Most of all, I would like to thank my partner and future wife, Emma Willett, for her love and care, for her suggestions and comments, and, above all, for her patience.

Contents

Declaration of Authorship v

Abstract vii

Acknowledgements ix

Glossary xli

1 Introduction 1

1.1 Framework 2

1.2 Interferometry 4

1.2.1 Michelson Interferometer 4

1.2.2 Fabry-Perot Interferometer 7

1.2.3 Gaussian Beams 10

1.2.4 Cavity Mode and Stability 12

1.2.5 Locking Scheme 14

Pound-Drever-Hall Locking 15

1.3 Gravitational Wave Detection 17

1.4 Thesis Structure 21

2 Sensitivity of Interferometric Measurements 25

2.1 Quantum Noise 27

2.1.1 Approximate Derivation 27
Shot Noise 27
Radiation Pressure Noise 28
Standard Quantum Limit 30
2.1.2 Quantum Mechanical Derivation 31
Shot Noise 32
Radiation Pressure Noise 33
Standard Quantum Limit 34
2.2 Seismic Noise 35
2.2.1 Passive Isolation 35
Spring-Based Isolation 37
Pendulum-Based Isolation 42
Non-Linear Coupling 45
2.2.2 Suspensions 47
2.2.3 Damping 48
2.2.4 Active Isolation 54
2.3 Thermal Noise 56
2.3.1 Suspension Thermal Noise 59
2.3.2 Substrate Brownian Noise 62
2.3.3 Coating Brownian Noise 65
2.4 Laser Noise 67
2.4.1 Laser Frequency Noise 68
2.4.2 Laser Intensity Noise 70
2.5 Summary 72

3 Macroscopic Quantum State: Motivation and Suspension **75**

3.1 Outline . . . 76

3.1.1 Contemporary Sensitivity . . . 76

3.1.2 SQL in Future Devices . . . 77

3.1.3 Beating the SQL . . . 78

3.1.4 Experimental Proposal . . . 80

3.1.5 Experimental Schematic . . . 83

3.2 Seismic Isolation Overview . . . 84

3.3 Optical Table Isolation . . . 86

3.4 Cryostat Frame . . . 88

3.5 Internal Suspension . . . 91

3.5.1 Horizontal Performance—Separate Suspension . . . 91

3.5.2 Horizontal Performance—Combined Suspension . . . 98

3.5.3 Vertical Performance . . . 100

3.5.4 Ancillary Degrees of Freedom . . . 105

3.5.5 Damping . . . 107

3.6 Practical Design . . . 110

3.6.1 Design Parameters and Transfer Functions . . . 113

4 Macroscopic Quantum State: Thermal Noise and Results **117**

4.1 Thermal Considerations . . . 117

4.1.1 Cooling Efficiency . . . 118

4.1.2 Upper Suspension . . . 121

4.1.3 Substrate Choice . . . 124

4.1.4 Test Mass Suspension . . . 125

4.1.5 Coating Choice . . . 126

4.2 Laser and Cavity 127
4.3 Final Design Noise Budget 129
4.4 Status and Results—Room Temperature 129
4.4.1 Experimental Layout 130
4.4.2 Room-Temperature Noise Budget 131
4.4.3 Laser Frequency—Actuation and Noise 134
4.4.4 Cavity Stability and Alignment 135
4.4.5 Lock Acquisition 136
4.4.6 Cavity Noise Spectrum 140
4.5 Concluding Remarks and Future Work 142

5 Optomechanical Filters for Interferometric Measurements 145
5.1 Introduction 146
5.2 From single-mode approximation to full analysis 149
5.3 Frequency-domain analysis 156
5.3.1 Formalism 157
5.3.2 Noise Spectral Density 163
5.3.3 Stability Analysis 165
5.4 Time-domain analysis 168
5.4.1 Discretisation of Equations 169
5.4.2 Connecting the System 171
5.4.3 Step-Response Stability Analysis 173
5.4.4 Numerical Noise Spectral Density 175
5.5 Optimal readout scheme 177
5.6 Conclusions 179

6 Modelling Classical Noise in High-Frequency Detectors **181**

6.1 Introduction . . . 182

6.2 Collapse Waveform . . . 185

6.3 Detection strategies . . . 190

6.4 Quantum noise and quantum limit from optical losses at the FSR . . . 193

6.5 Classical noises . . . 199

6.5.1 Residual gas noise . . . 199

6.5.2 Mirror thermal noise . . . 200

6.6 Modal Analysis . . . 203

6.7 Summary . . . 209

7 Compact Interferometry for Precision Displacement Measurement **213**

7.1 Introduction . . . 214

7.2 Optical layout . . . 218

7.3 Experimental results . . . 221

7.3.1 Environmental disturbances . . . 221

7.3.2 Laser frequency noise . . . 223

7.3.3 Down-converted noises . . . 224

7.3.4 Linearity . . . 225

7.4 Conclusion . . . 227

8 Conclusion **231**

A Relations in Interferometry **237**

A.1 Michelson Interferometer—Output Power . . . 237

A.2 Fabry-Pérot Interferometer—Output Power . . . 239

B Relations in Interferometer Sensitivity **241**

B.1 Offsets in Linear Systems 241

B.2 Direct Approach to Thermal Noise—Point-Mass Suspension 242

C Cryogenic Experiment **245**

C.1 Optimised Blade Design 245

C.2 Vertical Wire Stiffness 247

D Technical Drawings **249**

D.1 Cryostat Support Frame 249

D.2 Room-Temperature Suspension Prototype 252

List of Figures

1.1 Panel (a) shows the schematic diagram of a simple Michelson interferometer. Panel (b) shows the response with differential distance ($x_2 - x_1$ in Eq. 1.6). The annotations define the bright and dark fringes, which correspond to the state of constructive and destructive interference, respectively. 5

1.2 Panel (a) shows the schematic diagram of a simple Fabry-Pérot interferometer. Panel (b) shows the response (power measured at the two photodiodes) of a nominal Fabry-Pérot interferometer. The top plot shows the reflected power and the bottom plot shows the transmitted power. The parameters that give this particular response curve are: $T_1 = 20\%$, $T_2 = 40\%$, and $\lambda = 1550\,\text{nm}$. 9

1.3 Close up of a two-mirror-cavity error signal (lower plot) in the region around a single resonance. The transmitted power for the given region is shown in the top plot and corresponds to an enlarged region around the first resonance in Fig. 1.2b. 17

1.4 Centre: Schematic Diagram of a contemporary gravitational-wave detector featuring the key optics—the power recycling mirror (PRM), signal recycling mirror (SRM), the input test mass (ITM), and the end test mass (ETM). Surrounding the central interferometer schematic is a representation of a gravitational wave passing through a simple Michelson interferometer with the cyclical effect on the displacement of the arm optics shown. 20

2.1 Schematic diagram of the relevant fields and parameters in the quantum mechanical derivation of quantum noise in a Fabry-Pérot interferometer. 32

2.2 Quantum noise in an example system separated into the two constituent contributions of quantum shot noise and quantum radiation pressure noise. The particular system being modelled consists of a tuned Fabry-Pérot interferometer of length 10 cm with 100 g mirrors and an input coupler transmissivity of 10 ppm. The input laser power has been tuned to place the frequency at which the total noise intersects with the standard quantum limit at 100 Hz. 36

2.3 Panel (a) shows the schematic diagram of a simple spring-based isolation system. Panel (b) shows the Bode plot of the transfer function for the vertical displacement at the attachment point to the displacement of the mass at the bottom of the spring. The three areas indicated correspond to regions where the motion is (i) directly transferred from suspension point to mass, (ii) amplified by the resonant oscillations of the mass-spring system, and (iii) suppressed by the naturally diminishing response of the system above its resonance. 38

2.4 Schematic diagram of a simple pendulum-based isolation system. 42

2.5 Magnitude of the transfer functions of a double-stage and single-stage suspension. Both suspensions modelled consist of a total of 2 kg of mass suspended over a total length of 2 m. In the double-stage case, the mass and length are equally distributed between the two stages. One point of note is the location of the resonances, where the lower resonance of the double-stage suspension has been slightly shifted with respect to the standard 2 m pendulum resonance due to the interaction between the two stages. Another aspect is the different high-frequency slope, which is steeper for the double-stage design, and thus provides more isolation. 49

2.6 Bode plot of the transfer function of a generic single-stage isolation system with a pole frequency at 1 Hz. The system is shown in the case of no damping, structural damping, and viscous damping. The quality factor of the two damped examples is set at 10. This figure serves to illustrate two key effects of the damping system. Firstly, the presence of either damping source ensures a finite response of the system at the resonance with a smooth and continuous phase transition between the two discontinuous levels of the undamped response. Secondly, the two damping strategies show differing behaviour at frequencies well above the resonance. Whilst the structural damping approaches the undamped system in both magnitude and phase behaviour, the viscous damping curve follows a shallower gradient and a correspondingly lower phase lag of only 90°. 52

2.7 Diagram of an elementary feedback system. In the context of seismic isolation, the plant refers to the passive suspension and the servo refers to an electronic controller that accepts input from a displacement sensor and outputs to a force actuator on the suspended mass. 54

2.8 Bode plot of the transfer functions for a passive system that is undamped, experiencing viscous damping, and actively damped through feedback. The passive curves correspond to the system in Fig. 2.6. The active curve corresponds to Eq. 2.37 with an applied viscous-damping-type servo transfer function and cut-off at high frequency with a simple, single-pole low-pass filter at 5 Hz. This (by no means optimal) system serves to illustrate the key effect that the high-frequency isolation scaling of the undamped system can be recovered with active control. 57

3.1 Timeline of experimental milestones. The timeline shows notable past events at the point they were achieved and an estimate of future milestones according to our predicted experimental progress. 81

3.2 Panel (a) shows a simplified sketch of the key aspects of the system. The cryostat is shown placed atop the table connected to the grund with rigid legs. The cryocooler head is connected to a compressor via helium lines. Panel (b) shows the schematic of the inner components of the cryostat system in relation to the optical system (thermal shield removed for simplicity). The suspension is shown in its place isolating the cavity at the bottom of the vacuum chamber from the vibrations of the ground and the cryocooler that couple in through the 10 K plate at the top. This plate acts as the heat sink for the conductive cooling of the cavity mirrors via the suspension wires. The cavity is pumped and read out via a laser and photodetectors, respectively, placed in-air. 83

3.3 Horizontal, seismically-induced displacement spectrum of a table top connected with rigid legs. Although this motion is not the most direct measurement of the ground, it serves the purpose of later analysis better, such as the performance analysis of the seismically isolating suspension, whose input is better represented by the motion of the table over the motion of the floor. 85

3.4 Computer-aided design (CAD) drawing of the cryostat support frame in two different configurations: passive (left), and actively controlled (right). The frame is shown in the context of the surrounding infrastructure, such as the table and the supported cryostat. The frame itself is colourised in red to highlight the components that belong to the frame assembly. The frame is designed to allow for easy transition between the two modes of operation, whereby the pillars connecting the two hexagonal support plates can be replaced with piezo actuators in situ. 89

3.5 Panel (a) shows a sketch representation of the suspension, highlighting the individual stages and key components such as the suspension wires and blades. The sketch defines the naming convention used to refer to the individual stages, that is to say, the cryostat cold plate (CRYO), the top stage mass (TOP), the upper intermediate mass (UIM), the penultimate mass (PUM) and the two test masses (TM). Panel (b) shows how the suspension design is abstracted into simplified models for Lagrangian analysis. These models are further separated into the model used for the analysis of vertical input motion on the left and the model used for horizontal input motion on the right. The figure also defines the parameters and coordinates used in the relevant equations. 92

3.6 Changes in the resonant frequencies of a two-stage pendulum-based suspension. The legend is common to both plots. The dashed lines represent the resonances of a single pendulum of the respective length. Panel (a) shows the effect of changing the ratio of the lower mass to the upper mass. Panel (b) shows the effect of changing the ratio of the lower suspension length to the upper suspension length. 97

3.7 Top view of the two different profiles of blade spring. The intermediate stage blades are designed to span almost the full diameter of the cryostat vacuum chamber, whilst the CRYO stage blades are designed to be arranged in a triangular configuration. Further details of the design are found in Appendix C.1. 103

3.8 Comparison of differently located damping implemented on the same suspension. The ground-upper, upper-lower and ground-lower plots refer to configurations (a)–(c) respectively. The RHS plot is a zoomed view around the resonances shown in the LHS plot. 110

3.9 Design drawing of the room-temperature prototype suspension from a top- and bottom-angled viewpoint. This design incorporates all of the aspects discussed in this chapter, with four stages of horizontal isolation and three stages of vertical isolation. The horizontal stages consist of the three ring-like, suspended masses and the final, test mass stage for each cavity mirror. A more detailed breakdown of the suspension design including the assembly procedure of key components is shown in Appendix D.2. 112

3.10 Key transfer functions into the DCAV readout channel. The single-suspension transfer function shows the result of the simplified model in Section 3.5.1 with no damping or elastic restoring force present. The horizontal-horizontal transfer function shows all effects combined. The ancillary degrees of freedom have the locations of their resonances represented. 115

4.1 Heat extracted at the 10 K plate inferred from temperature measurements of the cold plate and sample plate connected via copper rods. The heat extraction power does not change significantly over with temperature except at the two extremes of the temperature range. The temperature is swept over from high to low, which means that in the initial stage the cryocooler pump was not operating at maximum power. Therefore, the drop in heat load at high temperature is due to transient behaviour. The low-temperature fluctuations represent noise due to insufficient data resolution in this region. According to the manufacturer specification, the heat extraction should begin to decline at 20 K, reaching zero at around 10 K. 120

4.2 Noise budget for the final, cryogenic design capable of reaching the SQL and enabling investigations of macroscopic quantum mechanics. 129

4.3 Optical layout of the experiment for the room-temperature prototype in the configuration discussed throughout Section 4.4. The two cavities are controlled independently with two separate but identical control systems. The cavity displacement extraction via the beating of the two cavities is not yet present in this layout. 131

4.4 Detailed layout of the locking electronics, shown as (4) in Fig. 4.3. Discussion of the parameters of the components can be found in Section 4.4.5. 132

4.5 Photo of the experimental set-up at the time of the measurements laid out in Section 4.4. The photo shows the yellow cryostat on its tetrahedral support frame in the passive configuration, as shown in Fig. 3.4. The optical components in the lower portion of the photo correspond to the components laid out on the RHS of the vacuum chamber in Fig. 4.3. The electronics in both lower corners of the photo make up the locking servo laid out in Fig. 4.4. 132

4.6 Photo of the suspension installed in its operating position within the vacuum chamber interior. The suspension is assembled as shown in Fig. 3.9, with the three ring-like masses clearly identifiable and one of the test masses visible just below the PUM. The blade springs are shown in the upwards pre-curved configuration under varying levels of load before fine-tuning the balance. 133

4.7 Noise budget for the room-temperature prototype of the experiment capable of operating at a QRPN-dominated level of sensitivity. 134

4.8 Panel (a) shows the laser frequency noise ASD as measured by a beat measurement between the two near-identical lasers. The noise of a single laser is estimated by dividing the beat note ASD by a factor of $\sqrt{2}$, which is included in the plot. Panel (b) shows the laser response bode plot for the PZT (slow channel) and current (fast channel) drives. 136

4.9 Bode plot of the open-loop transfer function constructed from the individual transfer functions of the PZT and current drives. The transfer functions are constructed from the cavity response, the particular responses of each channel (shown in Fig. 4.8b) and a further low-pass applied to the PZT. 138

4.10 Bode plot of the ratio of the PZT and current transfer functions shown in Fig 4.9. The cross-over between the two transfer functions occurs at the 0 dB crossing and is indicated with the vertical dashed line. The corresponding location is indicated on the phase plot and determines the phase margin at the blend. The relative gain is optimised to place the cross-over frequency at the location of the maximum phase margin. . . 139

4.11 Cavity displacement ASD with principal noise sources modelled for comparison. The residual seismic noise spectrum is estimated from the table motion spectrum multiplied through by the directly coupled horizontal-horizontal (H-H) and the cross-coupled vertical-horizontal (V-H) transfer functions. The key result is the relatively close agreement on the slope of the pre- and post-resonance region of the seismic noise and the agreement on the location of many of the resonances. Towards the higher frequency region, the laser frequency noise begins to dominate over the cavity spectrum as expected. 141

5.1 A schematics of the interferometer configuration with the signal-recycling cavity (SRC) as the optomechanical filter. Both the main arm cavities and the SRC resonant at ω_0, but they have different free spectral ranges due to the different cavity lengths. One of the mirrors in the SRC acts as the mechanical degree of freedom that resonant at ω_{m_0}. It couples to the signal sidebands around ω_0 via the radiation pressure due to the blue detuned SRC pumping field at $\omega_0 + \omega_m$. Note that $\omega_{m_0}/\omega_m \approx 1$, with slight difference caused by the optical spring effect due to the blue-detuned pumping (will be explained later near Eq. 5.17). ITM: input test mass, SRM: signal-recycling mirror. 148

5.2 Idealised mode interaction structure of the optomechanical system illustrated in Fig. 5.1. **(a)** The sWLC scheme. Here $\hat{a}$ represents the differential mode of the arm cavity which couples to the GW strain h, $\hat{b}$ represents the SRC (filter) mode, and $\hat{c}$ represents the mechanical mode of the suspended mirror. The coupling between $\hat{a}$ and $\hat{b}$ ($\hat{b}$ and $\hat{c}$) is characterised by ω_s (G). Mode $\hat{b}$ is coupled to the external ingoing field $\hat{b}_{\rm in}$, which carries the vacuum noise, and the outgoing field $\hat{b}_{\rm out}$, which carries the signal and will be read out. **(b)** The uWLC scheme for comparison. In the intracavity readout scheme presented in Ref. [163], the external fields couple to the arm cavity mode $\hat{a}$, rather than the filter cavity mode $\hat{b}$. **(c)** The frequency-domain mode structure under single-mode (for both $\hat{a}$ and $\hat{b}$) and resolved sideband (for $\hat{b}$) approximations. The parametric interaction $\hat{b}^\dagger\hat{c}^\dagger$ is realized by the optomechanical coupling under blue-detuned pumping by the mechanical resonant frequency. Note that all mode operators are defined in the rotating frame of frequency ω_0, and Ω is the sideband frequency with respect to it. 150

5.3 The signal-referred shot noise spectral density for different optomechanical interaction strengths G. Different curves are plotted based on Eq. 5.5 with the relation between G and ω_s shown near each line. 152

5.4 The fields involved in the complete analysis of the scheme shown in Fig. 5.1. The right figure is a simplified version of the left one when only looking at the differential mode of the two arms, which is adopted in the analysis in the main text. In the right figure, the sign convention for the mirror reflectivity is also shown. The arm cavity and the SRC length are both tuned to be integer numbers times the wavelength of the carrier at frequency ω_0. 156

5.5 Coupling between the optical modes and the mechanical modes represented in frequency domain. **(a)** Without optical spring compensation, the parametric interaction will be shifted by a small frequency (not to scale) which will greatly ruin the designed filter response. **(b)** With optical spring compensation, the correct sidebands are coupled. It is achieved by making the value of blue-detuning to be slightly larger than the bare mechanical resonant frequency, as described in Eq. 5.17. Insert: the best scenario for resolved-sideband is to set $2\omega_m$ around half FSR_b. It will make ω_m larger than γ_b as much as possible and, at the same time, make the influence from the higher FSR as little as possible. **(c)** The full circulation loop within the signal channel (sidebands around ω_0, represented by the red box in the upper panel) and the idler channel (sidebands around $\omega_0 + 2\omega_m$, represented by the blue box in the lower panel). The idler channel is ignored in the idealised Hamiltonian analysis. 158

5.6 The signal-referred shot-noise spectral densities from the full analysis (red and blue) in comparison with the one obtained from idealised Hamiltonian (red dotted curve). Compensating the frequency shift of the mechanical oscillator due to the optical spring effect has a significant influence on the sensitivity. 164

5.7 The noise spectral density for the idler channel, which can be read out by using a local oscillator at $\omega_0 + 2\omega_{m_0}$ (blue), in comparison with the signal channel (red). At low frequencies, the signal information contained in the idler channel is even more than that contained in the signal channel. 164

5.8 The plots for the determinant of the open-loop transfer matrix, given the nominal parameters. The left plot shows the case without compensating the optical spring, while the right one has the optical spring compensated and the contour does not enclose the origin. The opacity of the curve is intentionally made smaller when the magnitude of the frequency is large, which is to highlight the relevant low frequencies. 167

5.9 The plots for the determinant of the open-loop transfer matrix for $G = \omega_s$ (left) and $G = 1.01\,\omega_s$ (right), both with the optical spring shift compensated. For clarity, we only show a tiny regime around the origin, but we have checked that only the right plot has the origin enclosed when zooming out. 168

5.10 Field interactions at a mirror. 170

5.11 Logical flowchart of the time domain simulation. 172

5.12 In-loop step-response of the position of the undamped mechanical oscillator. **(a)** When the pump field is offset by ω_{m_0} only. **(b)** When the pump field is further offset by compensating the optical spring. The lighter solid fill contains rapid oscillations with the dark bounding lines representing the amplitude envelope. The inserts of (a) and (b) show the zoomed-in views of two sections in the time-series, where the individual oscillations can be seen. 174

5.13 Transfer functions for optical gain with optical-spring compensated. The pump-off ($G = 0$) gain is shown in contrast to the pump-on gain for **(a)** a finite mechanical quality factor of 5000, and **(b)** an undamped oscillator. The numerical, time-domain approach shows very good agreement with the analytical, frequency-domain approach. 176

5.14 Optical gain transfer functions for the signal and idler channels in the undamped, optical-spring-compensated system. 176

5.15 The normalized noise spectrum for the signal and idler channels. The signal channel (red) recovers the same response as the pump off case (magenta) which is limited by the bandwidth of the cavity at frequencies of a few kHz and above, which shows there is no degradation of the signal or the bandwidth of the cavity. The idler channel (blue) has greater sensitivity below 30Hz. The ideal blend (black) has greater sensitivity than both the signal and idler channels across the blend frequency range of approximately 25 Hz (green star) 179

6.1 Panel (a): An example binary neutron star merger waveform from Ref. [197] (in particular the Newman-Penrose scalar ψ_4 from the simulation with the HB EOS). In this case the post-merger remnant collapses to a black hole after $\sim 10\,\mathrm{ms}$. There is an $\ell = 2$, $m = 0$ gravitational wave component associated with the collapse (inset), though at this point it is still very small compared to the $\ell = 2$, $m = 2$ component which dominates throughout. Panel (b): The $\ell = 2$, $m = 0$ component of the waveform from the collapse of unstable, uniformly rotating neutron stars. We show two examples with the "ENG" and "MS1" EOSs and different star masses. In the latter case, the $m = 0$ component is about two-thirds the amplitude found in the left panel. We also find that the pre- and post-peak portions of the collapse waveform contribute roughly equally to the total power, and have similar frequencies. 186

6.2 The square root of the antenna power pattern of a gravitational wave detector for signals from sky location described by θ and ϕ at frequency of the arm cavity FSR. 192

6.3 The left part of the figure shows two gravitational wave detector configurations, the 20 km HF detector and the 34 km FSR detector, which are optimal for detecting the peak mode of the post-merger waveform and the 20 mode of the collapse waveform. The right part of the figure shows the strain sensitivity of the two configurations in comparison with *Advanced LIGO*, Einstein Telescope and Cosmic Explorer. The strain sensitivity of the FSR detector is for sources at $\theta = 45°, \phi = 45°$, where the antenna response is at a maximum. The arm cavity powers of the HF and FSR detectors are 4 MW and 10 MW, respectively. The assumed observed squeezing levels are 10 dB. The other detector parameters are shown in Table 6.1. The amplitudes of the peak mode and 20 mode with parameters listed in Eq. 6.2 and 6.3 are shown as dashed and dot-dashed curves, respectively. 193

6.4 (1) Left: The SNR of the HF type detector for the peak mode of the post-merger waveform, and the 20 and 22 modes of the collapse waveform as a function of arm length. The waveform parameters are given as the "canonical" values in Eq. 6.2 and 6.3 and a delay time of 10 ms. The maximal SNR (2.3) for the peak mode is at around 20 km marked with the star, which is chosen for the arm length of the HF detector. The other detector parameters are listed in Table. 6.1; Right: With the HF detector, the SNR for the peak mode of the post-merge waveform which has a 30% variation in the qausinormal frequency f and a 20% variation in the quality factor. The white star corresponds to the star in the left part. (2) Left: The SNR of the FSR type detector for the 20 and 22 modes of the collapse waveform as a function of arm length. The maximal SNR (0.8) for the 20 mode is at around 34 km marked with the star, which is chosen for the arm length of the FSR detector. The waveform parameters are given as the "canonical" values in Eq. 6.3 and a delay time of 10 ms. The other detector parameters are listed in Table. 6.1; Right: With the FSR detector, the SNR of the 20 mode of the collapse waveform which has a 30% variation in the qausinormal frequency f_{20} and a 20% variation in the quality factor. The white star corresponds to the star in the left part. 194

6.5 The SNR of the HF detector to peak mode of the post-merge waveform and FSR detector to the 20 mode of the collapse waveform as a function of delay time. The solid lines are for waveform parameters given in Eq. 6.2 and 6.3. The shaded areas give boundaries of minimal and maximal SNR taking into account the variation in the qausinormal frequency, the quality factor and a factor of 2 variation of the amplitudes of waveforms. At $\tau = 10\,\mathrm{ms}$, the SNR for peak mode is in the range of 0.5–8.6, for 20 mode is in the range of 0.3–1.9. 195

6.6 The diagram shows the coupled arm cavity and SRC. The SRC can be converted into a signal amplifier around the arm cavity FSR frequency with frequency dependent negative dispersion inside for cancelling the accumulated phase of the sidebands reflected from the arm cavity. The amplification factor is $1/\sqrt{2T_{\mathrm{SRM}}}$. 197

6.7 Noise Budget of the HF detector. The detector strain sensitivity is for optimal orientated sources at $\theta = 0°, \phi = 0°$. The detector parameters are listed in Table. 6.1. 201

6.8 Noise Budget of the FSR detector. The detector strain sensitivity is for optimal orientated sources at $\theta = 45°, \phi = 45°$. The detector parameters are listed in Table. 6.1. 202

6.9 The shape of the three lowest-frequency modes with non-zero beam overlap for an optimised aspect ratio of 1.75. Modes (a) and (c) are approximately analogous to the fundamental mode of a vibrating, circular membrane, with the oscillations of the central anti-node being in-phase and in opposition for the two modes respectively. Mode (b) corresponds to a radially symmetric expansion and contraction along the cylinder's axis of symmetry. The resonant frequencies specified under each mode are for a test mass radius of 10 cm. 206

6.10 The resonant frequencies of the three modes shown in Fig. 6.9 over a range of aspect ratios. Each data point was obtained through the FEA modal analysis performed individually for each new substrate geometry defined by the aspect ratio. Due to the time-consuming nature of this process, the parameter space is rather sparsely sampled. Despite this limitation, we can determine the key result, which is the approximate aspect ratio of 1.75 at which the two lowest resonant frequencies coincide. 207

6.11 Thermal noise in displacement contributed by the three highlighted modes in Fig. 6.9. The red dashed line shows the total thermal noise from the three modes considered and the black dashed line shows the result of Eq. 2.45 applied to this system. The total thermal noise from just the first three modes is already approaching the analytic model in black at low frequencies. The convergence between the two thermal noise approaches is likely much closer than the figure suggests due to the infinite half-plane model used in the creation of the black trace, which has the effect of overestimating the true thermal noise. Around and after the substrate resonances, the modal analysis captures the effect of the resonant peaks on the thermal noise which the direct approach misses completely. 209

7.1 Layout of the experiment and of the digital signal processing. A digitally sourced modulation is applied to the laser, which is then connected to both sensing heads. The raw DC diode signal is digitised and demodulated in the digital domain at multiple harmonics. An algorithm is applied to extract the displacement from the resulting signals. 218

7.2 Panel (a) shows the DC sensitivity of the sensing heads without the modulation-demodulation process (calibration factor is $2.85 \times 10^{15}\,\mathrm{Hz\,m^{-1}}$ in the setup). The DC sensitivity is compared with the laser frequency noise that was independently measured by beating two identical lasers. The two curves are consistent at low frequencies (below 1 Hz). The dashed lines represent the inferred noise levels from the two largest noise sources near the modulation frequency: laser amplitude noise and ADC noise. The former noise is suppressed above 7 kHz by the anti-aliasing filter. The laser frequency noise is below the level of $100\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ and does not limit the spectrum around the modulation frequency. Panel (b) shows the total sensitivity of the sensing head with the heterodyne readout and frequency noise stabilisation (FSS). The FSS-on curve represents the extend of the sensitivity increase we achieve with the frequency stabilisation. The curve is limited by the residual air currents below 100 mHz. The dashed lines represent the noise limits of several prominent, incoherent noise sources and show that our new noise floor is consistent with the laser amplitude noise limit at 100 mHz–10 Hz. The BOSEM sensitivity curve [50] allows for a comparison between the sensitivity of the sensing head and contemporary sensing devices. 222

7.3 Panel (a) shows the residual displacement from non-linear coupling observed when performing a sweep over the laser wavelength using temperature tuning. The circular fit represents a phase extraction algorithm that assumes balanced harmonics. The elliptical fit shows the improved linearity that we obtained through a more advanced algorithm that applies ellipse fitting to the Lissajous figure. Panel (b) shows the phase spectrum of a 1 Hz excitation of the laser wavelength using the PZT. The measured data represents the upper limit for the distribution of signal to higher harmonics due to non-linear coupling. The simulated curve represents our estimate of the spectrum that would correspond to the elliptical fit in panel (a). 225

A.1 Definitions of fields in a Michelson interferometer used in the output power derivation. 238

A.2 Definitions of fields in a Fabry-Pérot interferometer used in the derivation of reflected and transmitted power. 238

List of Tables

3.1 Parameters of the prototype suspension 114

4.1 Cavity parameters of the initial room-temperature design. 128

5.1 A list of parameters and nominal values 153

6.1 Parameters of the detector 199

C.1 Optimisation of Parameters 246

Glossary

Active Isolation A scheme or assembly designed to provide a buffer (isolation) between environmental vibrations and a test mass using feedback and control (active). See Section 2.2.4. 42, 54

Advanced LIGO The pair of second-generation gravitational-wave detectors located in Hanford, Washington and Livingston, Louisiana currently operating and responsible for the first-ever detection of gravitational waves. They are an upgrade of the first-generation Laser Interferometer Gravitational-Wave Observatory (LIGO) instrument. 21, 26, 47, 53, 54, 60–62, 76, 77, 79, 82, 103, 125, 126, 145, 182, 184, 200, 214–216, 221, 227, 229, 233, 235, 236

Advanced Virgo The second-generation gravitational-wave detector located near Pisa, Italy and the third instrument to detect a gravitational wave signature. It is an upgrade of the first-generation Virgo instrument. 21, 47, 55, 76, 79, 182, 184, 214

Euler-Lagrange Equation The system of differential equations that minimise the action of a mechanical system, used in conjunction with the system Lagrangian to obtain the equations of motion. See Eq. 2.13. 38–40, 94, 241

Fabry-Pérot Interferometer An interferometer with a single cavity formed by two partially reflective mirrors based on the round-trip interference of the circulating

electric field. This design is explored in Section 1.2.2. 4, 7–10, 14, 31, 32, 68, 72, 239

GEO600 The gravitational-wave detector located near Hanover, Germany, which has served and continues to serve as an essential test facility for techniques and instrumentation designed to improve the performance of gravitational-wave observatories. 47

KAGRA The Kamioka Gravitational Wave Detector (KAGRA) is the first operational cryogenic gravitational-wave detector and is located near Hida, Japan. 21, 77, 124, 184

Lagrangian A quantity that characterises the state of a physical system. For a mechanical system, it is related to the principal sources of energy (kinetic and potential) and is given by Eq. 2.10. 37, 39, 42–47, 59, 93, 94, 99, 104, 106, 107, 241

Laplace Transform An integral transform given by Eq. 1.1, which is commonly used to analyse the complex frequency response of linear time-invariant systems. It can be related to the Fourier transform by considering only the imaginary component of the complex frequency, $s = i\Omega$. 2, 40, 44, 58, 241

LISA The Laser Interferometer Space Antenna (LISA) is a future space-based gravitational-wave detector designed with longer interferometer arms designed with a detection band that extends significantly below that of terrestrial observatories. 35

Loss Angle The imaginary component of the complex Young's modulus, which quantifies the fractional energy loss of a mechanical system per period of oscillation. 50, 53, 58, 59, 62, 64, 66, 124, 126

Linear Time-Invariant A system whose time-evolution is only governed by a set of linear ordinary differential equations and with no explicit dependence upon time. 2, 3, 14, 26, 36, 37, 43, 45

Michelson Interferometer An interferometer with two arms based on the splitting and recombination of light by a central beamsplitter. This design is explored in Section 1.2.1. 4, 6–8, 21, 71, 78, 237

Passive Isolation A scheme or assembly designed to provide a buffer (isolation) between environmental vibrations and a test mass without the need for a powered control system (passive). See Section 2.2.1. 42, 47, 54

Pound-Drever-Hall A scheme for generating an error signal using radio frequency modulation of a carrier field that can be used to maintain optical cavities as their resonant condition with the laser light. 15, 69, 127, 134, 136

Quantum Radiation Pressure Noise The contribution to the total quantum noise that is derived from the fluctuation in radiation pressure incident upon the mirrors of the optical cavity. See Section 2.1.2. 27–29, 33, 34, 71, 75, 77, 82, 128, 131, 143

Quantum Shot Noise The contribution to the total quantum noise that is derived from the photon counting uncertainty (akin to other forms of shot noise). See Section 2.1.2. 27–29, 31–34, 70, 76

Standard Quantum Limit The boundary that governs the highest level of sensitivity that can be attained in a conventional interferometric measurement. See Section 2.1. 30, 34, 35, 71, 75–80, 82, 84–86, 90, 117, 118, 127–129, 131, 142

State Space A linear algebraic representation of a linear time invariant system with the time evolution of the output quantities and internal state variables determined by the set of $\boldsymbol{ABCD}$ matrices as given in Eq. 2.16 and 2.17. 39, 40, 94, 96, 99, 107, 113

Suspension An assembly, based usually upon a combination of pendulums and springs, that provides isolation of the suspended components from environmental disturbances that couple via the attachment point at the top of the suspension. 47, 48, 50, 51, 53–55, 59–64, 76, 80, 82–84, 86, 88, 90–92, 96–99, 101, 103, 107–110, 112, 117–121, 123, 129, 131, 134, 135, 139, 140, 142, 143, 252

Transfer Function A frequency-domain representation of the input-output response of a linear time-invariant system. The input and output quantities are related via Eq. 1.2. 3, 32, 39–41, 43, 44, 46–48, 50, 51, 55, 56, 85, 87, 94, 99–101, 105, 106, 109, 113, 135, 137, 141

To Emma:

Yours Yesterday

Yours Today

Yours Tomorrow

Chapter 1

Introduction

In 1887, *Albert A. Michelson* and *Edward W. Morley* failed to demonstrate the existence of the luminiferous aether [1]. However, their failure has been eclipsed by their development of interferometry as a tool of scientific discovery. Interferometry has pushed the boundaries of precision measurement, provided the means to probe quantum mechanical phenomena, and even opened a whole new window through which to observe the cosmological history of the universe. I hope that my failures will one day also lead to revolutionary advancements across multiple fields of physics.

Much of contemporary scientific discovery is concerned with the detection of weak signals or measurement of known signals to ever-increasing levels of precision. This is evidenced by the many advances made in the experimental demonstration of quantum mechanics—the standard for what constitutes a small scale. On the opposite end of the scale, astrophysical phenomena, although energetic and prominent at the source, often result in weak signals detected on Earth. From the discovery of cosmic microwave background anisotropy to the detection of gravitational waves (GWs), here one can also see the success of precision measurement.

Firmly embedded within the broad topic of precision measurement is the technique of laser interferometry. Interferometry constitutes the gold standard for displacement

measurement, with an excellent level of precision attainable, as well as versatility in its deployment. Although there is a common thread of GW instrumentation throughout the following chapters, this should not be considered as the limiting scope of this work. This thesis should be primarily understood as a collection of diverse applications of interferometry, spanning from novel enhancements of the sensitivity of existing devices, through to expanding the boundaries of fundamental physics research.

This introductory chapter lays out the basic properties of interferometry that underlie every application discussed in later chapters. It is designed to provide simple examples and define the notation and analytical conventions that I use in the analysis of interferometric devices. I also motivate the suitability of interferometers as displacement sensors broadly and as GW detectors in particular. The expert reader may choose to turn directly to the chapter of interest as each can be understood in isolation.

1.1 Framework

Much of the ensuing investigative framework is focused on the analysis of linear time-invariant (LTI) systems [2] and their performance in the frequency domain. To this end, I adopt the standard approach of using the Laplace transform defined as [3]

$$F(s) = \int_{J_0}^{\infty} f(t) \exp(-st) \mathrm{d}t, \tag{1.1}$$

where s denotes the complex frequency space. The Laplace transform can be linked to the more conventional Fourier transform via the substitution $s = i\Omega$. Unless otherwise specified, An upper-case quantity, such as $F(s)$ or $F(\Omega)$, should be understood as the frequency-domain representation of the time-domain quantity $f(t)$, denoted in lower-case. The Laplace transform is very useful in the analysis of LTI systems as the

output, $y(t)$, of the system can be related to the input, $x(t)$, in frequency space via

$$Y(s) = T(s)X(s), \tag{1.2}$$

where $T(s)$ is the transfer function, a characteristic property of the system. Eq. 1.2 reveals a key property of LTI systems, which is that a signal at any frequency, s or Ω, maps onto the same frequency in the output.

The sensitivity of systems presented herein is obtained via the standard application of the power spectral density (PSD) that can be defined in numerous (but equivalent) ways, such as [4]

$$S_{xx}(\Omega) = \int_{-\infty}^{\infty} R_{xx}(t) \exp(-i\Omega t) \mathrm{d}t, \tag{1.3}$$

where $R_{xx}(t)$ is the *autocorrelation function* of the signal $x(t)$. The PSD receives the units of $[\mathrm{x}]^2/\mathrm{Hz}$. The featured sensitivity plots often contain the amplitude spectral density, which is simply given by the square root of the PSD. The related cross spectral density (CSD) between two quantities, $x(t)$ and $y(t)$, is denoted by $S_{xy}(\Omega)$ or $S_{xy}(f)$. A frequently used standard relation between PSDs is

$$S_y(\Omega) = |k|^2 S_x(\Omega), \tag{1.4}$$

where k is the constant of proportionality in the relation $y(t) = kx(t)$. More generally, this relation can be expressed as

$$S_y(\Omega) = |T(\Omega)|^2 S_x(\Omega), \tag{1.5}$$

where $T(\Omega)$ is defined in Eq. 1.2. In some scenarios, the total PSD of a quantity is provided with a breakdown of each individual contribution. In such a case the PSD is

supplemented with a subscript following a comma that denotes the particular noise, for example, $S_{x,\mathrm{th}}$ that denotes the thermal noise contribution to x.

1.2 Interferometry

Many aspects of the behaviour of light can be explained through its treatment as an electromagnetic wave. Here, the most interesting is the ability of light to interfere, that is to say, to produce patterns through the superposition of multiple waves, which happen to be highly sensitive to microscopic perturbations. Interferometers are devices which allow the transformation of a physical quantity, such as displacement, into a light signal via the precisely controlled manipulation of the light's phase. In this sense, interferometers act as phase-to-intensity converters.

Interferometry is a tool with a broad range of applications due to the diversity of interferometer design. The multiple waves that generate the interference pattern can be obtained by, for example, splitting light through a partially reflective surface or distorting them through diffraction. The interference can arise from the waves travelling along different optical paths or by being confined over several round trips in a closed loop. Due to the breadth of the topic, it is not useful to speak in generalities. Therefore, I will focus directly on the two key interferometer types that will be utilised throughout this work. These are the Michelson interferometer and the Fabry-Pérot interferometer.

1.2.1 Michelson Interferometer

One of the first uses of the Michelson interferometer was in the well-known *Michelson-Morley* experiment of 1887 [1]. The interferometer topology is laid out in Fig. 1.1 and consists of a light source, a central beamsplitter, and two 'arms' formed by the

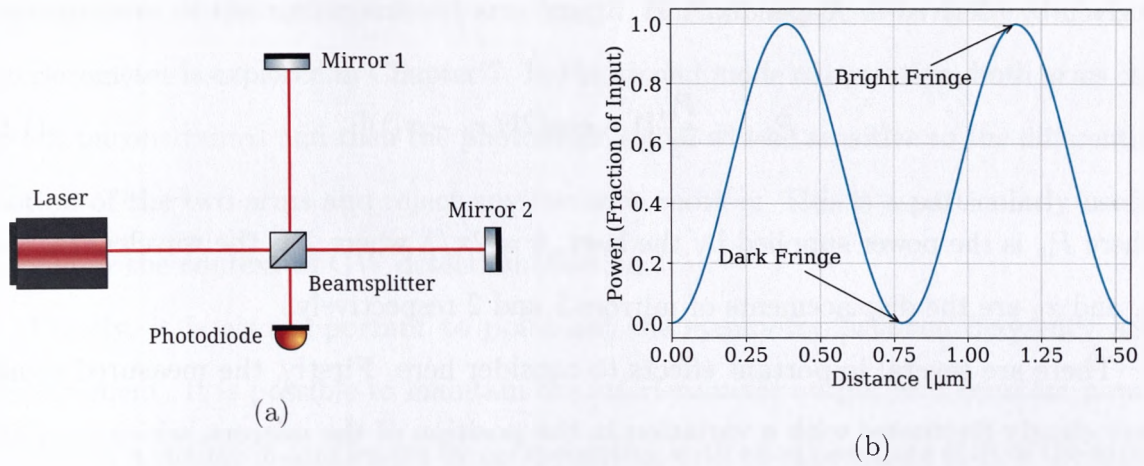


FIGURE 1.1: Panel (a) shows the schematic diagram of a simple Michelson interferometer. Panel (b) shows the response with differential distance ($x_2 - x_1$ in Eq. 1.6). The annotations define the bright and dark fringes, which correspond to the state of constructive and destructive interference, respectively.

beamsplitter and a highly reflective mirror on each end. In the conventional configuration, the light field enters from the left-hand side (LHS). The field is split at the beamsplitter, usually with equal parts in reflection and transmission, and the two field components travel along their corresponding arm, where they reflect and return to the beamsplitter to recombine. The source of interference is the different paths travelled by the two split field components.

The utility of interferometry within the context of this work is in its displacement-measuring capability. Consider the effect of a displacement of the two end mirrors of each interferometer arm. Performing the analysis requires a straightforward application of wave optics and the standard approach to light interaction with an optical component [5]. The perturbation of the light field generates an interference pattern, which is detected by the photodiode in Fig. 1.1a. The resulting optical power incident on the photodiode

is given by (derived in Appendix A.1)

$$P_{\text{out}} = \frac{P_{\text{in}}}{2}[1 - \cos(2k(x_2 - x_1))], \tag{1.6}$$

where P_{in} is the power supplied by the laser, $k = 2\pi/\lambda$ where λ is the wavelength, and x_1 and x_2 are the displacements of mirrors 1 and 2 respectively.

There are several important effects to consider here. Firstly, the measured signal very clearly fluctuates with a variation in the position of the mirrors, which qualifies the Michelson interferometer as a candidate for displacement sensing. Conventionally, a linear response is sought in the displacement sensor with a constant conversion factor, in this case from distance in metres to optical power in Watts. Due to the evidently non-linear response (see Fig. 1.1b),consider only small displacements confined to a narrow region around some average value. For example, considerering only the region around the mid-point between maximum and minimum power (where the gradient of Eq. 1.6 is maximised), a Taylor expansion can be applied to linear order, which gives an optical response of kP_{in} W m^{-1}. For a nominal system consisting of a 10 mW, 1550 nm laser, and an InGaAs photodiode with a dark current of 0.1 nA and responsivity of 1 A W^{-1}, the lowest static displacement that can be hypothetically detected is 2.5 fm. This rather impressive resolution gives an indication of the utility of interferometry in precision metrology. Ultimately, the displacement sensitivity that is achievable in a realistic system is limited by a host of noise sources whose study and mitigation encompass a significant part of this work.

An important property to consider is the fact that the measured signal is a function of the difference between the arm displacements. This allows the Michelson interferometer to be used in two distinct ways. If one arm can be effectively and reliably constrained to a fixed length, the interferometer can be used as a precise displacement sensor for the

measurement of the unconstrained arm length. Such an application of the Michelson interferometer is explored in Chapter 7. In the second mode of operation, both arms can be left unconstrained and then the photodiode signal will be sensitive to the differential motion of the two arms and reject any common motion. This is a particularly useful feature in the context of GW detection (Section 1.3).

Finally, it is also important to point out the symmetry between frequency and displacement. It is possible to maintain the interferometer output at a constant power level upon a change in arm length by compensating with an appropriate shift of the input laser frequency (i.e. keeping the term kx constant and noting that $k = 2\pi f_0/c$, where f_0 is the laser frequency). Consequently, it is impossible to separate power fluctuations that originate due to displacement from those that originate due to fluctuations in the laser frequency. This effect is a key source of uncertainty in all interferometric measurements and several sections throughout the following chapters are devoted to discussing it in detail. Nonetheless, the symmetry between frequency and displacement can be useful, for example, in the case of maintaining cavity lock, as laid out in Section 1.2.5.

1.2.2 Fabry-Perot Interferometer

The second type of ubiquitous device that I consider is the Fabry-Pérot interferometer. Developed by *Charles Fabry* and *Alfred Pérot* in 1899 [6], it is also sometimes called an *etalon*[1]. The topology of this type is very simple, consisting of merely two face-on mirrors as shown in Fig. 1.2a. Such a simple design is often an elementary building block within the design of more complex interferometer systems. It is thus also generically referred to as an optical cavity, due to the confinement of the EM field within the space between the two mirrors. The partial reflection from each mirror results in a

[1]The etalon terminology is usually used when mirrors are fixed and the interferometer is used as a frequency reference

round-trip interference of the field entering the cavity with a phase-delayed component of itself. By considering the superposition of two waves of equal frequency, it is clear that the resulting wave increases in amplitude when the two incoming waves are in-phase. The iteration of this process over subsequent round-trips within the cavity causes the intra-cavity field to build up over time. This constructive interference is reached over a tightly constrained range of operation that forms a 'resonance' in analogy with the standing wave criterion in acoustics.

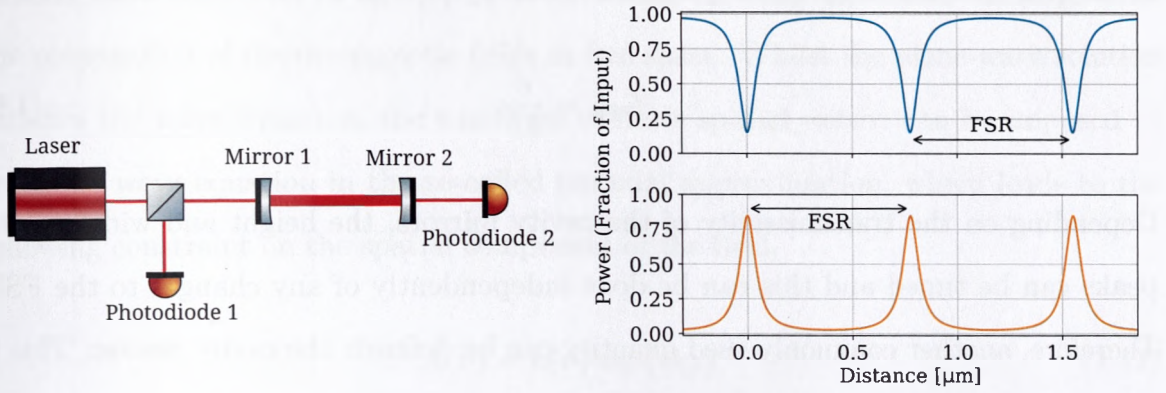
The input field to a Fabry-Pérot interferometer is conventionally generated from a laser placed on one side of the cavity. Due to the partial reflectivity of both cavity mirrors, there are two fields that can be probed, one on each end of the cavity. These signals are defined to be in reflection and transmission based on the position of the photodetector relative to the input field. Continuing with the standard field interaction approach [5], the two signals dependent on the cavity length, L, (derived in Appendix A.2) are given by

$$P_{\text{refl}} = \frac{R_1 + R_2 - 2r_1 r_2 \cos(2kL)}{1 + R_1 R_2 - 2r_1 r_2 \cos(2kL)} P_{\text{in}}, \tag{1.7}$$

$$P_{\text{tran}} = \frac{T_1 T_2}{1 + R_1 R_2 - 2r_1 r_2 \cos(2kL)} P_{\text{in}}, \tag{1.8}$$

where P_{in} is the power supplied by the laser, T and R correspond to the power transmissivity and reflectivity of the mirrors respectively, and subscripts 1 and 2 refer to properties of mirrors 1 and 2. The lower case $r = \sqrt{R}$ refers to amplitude reflectivity.

From this result, some similarities and key differences can be derived between the two interferometers presented thus far, some of which can be well illustrated by considering the comparison between the signals of the two interferometers shown in Fig. 1.1a and 1.2a. In similarity with the Michelson interferometer, the Fabry-Pérot



(a) (b)

FIGURE 1.2: Panel (a) shows the schematic diagram of a simple Fabry-Pérot interferometer. Panel (b) shows the response (power measured at the two photodiodes) of a nominal Fabry-Pérot interferometer. The top plot shows the reflected power and the bottom plot shows the transmitted power. The parameters that give this particular response curve are: $T_1 = 20\%$, $T_2 = 40\%$, and $\lambda = 1550\,\text{nm}$.

interferometer is capable of acting as a displacement sensor. Additionally, as the response is related to a change in the cavity length, the Fabry-Pérot interferometer, strictly speaking, also responds to a differential measurement—in this case, the difference between the displacements of the two cavity mirrors. However, consider the sharpness and narrowness of the Fabry-Pérot interferometer response in Fig. 1.2b, where it is clear that (a) the Fabry-Pérot interferometer makes for a more sensitive device (or at least has the potential) and (b) the Fabry-Pérot interferometer has a substantially smaller range over which it can operate.

For such a configuration there are two key quantities of interest. As is to be expected from the presence of the trigonometric functions in Eq. 1.6–1.8, both interferometer types exhibit a repeating response with displacement and, analogously, with frequency. The frequency span of a single repetition in a Fabry-Pérot interferometer is referred to

as the free spectral range (FSR), which is given by [5]

$$f_{\mathrm{FSR}} = \frac{c}{2L}. \tag{1.9}$$

Depending on the transmissivity of the cavity mirrors, the height and width of the peaks can be tuned and this can be done independently of any changes to the FSR. Therefore, another commonly used quantity can be defined: the cavity *finesse*. This is a measure of the relative width of the resonant peak, given by [5]

$$\mathcal{F} = \frac{f_{\mathrm{FSR}}}{f_{\mathrm{FWHM}}} = \frac{\pi}{2\arcsin\left(\frac{1-r_1r_2}{2\sqrt{r_1r_2}}\right)} \approx \frac{\pi}{1-r_1r_2}, \tag{1.10}$$

where f_{FWHM} is the full-width frequency at half-maximum of the peak power and the approximation holds for reflectivities close to unity.

Fabry-Pérot interferometers see widespread use in areas of precision metrology, particularly due to their versatility (they can be incorporated into more complex optical systems) and their very sharp response to displacement fluctuations. The analysis of multiple systems consisting of Fabry-Pérot interferometers is performed throughout several of the following chapters. The types of noise present in such a configuration are discussed in Chapter 2, whilst practical examples in the context of models, simulations, and experimental work are presented in Chapter 3, 4, and 5.

1.2.3 Gaussian Beams

Although much of the analysis performed thus far can be explained through plane-wave optics alone, there are aspects of a real laser beam beyond the plane-wave treatment that are essential to understand when designing a stable interferometric system. A more

general wave treatment of light requires revisiting the wave equation, which governs the propagation of electromagnetic fields in free space. Whilst the plane-wave solution satisfies the wave equation, the condition of finite spatial extent can be imposed to solve the wave equation in the so-called paraxial approximation, which leads to the following constraint on the spatial component of the field,

$$A(\boldsymbol{r}) = u(\boldsymbol{r}) \exp\{ikz\}, \tag{1.11}$$

where $u(\vec{r})$ is a slowly varying component of the complex amplitude for a wave propagating in the $\boldsymbol{z}$ direction. This approximation leads to the paraxial Helmholtz equation,

$$\nabla_{\perp}^2 u + 2ik\frac{\partial u}{\partial z} = 0, \tag{1.12}$$

where $\nabla_{\perp}^2$ is the transverse component of the usual Laplacian operator. The full solution to the wave equation under such constraints consists of a sum over an infinite series of particular solutions that can receive arbitrary weights to construct any shape that satisfies Eq. 1.12. This decomposition into a sum over different modes can be achieved in different modal bases, one of the most prominent of which is the Hermite-Gauss basis, which yields the following solution to the electric field [7],

$$\begin{aligned} E(x,y,z,t) = \sum_{n,m} E_{n,m} \frac{w_0}{w(z)} H_n\left(\frac{\sqrt{2}x}{w(z)}\right) H_m\left(\frac{\sqrt{2}y}{w(z)}\right) \exp\left(-\frac{x^2+y^2}{w(z)^2}\right) \\ \exp\left(i\left[\omega_0 t - kz - \frac{k(x^2+y^2)}{2R(z)} + \psi_{n,m}(z)\right]\right), \end{aligned} \tag{1.13}$$

where $E_{n,m}$ is a complex scaling factor of the mode's amplitude, w_0 is the beam waist[2], $w(z)$ is the beam radius, H_i are the Hermite polynomials, $R(z)$ is the beam's radius of curvature, and $\psi_{n,m}(z)$ is the Gouy phase, an additional term that contains the phase change over the beam focus.

The fundamental mode ($n = 0$, $m = 0$) has a Gaussian spatial distribution, giving rise to the term Gaussian optics to describe the specific analysis of such spatially distributed beams. As lasers generally produce spatially confined beams that satisfy the paraxial approximation, the results of Gaussian optics apply in the analysis of laser interferometers as well.

1.2.4 Cavity Mode and Stability

An important practical consideration in the operation of any interferometric device is its robustness to perturbations from its ideal operating point. In the naive case of two planar cavity mirrors, the interferometer is highly sensitive to small angular misalignments. Consider the reflection of rays inside a misaligned planar cavity. As no line can be drawn between the two mirrors that meets both mirror surfaces at a right angle, the reflected rays cannot be confined along a single cavity axis each round trip. In such a case, the cavity will fail to resonate even if the longitudinal distance has been correctly set to satisfy the resonant condition.

The solution to this is to refocus the beam using curved cavity mirrors. However, this introduces a key condition on the input field's transverse components, as a curved mirror can alter the mode content of the beam. In simple terms, the resonant condition can be further elaborated to require a round-trip reproduction of the spatial mode

[2]Beam radius is the distance in the transverse direction at which the field amplitude changes from the mid-beam amplitude by a factor of $1/e$. The beam waist is the beam radius at the narrowest point of the beam.

shape as well as the longitudinal phase. This means that a particular, stable cavity has a characteristic transverse field configuration that satisfies the resonant condition, which is called the cavity mode.

Deriving the cavity mode is a standard exercise in Gaussian optics, and leads to the condition that the input beam's mode must match the specific mode that can be supported by the cavity. Of particular importance to this stability condition are the radii of curvature of the cavity mirrors in relation to the cavity length. A stability parameter for each mirror can be defined, given by [7]

$$g_i = 1 - \frac{L}{\mathcal{R}_i}, \tag{1.14}$$

where $\mathcal{R}_i$ is the radius of curvature of mirror i. For the complete cavity, the stability criterion becomes

$$0 \leq g_1 g_2 \leq 1. \tag{1.15}$$

Calculations of cavity stability are a standard yet essential part of interferometer design and will be discussed in more detail in the relevant chapters. There is a set of common, named cavity configurations, such as confocal, $\mathcal{R}_1 = \mathcal{R}_2 = L$; concentric, $\mathcal{R}_1 = \mathcal{R}_2 = L/2$; hemispherical, $\mathcal{R}_1 = L, \mathcal{R}_2 = \infty$. All of these cavities are on the boundary of stability and thus realistic cavities will usually be near one of these canonical configurations with some additional stability margin to protect from small perturbations. There are reasons for choosing each configuration based on specific experimental constraints/requirements, such as beam waist size and location, beam size at the mirrors, cavity response to angular tilts of the mirrors, or one of the mirrors must be planar.

1.2.5 Locking Scheme

As expressed through Eq. 1.7 and 1.8 it is often necessary to consider a very narrow range of operation over which a sufficiently linear response can be obtained of the output of an interferometer to displacement. Small amplitude signals within particular parts of the displacement spectrum are often the most interesting, but if the root mean square (RMS) displacement is too large, the linear approximation breaks down nonetheless and thus the whole endeavour suffers from excess noise and other non-linear effects. In the case of the sharp and narrow region around the resonance of a Fabry-Pérot interferometer, the issue of short displacement range is even more troublesome.

However, it is possible to constrain the operating range to the narrow region around the resonance. This can be achieved through the implementation of an appropriate control scheme that uses the principles of feedback/feedforward to reduce the RMS of the signal produced by the cavity. The process of maintaining a cavity at the optimal operating point is known as 'locking' and the process by which this state is achieved is referred to as 'lock acquisition'. Aside from the specifics of the application, the basic concepts within achieving and maintaining cavity lock, such as controllability and stability, can be readily understood with reference to standard controls theory (e.g. [8]) and thus will not be laid out here.

Within the scheme of a conventional feedback loop, the error signal plays a key role as the quantity that defines a measure of the 'distance' that the system is away from the optimal set-point. As part of the LTI formalism, the error signal should be linear in the measured quantity to ensure ease of analysis and mitigate the additional noise that arises from non-linear effects. The one area of the cavity locking process that is worthwhile to discuss further is the specific technique for obtaining this error signal in the context of an optical cavity. The raw power signal detected at the photodiode

cannot be directly used as an error signal for maintaining the cavity on resonance, as the resonance, by definition, occurs at a stationary point. It thus exhibits a quadratic drop in cavity power for small perturbations. Therefore, a more advanced technique is required that produces a linear signal, which can be fed into the feedback control loop.

Pound-Drever-Hall Locking

One such widely-used technique is the Pound-Drever-Hall (PDH) locking scheme [9]. In simple terms, this scheme involves modulating the input electric field's phase/frequency at some known and controlled frequency, resulting in an input field

$$E_{in} = E_0 \exp(i[\omega_0 t + m\cos(\omega_m t)]), \tag{1.16}$$

where ω_0 is the frequency of the carrier field, m is the modulation index, a dimensionless quantity representing the depth of the modulation, and ω_m is the frequency of the modulation. From Eq. 1.16, it is clear that the intensity of the light does not change as a result of the modulation and requires interference with another field to be observed at a photodetector. Usually, the modulation index is small, as the depth of the perturbation of the carrier field required for locking does not have to be large. In such a case, where $m \ll 1$, the field can be approximated as

$$E_{in} \approx E_0 \exp(i\omega_0 t)[1 + i\frac{m}{2}(\exp(-i\omega_m t) + \exp(i\omega_m t))]. \tag{1.17}$$

Written in this way, it is clear that the field contains three frequency components, one at ω_0 and two at $\omega_0 \pm \omega_m$. Therefore, a weak modulation generates two so-called sidebands, one at either side of the carrier.

The modulation sidebands are often placed in the radio frequency (RF) band, at frequencies above the cavity bandwidth, and specifically chosen to not coincide with any multiple of the cavity FSR. This condition prevents the RF sidebands from being resonantly enhanced by the cavity, otherwise both the carrier and the sidebands would receive a common phase change from the cavity interaction, and the measured power would be invariant to the modulation, as is the case in Eq. 1.16. However, with non-resonant sidebands (situated in the relatively flat region of Fig. 1.2b), the sideband power is only very weakly dependent on the cavity length. Therefore, the RF sidebands effectively act as a local oscillator for the signal field reflected from the cavity at the carrier frequency. This generates a beat signal at the photodiode, which results in a component of the detected power at the modulation frequency, ω_m. Photodetectors exist with a sufficiently broad bandwidth to detect the beat signal in the RF band, meaning that the final stage of demodulation from ω_0, can be performed electronically. This yields a signal that contains a closely linear region around the resonance of the cavity, with a zero crossing when the carrier is on resonance, as shown in Fig. 1.3.

As mentioned in Section 1.2.1, there is a symmetry between displacement and frequency. Strictly speaking, keeping the cavity locked on resonance only requires that the term kL in Eq. 1.7 and 1.8 is kept at an integer multiple of π. This can be achieved through either an actuation on the mirror positions through, for example, a piezoelectric transducer, or through an actuation on the laser frequency. One effect can be used to compensate for the other. For example, a free-swinging cavity whose displacement is not controlled can be kept locked through laser frequency actuation, which is the principle applied in Chapter 4. Similarly, a drifting laser frequency can be compensated using a sufficiently quiet reference interferometer, as is the case for the frequency stabilisation servo implemented in Chapter 7. Conversely, if the laser

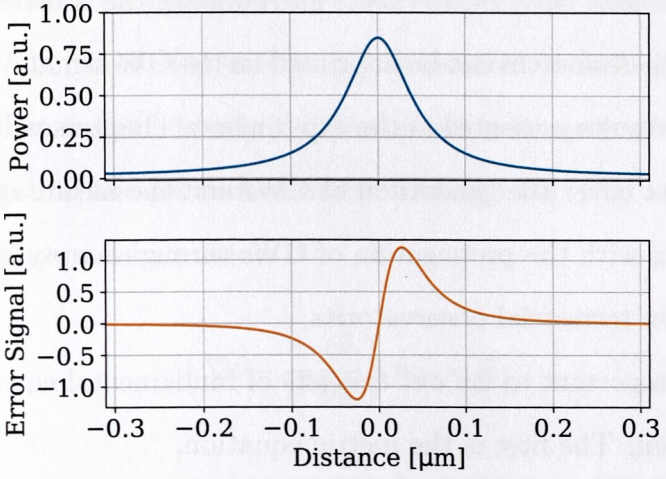


FIGURE 1.3: Close up of a two-mirror-cavity error signal (lower plot) in the region around a single resonance. The transmitted power for the given region is shown in the top plot and corresponds to an enlarged region around the first resonance in Fig. 1.2b.

is the pre-stabilised (by, for example, locking to an atomic absorption line [10]), or intrinsically quiet, then the cavity's length can be stabilised to the laser. Crucially, the locking servo does not distinguish between signal and noise. Therefore, it does not directly boost signal-to-noise ratio (SNR) but rather suppresses the high RMS of the readout signal and thus maintains the interferometer within the linear operating range.

1.3 Gravitational Wave Detection

Thus far the technique of interferometry has been discussed in its application to displacement measurement. However, given the frequent reference within the ensuing chapters to GW astronomy, the knowledge and experience accumulated in this field, and the presentation of direct contributions to this field, it is worthwhile to discuss the basic principles of GW detection.

There is a significant body of literature surrounding the sources of GWs and what information about these sources can be imprinted on the GW signal. As the astrophysical implications of the results presented in the experimental chapters will only be mentioned in passing, I will not cover the generation of GWs and the nature of their sources here. Instead, I will begin with the propagation of GWs through empty spacetime and focus on their detection in terrestrial observatories.

To begin, it is important to lay out two sets of fundamental equations that describe spacetime in vacuum. The first is the metric equation,

$$\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu, \tag{1.18}$$

where $\mathrm{d}s$ is the differential line element, $\mathrm{d}x^\mu$ is the coordinate differential and $g_{\mu\nu}$ is the spacetime metric, which can be reduced to the Minkowski metric $g_{\mu\nu} = \eta_{\mu\nu} = \mathrm{diag}\{-1, 1, 1, 1\}$ in flat spacetime (I have adopted the MTW[3] sign convention in natural units). The second set of equations are the Einstein field equations in vacuum, given by [11]

$$R_{\mu\nu} = 0, \tag{1.19}$$

where $R_{\mu\nu}$ is the Ricci tensor. The Ricci tensor can be connected back to the metric tensor via the definitions

$$R_{\mu\nu} = \frac{\partial \Gamma^\sigma_{\mu\nu}}{\partial x^\sigma} - \frac{\partial \Gamma^\sigma_{\mu\sigma}}{\partial x^\nu} + \Gamma^\sigma_{\mu\nu}\Gamma^\rho_{\sigma\rho} - \Gamma^\sigma_{\mu\rho}\Gamma^\rho_{\nu\sigma}, \tag{1.20}$$

$$\Gamma^\sigma_{\mu\nu} = \frac{1}{2} g^{\sigma\rho}\left(\frac{\partial g_{\rho\mu}}{\partial x^\nu} + \frac{\partial g_{\rho\nu}}{\partial x^\mu} - \frac{\partial g_{\mu\nu}}{\partial x_\rho}\right), \tag{1.21}$$

where $\Gamma^\sigma_{\mu\nu}$ are the Christoffel symbols and $g^{\mu\nu}$ is the inverse metric.

[3] Misner-Thorne-Wheeler

The solution to Eq. 1.19 allows for a perturbation to the Minkowski metric of the form $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, where the metric perturbation is small, $|h_{\mu\nu}| \ll 1$, and the expansion is limited to terms linear in the perturbation (the so-called *linearised theory of gravity*). Propagating the perturbed metric through Eq. 1.20 and 1.21 shows that $h_{\mu\nu}$ follows the familiar wave equation

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right) h_{\mu\nu} = 0 \tag{1.22}$$

The solution to the metric perturbation can be written in the so-called *transverse traceless gauge*, which imposes certain restrictions on the properties of the metric components. Consider the metric components individually, where t represents the time index and x, y and z represent the three spatial directions, the plane wave solution to Eq. 1.22 results in the following constraints [11]

$$h^{TT}_{xx} = -h^{TT}_{yy} = h_+ \cos(\omega t - kz), \tag{1.23}$$

$$h^{TT}_{xy} = h^{TT}_{yx} = h_\times \cos(\omega t - kz), \tag{1.24}$$

$$h_{\mu\nu} = 0 \qquad \text{otherwise.} \tag{1.25}$$

Here the wave is propagating in the z direction and the superscript TT denotes the transverse traceless gauge. The wave amplitude is given by h_+ and $h_\times$, which denotes the two possible wave polarisations.

A notable consequence of this result is that the perturbation to spacetime caused by a propagating GW occurs in both the transverse directions and only in the transverse directions. The passage of a GW between two spatially separated bodies will perturb the distance between them and this change in distance can be measured by the change in the arrival time of a light pulse travelling between the two bodies [12]. In recognising

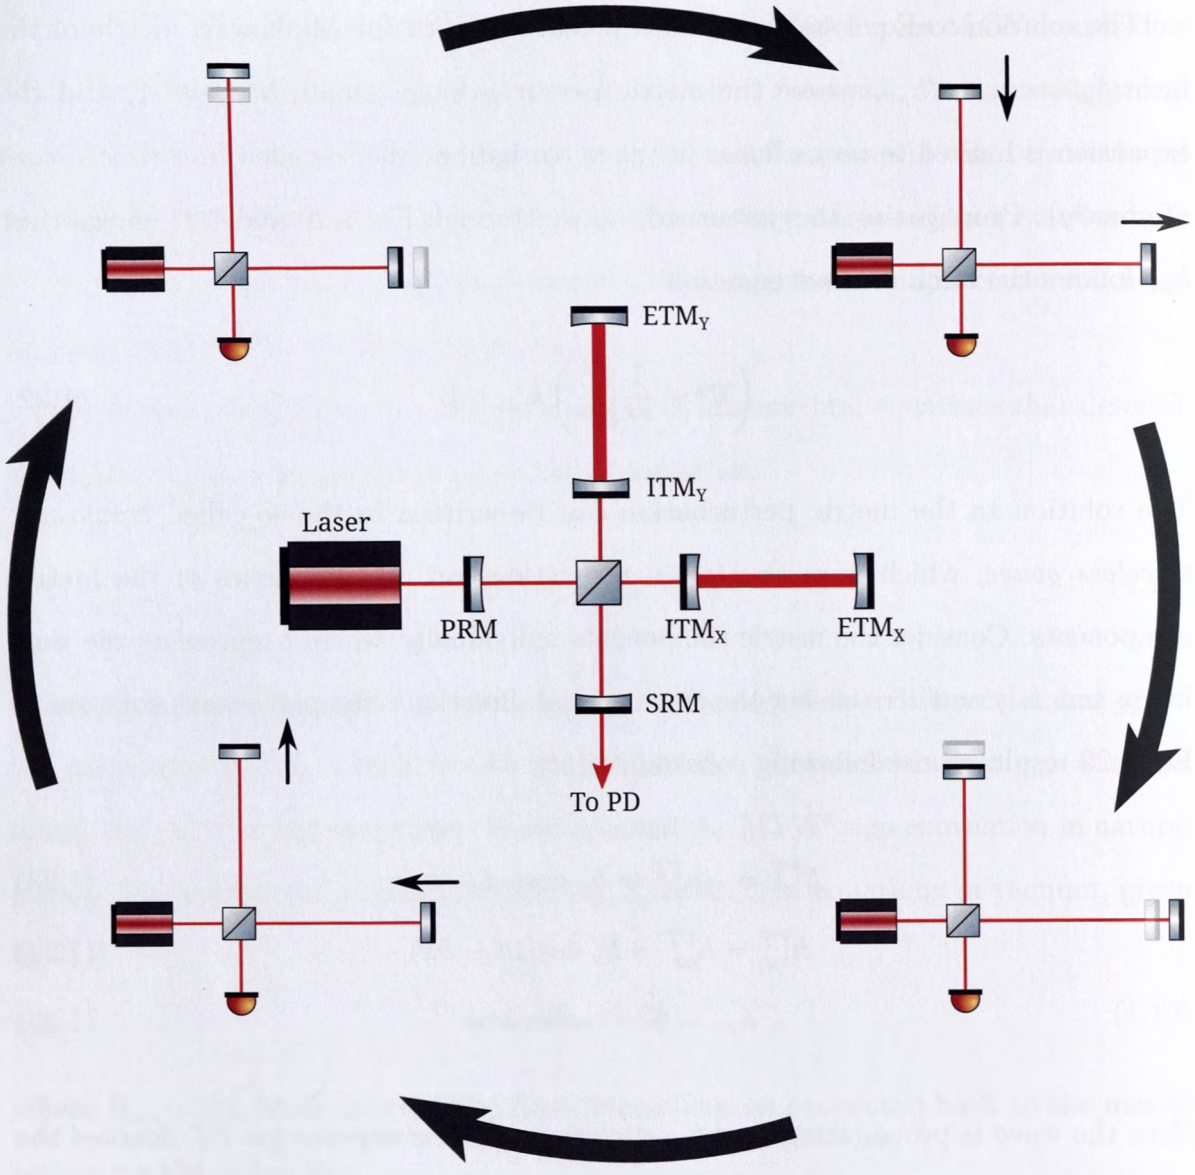


FIGURE 1.4: Centre: Schematic Diagram of a contemporary gravitational-wave detector featuring the key optics—the power recycling mirror (PRM), signal recycling mirror (SRM), the input test mass (ITM), and the end test mass (ETM). Surrounding the central interferometer schematic is a representation of a gravitational wave passing through a simple Michelson interferometer with the cyclical effect on the displacement of the arm optics shown.

this feature of GWs, it becomes clear that the interferometers presented in Section 1.2 (being devices that specifically measure the separation between two optical components

with high precision) are ideal candidates to detect them.

The Michelson interferometer, with its ability to measure displacements in orthogonal directions, is particularly well suited to detect the differential, transverse perturbation of spacetime by GWs. A representation of how a Michelson interferometer responds to a passing GW is shown in Fig. 1.4. Due to the very small amplitude of typical GW signals ($h < 10^{-20}$), it is highly desirable to increase the response of the interferometer even at the expense of a reduction in range. Therefore, the next logical step is to enhance the base design with the addition of resonant optical cavities, which sharpen the interferometer response (Section 1.2.2). These considerations are at the heart of contemporary GW detector design and the schematic shown at the centre of Fig. 1.4 has identifiable similarities with all existing GW observatories, such as *Advanced LIGO* [13], *Advanced Virgo* [14] and *KAGRA* [15]. The success of this approach to GW detection was proven on 14$^{\text{th}}$ September 2015 with the first detection of a GW signal detected from the merger event of two black holes [16].

Due to the exceptionally small amplitude of GW signals on Earth, their detection is a challenging task in precision interferometry. Achieving such sensitivity relies on good knowledge of the limiting noise sources and advanced strategies in their suppression that are at the frontiers of scientific research in this field.

1.4 Thesis Structure

Chapter 2 provides some essential theoretical background for the experimental chapters that follow. In it, I discuss how different types of noise affect the displacement sensitivity of interferometric measurements. Particular focus is placed on the origins of quantum noise, suppression of seismic noise, and the modelling of thermal noise.

Chapters 3 and 4 together present a table-top experiment for the measurement of quantum perturbations over the most sensitive frequency band, centred at 100 Hz. Chapter 3 provides the motivation for the experiment and places it within the context of the contemporary theoretical understanding and current/planned experimental efforts. The chapter further contains a detailed exploration of the seismic isolation techniques, particularly the work on the internal, passive seismic suspension. Chapter 4 focuses on the aspects of the design pertaining to thermal noise and its suppression. This chapter also presents the current status of the experiment with relevant experimental results.

Chapters 5–7 present my contributions to individual, published projects spanning from theoretical modelling and simulations, through to experimental work and laboratory measurements. Chapter 5 presents a novel approach to analysing a quantum amplifier capable of surpassing the conventional shot-noise-limited sensitivity in contemporary GW detectors. This approach utilises a numerical time-domain simulation in support of an expanded frequency-domain analysis.

Chapter 6 contains my contribution to a proposal for a high-frequency detector capable of investigating GW signals from the collapse of post-merger supernovae remnants. It presents the modelling of classical noises at frequencies above the regime usually considered within the band of current detectors.

Chapter 7 presents a compact, multi-fringe interferometric displacement sensor with a level of sensitivity that surpasses displacement sensors currently deployed on the test-mass suspensions of GW detectors. This chapter contains the experimental layout, the sensitivity results of a fixed-mirror measurement and a discussion of the noise sources that limit the sensitivity including suggestions for their further mitigation where appropriate.

Chapter 8 summarises the work presented in each chapter and reiterates the overarching message of this thesis: the application of precision interferometry has broad use in quantum measurement and GW detection and its continued development is essential to our further understanding of fundamental problems in contemporary physics.

Chapter 2

Sensitivity of Interferometric Measurements

Interferometric devices have the potential to measure displacement with excellent levels of precision due to their sharp output response over a very narrow range. However, this potential can only be realised if the level of noise in the readout is sufficiently smaller than the signal. To ensure this is the case, several strategies can be employed. I place these strategies into informal categories as a logical way to think about the different approaches to noise mitigation that are deployed within subsequent chapters:

- noise reduction—directly reducing the level of noise, such as lowering the temperature to inherently reduce thermal fluctuations (Chapter 4 and 6) or locating the interferometer in an area with inherently lower environmental disturbances,
- noise suppression—implementing a technique that lowers the coupling of noise to the signal channel, such as isolating the interferometer from ground vibrations (Sections 2.2 and 3.2–3.6),
- noise cancellation—measuring the noise independently of the signal so that it can be subtracted/suppressed in post-processing or through feedback control

(Section 7.3),

- signal amplification—increasing signal response without a corresponding increase in noise (Chapter 5),

- optimal readout—manipulating the readout technique to make it less sensitive to the noise without directly removing it (Sections 3.1.3 and 5.5).

The sensitivity of a device refers to the ambient level of noise that is measured by it in the absence of signal, which ultimately determines the limit of the smallest measurable signal. As I am primarily concerned with LTI systems (Section 1.1), the sensitivity can be understood well in the form of a PSD (Eq. 1.3). This can be further supplemented by measured or modelled contributions from each known noise source, giving a noise budget. The noise budget for the *Advanced LIGO* design [13], as well as the sensitivity measurement from the third GW observing run (O3) [17], provide good examples of modelling and independently measuring the impact of multiple noise sources. This, in turn, allows for the prediction and verification of which noise sources are the dominant ones at particular frequencies and to consider the optimal strategies for mitigation.

Generally, the sources of noise in an interferometric displacement measurement fall into two distinct types. These can be described as (a) those that induce an unwanted displacement in the test mass[1], and (b) those that affect the field inside the interferometer in a way that is indistinguishable from displacement[2]. I consider three particular noise sources in detail—quantum noise, seismic noise, and thermal noise. These noises, in a sense, comprise the fundamental limitations of the interferometer sensitivity, as they are dictated by the location, material properties or a fundamental

[1] e.g. seismic noise, suspension and substrate thermal noises

[2] e.g. laser frequency noise, quantum shot noise

aspect of nature, and are not easily overcome. I conclude the chapter with a brief discussion of the key noises generated by the laser and how they can be suppressed.

2.1 Quantum Noise

The quantum nature of light imposes a limit upon the precision of any measurement of its properties [18]. This is widely understood as the application of the Heisenberg uncertainty limit to the field inside the interferometer [19,20]. Although often treated as a single phenomenon, the quantum noise can be separated into two conceptually distinct sources, that is the quantum shot noise (QSN) and the quantum radiation pressure noise (QRPN). The two contributions generate noise in the displacement readout via an independent effect on the field in the case of the former, and by inducing an unwanted displacement of the test mass in the case of the latter.

2.1.1 Approximate Derivation

To motivate the existence of the two quantum noise components and intuitively understand their behaviour, one can consider the measurement of laser light as a photon counting experiment. This is a process followed by Ref. [21] and provides a good order-of-magnitude approach. In this section I will outline this approach in much the same way, as it is a central concept in the experiment presented in Chapter 3 and 4 and thus it is worthwhile to reproduce it here.

Shot Noise

The QSN is fundamentally similar to other familiar forms of shot noise, such as electron shot noise in circuits [22]. To understand this effect, consider light reflecting

off a mirror, measured by an idealised 'phase detector'. Given an average arrival of N_γ photons over a time interval τ, the detection can be treated as following a Poisson distribution with a fluctuation of $\sqrt{N_\gamma}$. The phase is related to the average arrival time of the photons, so it can be argued that the uncertainty on the phase, $\delta\phi$, is given by the fractional uncertainty on the number of photons that arrive in the time, τ, that is to say, $\delta\phi = 1/\sqrt{N_\gamma}$. This can be relate to the phase noise PSD as $\delta\phi^2 \approx S_\phi/\tau$, where the assumption is made that the QSN has a white spectral density. As the phase is directly related to the displacement by $\phi = kx$, and $N_\gamma = P_0\tau/(\hbar\omega_0)$, this results in an approximate displacement PSD for the QSN

$$S'_{x,\mathrm{sh}} = \frac{\hbar c^2}{P_0\omega_0}, \tag{2.1}$$

where the primed quantity is used in this section to distinguish the approximate derivation from the later non-primed, accurate derivation. This simple treatment provides some important insights into the properties of the QSN. The inverse scaling with the macroscopic, average power in Eq. 2.1 correctly captures the fact that whilst the signal amplitude grows linearly with power, the noise amplitude grows only as the square root of the power, leading to the ultimately favourable scaling of the SNR. Additionally, Eq. 2.1 does not feature any physical or mechanical properties of the mirror. This reflects the fact that the QSN is purely a consequence of the quantum behaviour of the light and not dependent on the light-mirror interaction.

Radiation Pressure Noise

The second contribution to the total quantum noise, the QRPN, results from the radiation pressure generated by the field on the test mass [23]. Returning to the single-mirror photon picture once more, consider the irregular arrival of photons as

generating a fluctuating force on the test mass. This is given by $\delta F = \sqrt{N_\gamma}\hbar k/\tau$ in reference to the fluctuation in the total momentum imparted by the photons over the interval, τ. Applying the assumption that the test mass is free, that is to say, unconstrained by any rigid connections, one can use Newton's second law of motion (in its Fourier transformed form $F(\Omega) = -m\Omega^2 X(\Omega)$) to get the approximate contribution from QRPN

$$S'_{x,\mathrm{rp}} = \frac{P_0\hbar\omega_0}{m^2c^2\Omega^4} \tag{2.2}$$

This approximate derivation also contains useful insights. Particularly, this equation shows a proportional scaling with the laser power in contrast to the inverse scaling of the QSN in Eq. 2.1. In this case, an increase in the power does not improve the signal response but only serves to drive the mirror dynamics with greater amplitude, thus the increased power has a detrimental effect. Secondly, a property of the mirror, the mass, now features explicitly in the equation, which reflects the fact that the QRPN couples into the measurement explicitly via the light-mirror interaction.

The QRPN can be thought of as a necessary consequence of performing a measurement of the test mass position, which is enforced by the Heisenberg uncertainty principle [20]. The QSN and QRPN can be readily analogised to the quantum noise limitation of the Heisenberg microscope [24]. The former reflects the uncertainty on the position intrinsic to the sensitivity of the probe (EM field) and the latter reflects the disturbance of the system caused by the probe interacting with the system. For this reason, the QRPN is often referred to as 'back-action' noise.

Standard Quantum Limit

Combining the contributions in Eq. 2.1 and 2.2 gives the total, approximate quantum noise PSD

$$S'_x(\Omega) = \frac{\hbar c^2}{P_0\omega_0} + \frac{P_0\hbar\omega_0}{m^2c^2\Omega^4}. \tag{2.3}$$

For the measurement of a particular test mass, this equation can be minimised with respect to $P_0\omega_0$ to obtain

$$S'_{\mathrm{SQL}}(\Omega) = \frac{2\hbar}{m\Omega^2}. \tag{2.4}$$

This curve represents the limit of sensitivity achievable at a particular frequency in a conventional, interferometric measurement[3]. This so-called standard quantum limit (SQL) was formulated by *V. B. Braginsky* in 1968 and subsequently discussed in the framework of a minimal force measurement [19, 25]. The SQL evidently cannot be reached at all frequencies simultaneously. Rather, it is a bounding surface, which provides the sensitivity limit that cannot be surpassed by any conventional configuration and is reached by all such configurations at only one characteristic frequency, denoted by ω_{SQL}. This special point reflects the balance between the two contributions and reveals a very interesting property of interferometric measurement; the SQL is only a function of a single free parameter—the mass of the test mass—and is completely independent of the interferometer topology and optical parameters. This is despite the fact that the position-momentum uncertainty of the test mass itself does not directly enforce the SQL [26].

[3]That is to say, a measurement that does not make use of specific quantum-noise surpassing techniques, as discussed in Section 3.1

2.1.2 Quantum Mechanical Derivation

There are two fundamental issues with the simplified derivation in Section 2.1.1. Firstly, a conventional photo-detector only gives access to the intensity of the EM field, so the assumption that one can simply extract the light phase directly is rather unrealistic. The phase detection is made possible by the presence of an interferometer, whose response is not accounted for in the prior analysis. Secondly, the QSN is fundamentally a quantum phenomenon and thus treating it as a classical, random process is not accurate. Therefore, the precise treatment must consider the system in terms of quantum mechanical operators [27, 28]. Performing these derivations requires substantial familiarity with quantum optical conventions and analysis, which is unnecessary for the scope of this work.

Instead, it is possible to use fact that the quantum mechanical operators representing the electric field propagate through the interferometer in the same way as classical fields. This approach is clearly and thoroughly laid out in Ref. [29]. Consider the fields in a Fabry-Pérot interferometer, as shown in Fig. 2.1, which model the propagation of signal and noise to the readout in a realistic optical system. There are two key, interrelated properties of the quantum mechanical approach that will be asserted without proof[4]. Firstly, the complex electric field can be represented in terms of two quadrature operators, given by $\hat{a}_1 = (\hat{a}(\Omega) + \hat{a}^\dagger(-\Omega))/\sqrt{2}$ and $\hat{a}_2 = (\hat{a}(\Omega) - \hat{a}^\dagger(-\Omega))/\sqrt{2}$ (Eq. 6 in Ref. [29]), which can, in loose terms, be understood as the amplitude and phase perturbations of the field respectively. Secondly, losses, such as absorption of light by a photodiode, introduce an input, so-called vacuum field ($\hat{a}$ in Fig. 2.1) with corresponding PSD and CSD given by $S_{a_n,a_m} = \delta_{n,m}$ for $a_n, a_m \in \{a_1, a_2\}$ (Eq. 26 in Ref. [29]).

[4]These are standard results used within the field. Definitions and examples of their use can be found in e.g. Ref. [21, 27–29]

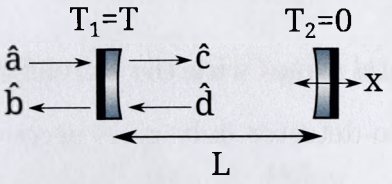


FIGURE 2.1: Schematic diagram of the relevant fields and parameters in the quantum mechanical derivation of quantum noise in a Fabry-Pérot interferometer.

Shot Noise

This formalism can be used to obtain the separate transfer functions of the signal and noise to the readout and, from these, it is possible to derive the quantum noise coupling to the displacement measurement. The full derivation for the specific case of a tuned Fabry-Pérot interferometer is fully covered in Ref. [30]. A conceptually similar approach is employed in the numerical simulation of quantum noise in Chapter 5. The resulting QSN level is given by

$$S_{x,\mathrm{sh}} = \frac{\hbar L^2(\Omega^2+\gamma^2)}{P_c T \omega_0}, \tag{2.5}$$

where $\gamma = Tc/(4L)$ is the cavity bandwidth. This result shows the effect of the cavity in two ways that differ from the result in Eq. 2.1. Most notably, there is now a frequency dependence to the QSN, which becomes significant above the cavity bandwidth. This was not an effect accounted for in the approximate derivation because of the absence of the cavity. Secondly, the low-frequency sensitivity, given by $S_{x,\mathrm{sh}}\big|_{\Omega\ll\gamma} \approx \hbar T c^2/(16 P_c \omega_0)$ has a slightly different form, once again due to the presence of the cavity. However, some key similarities remain. Notably, the spectral density is flat with the same inverse scaling with power and the absence of any mechanical properties of the cavity mirrors.

Of note is the behaviour of the QSN with mirror transmissivity in Eq. 2.5. The

cavity bandwidth and the noise floor below the cavity bandwidth are both proportional to the transmissivity. Generally, it is preferable to minimise the noise and maximise the cavity bandwidth to extend the noise floor over the broadest frequency band possible. However, due to the proportionality of both terms with transmissivity, both goals cannot be achieved simultaneously. If one takes the peak optical gain of the cavity to be in a sense proportional to the inverse of the QSN floor, it can be seen that the gain-bandwidth product is constant. This result, referred to as the Mizuno limit [31], poses a very formidable obstacle towards the design of broadband, low-noise interferometric measurement.

Finally, it should be noted that the rising noise above the cavity bandwidth is a consequence of the degradation in the signal response of the optical cavity rather than due to an intrinsic increase in the noise. The QSN injected into the cavity itself maintains a white spectral density as the namesake suggests. This distinction, whilst seemingly abstract here, will become crucial in the understanding of the quantum amplifier analysed in Chapter 5. The key concept to understand is that the noise in displacement can be reduced through two completely independent means: reducing the noise through a reduction/shaping of the input vacuum fluctuations, or increasing the signal response by amplifying the optical gain to displacement.

Radiation Pressure Noise

The quantum mechanical treatment of the radiation pressure effect involves the rigorous definition of the mirror's equation of motion [21,29]. The interesting distinction between the two quantum noise sources is that the QSN arises from the direct coupling of the input phase quadrature to the output phase quadrature, that is to say, $\hat{a}_2 \rightarrow \hat{b}_2$, whereas the QRPN cross-couples the input amplitude quadrature to the output phase

quadrature, that is to say, $\hat{a}_1 \rightarrow \hat{b}_2$. The accurately derived result (following the detailed approach in Ref. [30]) is given by

$$S_{x,\mathrm{rp}} = \frac{16 P_c T \hbar \omega_0}{m^2 L^2 \Omega^4 (\Omega^2 + \gamma^2)}, \tag{2.6}$$

which shows key differences due to the cavity response and key similarities in the low-frequency scaling compared to Eq. 2.2 in complete analogy with the result of the two QSN derivations.

Standard Quantum Limit

In the complete treatment, the total quantum noise, given by the combination of Eq. 2.5 and 2.6, can be written more concisely as [21, 29]

$$S_x(\Omega) = \left(\frac{1}{\kappa} + \kappa\right) \frac{S_{\mathrm{SQL}}(\Omega)}{2}, \tag{2.7}$$

adopting the notation of [21]. The parameter κ is given by

$$\kappa = \frac{\omega_0 T_1 P_c}{m L^2 \Omega^2 \left(\Omega^2 + \left(\frac{T_1 c}{4L}\right)^2\right)}, \tag{2.8}$$

and

$$S_{\mathrm{SQL}}(\Omega) = \frac{8\hbar}{m\Omega^2}. \tag{2.9}$$

The comparison between the QSN and the QRPN is shown in Fig. 2.2 where the SQL surface is also indicated. This figure concisely lays out the behaviour of the two noise sources discussed in their respective sections.

When written in the form of Eq. 2.7, it is still possible to identify the two quantum noise sources, where $S_{\mathrm{SQL}}/2\kappa$ is the QSN and $\kappa S_{\mathrm{SQL}}/2$ is the QRPN. In this form,

however, it is clear that the quantum noise is minimised at $\kappa = 1$. This means that the terms collected into S_{SQL} in Eq. 2.9 can be identified with the SQL concept derived in Eq. 2.4, except with a corrected numerical prefactor. Further, due to the explicit dependence of κ on frequency, Eq. 2.7 also confirms that the SQL can only be reached at only one frequency in a given detector configuration (this is evident in Fig. 2.2 also).

Thus, the insight obtained from the approximate derivation holds true even in the precise treatment. This definition of the SQL applies to a longitudinal position measurement but a similar result can be derived for other degrees of freedom, such as angular motion [32]. The implications of the SQL on quantum measurement, including efforts to surpass it are further discussed in Chapter 3.

2.2 Seismic Noise

Seismic noise is a very prominent noise source in many interferometric applications. It arises from the ambient ground vibrations and intermittent seismic events, which directly result in displacement of the test masses in terrestrial interferometric devices. In fact, this noise source would naturally be the most prominent over the whole detection band of contemporary GW detectors by several orders of magnitude if not for the extensive use of suppression techniques.

2.2.1 Passive Isolation

The most comprehensive way to reduce the effect of the ground's vibrations is to avoid placing the interferometer in direct contact with the ground. Taken to its logical conclusion, this leads to the consideration of not having it on Earth at all, which is the approach of space-based detectors such as *LISA* [33]. Given such a radical step is often

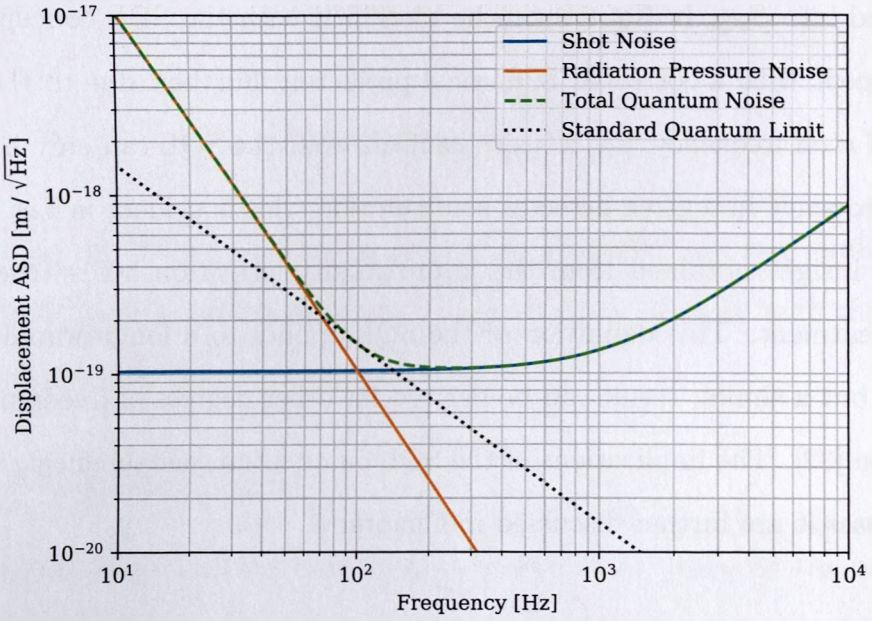


FIGURE 2.2: Quantum noise in an example system separated into the two constituent contributions of quantum shot noise and quantum radiation pressure noise. The particular system being modelled consists of a tuned Fabry-Pérot interferometer of length 10 cm with 100 g mirrors and an input coupler transmissivity of 10 ppm. The input laser power has been tuned to place the frequency at which the total noise intersects with the standard quantum limit at 100 Hz.

impractical and too expensive, one must instead consider the options for isolating the interferometer from the seismic noise on Earth.

Fundamentally, the detector must be in contact with the ground somehow. The purpose of seismic isolation is to provide a buffer between this contact point and the test masses[5], which attenuates the ground motion, and is often mechanical in nature. Two such commonly deployed components are the pendulum and the spring. Both of these components fit the category of LTI systems for motion that is sufficiently small

[5]the masses whose position is ultimately measured (often the mirrors of an interferometer)

in amplitude[6], which is particularly useful for the analysis common frequency-domain methods and standard tools of control theory can be utilised. The techniques and results described in this section will be extensively applied within the experimental design outlined in Chapter 3.

The framework I adopt in the analysis of seismic isolation begins with constructing the system Lagrangian,

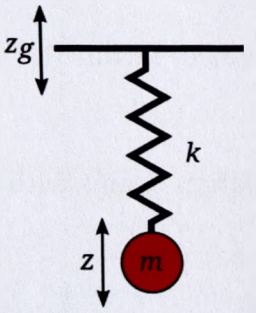
$$\mathcal{L}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = T(\boldsymbol{q}, \dot{\boldsymbol{q}}) - V(\boldsymbol{q}, \dot{\boldsymbol{q}}), \tag{2.10}$$

where T and V are the system's kinetic and potential energy respectively and both are functions of the generalised coordinates, $\boldsymbol{q}$, and their respective velocities.

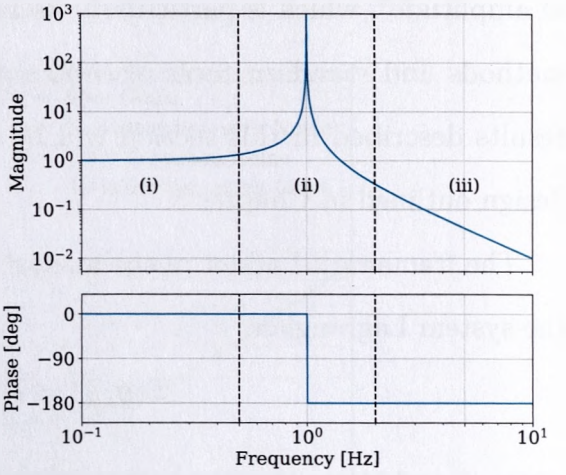
Spring-Based Isolation

First, consider the isolation capabilities of a spring as this is the simplest system to analyse. Springs can be used in vertical or horizontal configurations. In GW detectors, however, they are primarily used for the isolation of vertical motion [34]. In this light, consider a vertically oriented spring with a spring constant k, supporting a mass m as shown in Fig. 2.3a. The mass motion, z, is defined in an inertial frame rather than in the lab frame, which is itself experiencing vertical (non-inertial) motion, z_g, due to the ground vibration. A very useful property of the LTI treatment is that if two input drives do not couple to the same degrees of freedom, each driving term can be considered in terms of two independent Lagrangians. As vertical ground motion induces

[6]The definition of small in this case means motion that is confined to regions of linear response (for example within the elastic region of deformation of mechanical components) or, at least, regions where a linear expansion of the dynamics is accurate enough for the purposes of the experiment.



(a)



(b)

FIGURE 2.3: Panel (a) shows the schematic diagram of a simple spring-based isolation system. Panel (b) shows the Bode plot of the transfer function for the vertical displacement at the attachment point to the displacement of the mass at the bottom of the spring. The three areas indicated correspond to regions where the motion is (i) directly transferred from suspension point to mass, (ii) amplified by the resonant oscillations of the mass-spring system, and (iii) suppressed by the naturally diminishing response of the system above its resonance.

vertical motion in the spring[7], the kinetic energy reduces to

$$T = \frac{1}{2}m\dot{z}^2. \tag{2.11}$$

The potential energy can be written down with reference to Hooke's law as

$$V = \frac{1}{2}k(z - z_g)^2, \tag{2.12}$$

where it should be noted that the spring extension can be given by the difference, $(z - z_g)$, regardless of any static offsets between the two quantities (see proof in Appendix B.1). The equation of motion can subsequently be derived from the Euler-Lagrange (EL)

[7]Vertical ground motion can induce coupling into non-vertical degrees of freedom but this effect can usually be made small and treated separately (see later subsection for discussion of cross-couplings).

equations

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial \mathcal{L}}{\partial \dot{q}_i} - \frac{\partial \mathcal{L}}{\partial q_i} = 0. \tag{2.13}$$

For the simple case of a single spring, the system state can be described with one generalised coordinate, which results in only a single EL equation for the system. In a more complex system, the number of generalised coordinates can be naturally expanded to describe the system, which in turn will result in a coupled set of simultaneous equations resulting from the EL equation applied to each coordinate.

Returning to the simple case, the spring equation of motion (EOM) is given by

$$m\ddot{z} + k(z - z_g) = 0 \tag{2.14}$$

from the Lagrangian constructed from Eq. 2.11 and 2.12. Note that Eq. 2.14 can be rewritten in the form

$$m\ddot{z} + kz = F, \tag{2.15}$$

where the LHS represents the drive-free dynamics of the system and $F = kz_g$ is the external force applied by the moving ground. The form of Eq. 2.15 is suggestive of the general way to couple an arbitrary, external force into the EOM via the right-hand side (RHS) of the EL equation. Although I will primarily couple the driving force generated by ground motion via the definition of the coordinate system that led to Eq. 2.14, I will also use the latter approach in situations where this treatment is more appropriate, for example in Section 2.3.

At this stage, the analysis can continue in one of two ways. One can (i) construct a state space model or (ii) produce the transfer function by transforming into the frequency domain. In the case of (i), I adopt the standard state space representation of

a system, given by the equations

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{A}\boldsymbol{x}(t) + \boldsymbol{B}\boldsymbol{u}(t), \tag{2.16}$$

$$\boldsymbol{y}(t) = \boldsymbol{C}\boldsymbol{x}(t) + \boldsymbol{D}\boldsymbol{u}(t), \tag{2.17}$$

where $\boldsymbol{x}(t)$ is a vector containing the internal states, $\boldsymbol{u}(t)$ is a vector of all the external inputs, and $\boldsymbol{y}(t)$ is a vector of the system observables (outputs). $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$, and $\boldsymbol{D}$ are the matrices governing (a) the internal dynamics, (b) the input driving the system, (c) the linkage between observables and the internal states, and (d) the direct coupling of inputs to observables, respectively. There are many benefits to using this approach, particularly in the case of cumbersome systems with multiple degrees of freedom, such as that presented in Chapter 3. Due to the direct analogy between Eq. 2.16 and the system's EOMs, it is often much easier to obtain the state space representation from the EL equations than it is to solve a complicated system of simultaneous equations to obtain the transfer function. However, a drawback is that it may not be as clear to get an intuitive understanding of the system's behaviour by simply glancing at the state space representation. In such a case, it is often more useful to turn to the transfer function instead.

The transfer function can be obtained either from the state space model or from the original EOM. In the former case, this is particularly straightforward to achieve computationally, even for a highly complex system, if an analytic representation is not required. However, in the single spring case, it is simple enough to start with Eq. 2.14 and solve the system analytically. Transforming to the frequency domain via the Laplace transform yields

$$ms^2 Z + k(Z - Z_g) = 0, \tag{2.18}$$

which can be rearranged in the transfer function form

$$T(s) = \frac{Z(s)}{Z_g(s)} = \frac{k}{ms^2 + k}. \tag{2.19}$$

This formalism leads to some insight into the system's behaviour and why a spring acts as seismic isolation. Consider the steady state frequency response of the system by substituting $s = i\Omega$ into the equation. Fig. 2.3b shows the transfer function in frequency space for a nominal choice of parameters. The three regions labelled in Fig. 2.3b highlight the three features of interest for an isolation system. Region (i) corresponds to the low-frequency response of the isolation, which is often a region of flat response close to unity in magnitude. In this region, the spring transfers the ground motion directly to the mass with negligible attenuation or amplification. Feature (ii) corresponds to the region at and around the resonance of the system, which occurs at the root of the denominator of Eq. 2.19. Noting that the standard relation for the resonant frequency, ω_0, of a spring is given by ${\omega_0}^2 = k/m$, Eq. 2.19 can be rewritten in the form

$$T(\Omega) = \frac{{\omega_0}^2}{{\omega_0}^2 - \Omega^2}. \tag{2.20}$$

From this result, it is clear that the root of the system occurs at the resonant frequency. In region (ii), the motion is substantially amplified by the spring system. In the final region, (iii), the transfer function falls off as $\sim \Omega^{-2}$, corresponding to a reduction in the coupling between the motion of the ground and the resultant motion of the mass. It is in this sense that the spring acts as seismic isolation—at frequencies above the system resonance, a test mass connected to the spring experiences a residual level of seismic noise that is attenuated from the pure ground motion. This type of isolation entirely relies on the natural, mechanical response of the spring to external excitation. As such,

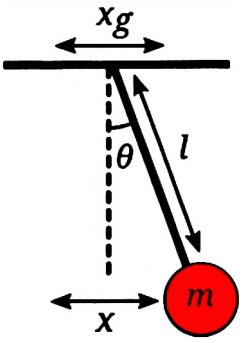


FIGURE 2.4: Schematic diagram of a simple pendulum-based isolation system.

it falls under the category of passive isolation, which is contrasted in Section 2.2.4 with its counterpart, active isolation.

Pendulum-Based Isolation

A similar treatment can be applied to the case of a pendulum. Consider a mass m suspended from a pendulum of length l as in Fig. 2.4. The pendulum is suitable for horizontal isolation and is ubiquitously used among the high-performance isolation stages of all contemporary GW detectors. As such, consider a horizontal ground vibration, x_g that displaces the suspension point of the pendulum. Unlike the spring example, the pendulum's motion lies within a plane formed by the vertical axis and one horizontal direction (motion in the perpendicular direction can be treated independently). Although the Lagrangian must now be constructed from both the x and z components of the pendulum motion, they can both be parametrised in terms of the angle, θ, to the vertical axis. This is due to the motion being confined to an arc, constrained by the length of the pendulum wire. Consequently, the system is fully described by the new

generalised coordinate, θ. The parametrisation is as follows:

$$x = x_g + l\sin(\theta), \tag{2.21}$$

$$z = z_0 - l\cos(\theta) \tag{2.22}$$

where it should be noted that the ground couples directly into the definition of the x coordinate. This approach gives θ a physical meaning—it represents the angle from vertical through which the tension in the pendulum wire acts on the pendulum mass. The ground motion can be instead coupled into the EOM via an external torque acting on the mass, but this approach makes the meaning of θ more opaque and the analysis less intuitive. Ultimately, both approaches yield the same transfer function from ground motion to the x displacement. The system Lagrangian thus becomes

$$\mathcal{L} = \frac{1}{2}ml^2\dot{\theta}^2 + ml\cos(\theta)\dot{\theta}\dot{x}_g + mgl\cos(\theta), \tag{2.23}$$

where terms not explicitly containing θ or its derivative have been omitted.

This equation is no longer explicitly linear in the state variable, θ, which results in the breakdown of the LTI analysis. However, it has been shown that the motion of the ground is very small[8], which means that a Taylor expand may be applied to linear order in the EOM (second order in the Lagrangian) in the unit-less quantities θ and x_g/l, and their derivatives. Applying this approximation yields the, now linear, EOM

$$l\ddot{\theta} + g\theta + \ddot{x}_g = 0 \tag{2.24}$$

[8]The typical ground motion has an RMS of $\sim$1 µm [35]. For a metre-scale pendulum, this results in an RMS angular displacement of $\sim$1 µrad, which satisfies the small angle approximation.

and this leads to the transfer function

$$T_\theta(s) = \frac{\Theta(s)}{X_g(s)} = \frac{s^2}{ls^2 + g}. \tag{2.25}$$

This can be used to obtain the transfer function to the Laplace transformed longitudinal position of the pendulum mass using the linear order expansion of Eq. 2.21, resulting in

$$T_x(s) = \frac{X(s)}{X_g(s)} = 1 + l\ T_\theta(s) = \frac{g}{ls^2 + g}. \tag{2.26}$$

By introducing the standard result for the resonance of a simple pendulum, ${\omega_0}^2 = g/l$, it can be shown that the RHS of Eq. 2.26 conforms to the result in Eq. 2.20. This confirms the use of a pendulum for seismic isolation in analogy with the spring system.

Thus far, the treatment of the pendulum has considered only the gravitational restoring force, which is valid in the case where the suspending wire has no stiffness. In the more general case, it is important to account for the anelasticity of the wire [36]. For small deformations of the wire, it is sage to assume the restoring torque is linear with the angular displacement, θ. This gives rise to the anelastic potential energy term in Eq. 2.23 of the form

$$V_{\text{an}} = \frac{1}{2} k_{\text{el}} \theta^2, \tag{2.27}$$

which is analogous to Hooke's law for a linear spring except that here k_{el} is the *angular* spring constant. Using anelastic theory, it can be shown that the spring constant is given by [37]

$$k_{\text{el}} = \frac{1}{2}\sqrt{mgYI_a}, \tag{2.28}$$

where Y is the Young's modulus and I_a is the second moment of area.

The Lagrangian can then be solved, augmented with the term in Eq. 2.27. This

time, the EOM is derived in terms of x in the small-angle approximation, which gives the form

$$m\ddot{x} + (k_g + k_w)x = F, \tag{2.29}$$

where $F = k_g x_g$ is the external driving force, $k_g = mg/l$ is the effective gravitational spring constant and $k_w = k_{\text{el}}/l^2$ is the linear spring constant of the wire's anelasticity. With the constants grouped in such a way, Eq. 2.29 is in the exact form of Eq. 2.15 with a total effective spring constant $k = k_g + k_w$. As is known from the previous analysis that the resonant frequency of such a system is given by ${\omega_0}^2 = k/m$, it is clear that the effect of the wire's anelasticity is to shift the resonance frequency above the natural resonance of the simple pendulum.

Non-Linear Coupling

It may be tempting to suggest that a pendulum might have utility in vertical isolation as well as horizontal isolation, given that the pendulum Lagrangian intrinsically contains a component in x and z. Restructuring the coordinate definitions in Eq. 2.21 and 2.22 with an applied vertical displacement, z_g, results in the EOM

$$l\ddot{\theta} + g\theta + \theta\ddot{z}_g = 0, \tag{2.30}$$

where it was necessary to expand to second order to couple in the vertical ground motion. This equation is crucially different to Eq. 2.24 despite its very similar appearance, as the ground motion couples into the equation of motion via a purely non-linear term, $\theta\ddot{z}_g$. Nonlinear effects, by definition, cannot be treated in the LTI framework and are rather an inconvenient source of error and noise within measurements. As such, a pendulum cannot be used as a standard source of vertical isolation. However, it is fortunate that

the non-linear nature of the vertical ground coupling in a pendulum is inherently smaller than the first-order, linear effect. Therefore, whilst the pendulum does not help with vertical isolation, it does not substantially hinder it either.

Another consideration is that even for the horizontally driven pendulum, the vertical coordinate of the pendulum mass is perturbed by the ground motion. After all, the vertical degree-of-freedom is crucial as it is the only source of potential energy in the Lagrangian. This effect of the horizontal drive exciting a vertical displacement is an example of so-called cross-coupling. In the general case, the six input degrees-of-freedom (three linear and three rotational) transfer to the corresponding test mass degrees-of-freedom via a matrix of transfer functions: $\boldsymbol{x} = \boldsymbol{M}\boldsymbol{x_g}$, where $\boldsymbol{M}$ is a 6×6 matrix. The direct coupling is given by the diagonal elements of $\boldsymbol{M}$ and the cross-coupling is given by the off-diagonal elements. Generally, cross-coupling is undesirable, as this generates an inter-dependence between the degrees-of-freedom, which makes the system more complicated. It may even completely defeat the benefit of using a notionally well-constructed isolation system. For example, a system with excellent horizontal-horizontal isolation may become dominated by vertical-horizontal coupling, which prevents it from achieving its design performance. In the pendulum case, this cross-coupling is also a non-linear effect, given by the need for a second order expansion of Eq. 2.22, leading to $z = z_0 - l(1 - \theta^2/2)$. This second-order dependence on θ leads to a second-order dependence on x_g. By the same argument from the small angle approximation, this cross-coupling is also inherently small and thus not a significant effect. An example of cross-coupling as a primary cause for concern in a precision metrology application can be found in Ref. [38].

2.2.2 Suspensions

The pendulum and spring form the fundamental building blocks of conventional, passive isolation systems. A complete system consists of a combination and 'chaining together' of both types of components. As outlined above, the pendulum is used for horizontal isolation and the spring for vertical isolation. The combination of these components forms a suspension. Prominent examples include the *GEO600* [39, 40] facility's triple suspension [41], the *Advanced LIGO* [13] quadruple pendulum [42] and the *Advanced Virgo* [14] *superattenuator* [43]. All of these suspensions can fundamentally be modelled as pendulums and springs using a technique analogous to the Lagrangian analysis discussed thus far [44].

The treatment of a suspension is merely an extension of the treatment of the individual pendulum/spring components. As an intuitive expansion of the single pendulum, consider suspending multiple pendulums, each new pendulum attached to the mass of the previous. The Lagrangian can be constructed for the full system from the total kinetic and potential energy contributed by each pendulum mass. Such analysis can become very cumbersome due to the large number of terms, even if the mathematics of the situation is no more complicated. Section 3.5 contains a detailed exploration of this analysis technique applied to a real system. Therefore, the technique will only be mentioned in general terms here.

Broadly, for n pendulums arranged in a suspension, each extra pendulum stage adds one further resonance to the system and the final transfer function becomes the product of n individual pendulum-like transfer functions. That is to say,

$$T(\Omega) = \prod_{i=1}^{n} \frac{{\omega_i}^2}{{\omega_i}^2 - \Omega^2}, \tag{2.31}$$

where ω_i represents each resonance of the system in turn. It should be noted that the coupling between the pendulum stages redistributes the resonances of the system. This means that the resonances of the complete system are not identical to the resonances given by the resonance formula as applied to each pendulum in isolation.

In the high frequency limit (above all resonances), Eq. 2.31 scales as Ω^{-2n}. To understand the utility of a multi-stage suspension, consider two possible alternatives for an isolation system. The first is a 2 kg mass suspended from a 2 m pendulum and the other is a two-part suspension with two 1 kg masses suspended from a 1 m length each. Although the total mass and length of the suspensions are the same, the two systems behave quite differently. Fig. 2.5 shows the different transfer functions of the two suspensions. Of particular importance is the difference in the high-frequency suppression of the two designs, where the transfer function of the two-stage suspension eventually drops below the single stage, thus making the former a superior option over the latter.

2.2.3 Damping

Thus far, I have considered an undamped system, that is so say one with no dissipative component in its equation of motion. The result of this are the numerical infinities that are present in the transfer functions at the resonances (at every ω_i in Eq. 2.31). A more complete treatment includes the effect of energy dissipation, which ensures a finite response across all frequencies, even at the resoannces.

In a sense damping is an undesirable phenomenon due to its effect on, for example, the level of thermal noise (Section 2.3). On the other hand, a system with very little damping will generate large oscillations at the resonant frequencies, which can introduce a host of other problems. The high amplification of ground motion at the resonances

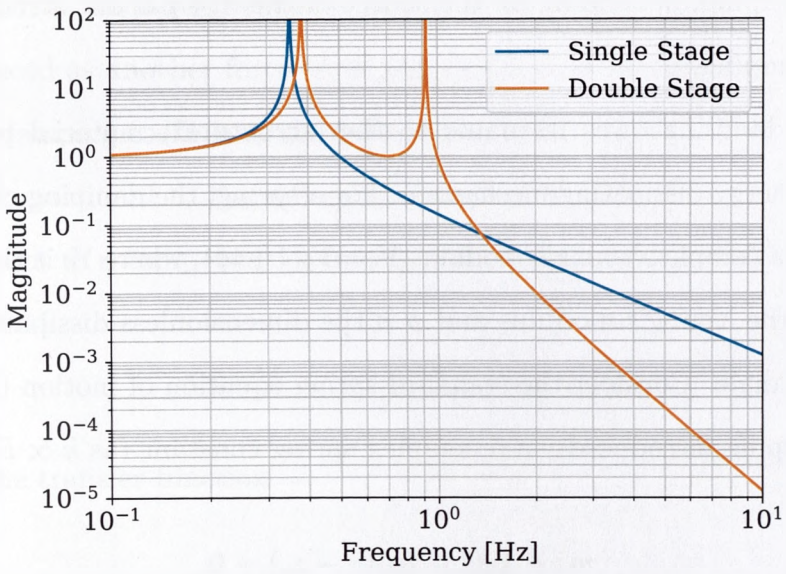


FIGURE 2.5: Magnitude of the transfer functions of a double-stage and single-stage suspension. Both suspensions modelled consist of a total of 2 kg of mass suspended over a total length of 2 m. In the double-stage case, the mass and length are equally distributed between the two stages. One point of note is the location of the resonances, where the lower resonance of the double-stage suspension has been slightly shifted with respect to the standard 2 m pendulum resonance due to the interaction between the two stages. Another aspect is the different high-frequency slope, which is steeper for the double-stage design, and thus provides more isolation.

may lead to the breakdown of the first-order approximation, such as that used in the pendulum-based system. This leads to an increase in the non-linear component of the motion and thus a potential increase in seismic noise [45]. A further negative effect is on the stability and controllability of the whole interferometer. In Section 1.2.5 I outlined the use of locking loops, which maintain the interferometer's sharp signal response. These schemes rely on the ability to respond sufficiently quickly to a control signal (such as the error signal outlined in Fig. 1.3). However, if the test masses swing too fast, this puts an excessive strain on the control loops, an issue that is relevant to the damping solution proposed in Section 3.5.5. Thus, it is often important to consider the level of damping and, if necessary, to strengthen it. There are two general ways of

introducing damping, depending on the source of the dissipation: structural damping and viscous damping.

Structural damping owes its name to the (structural) material properties that introduce the energy dissipating mechanism. For a spring, the damping can be modelled by introducing a complex Young's modulus $Y = Y_0(1+i\phi)$, where Y_0 is the familiar, real component of the Young's modulus and ϕ is the dimensionless dissipative component called the loss angle. Consider the standard spring equation of motion (Eq. 2.14) with structural damping introduced via a complex spring constant (as $k \propto E$)

$$m\ddot{z} + k_0(1 + i\phi)(z - z_g) = 0, \tag{2.32}$$

which results in the transfer function

$$T(\Omega) = \frac{{\omega_0}^2(1 + i\phi)}{{\omega_0}^2 - \Omega^2 + i{\omega_0}^2\phi}, \tag{2.33}$$

which already includes the standard result ${\omega_0}^2 = k_0/m$.

By considering the magnitude of Eq. 2.33 it can be shown that the response of a structurally damped system remains maximised at the undamped resonance, ω_0. However, the damping results in a finite response across all frequencies, with a maximum response of $|T(\omega_0)| = \phi^{-1}$. This measure of amplification can be identified as the so-called quality factor (Q). For a multi-stage suspension, each resonance may generically contain its own loss angle, which may also be a frequency dependent quantity. In such a case, a Q can be defined for each resonance, representing the height of the resonant peak and given by $Q_k = {\phi_k}^{-1}(\omega_k)$ for the k^{th} resonance.

The latter source of dissipation, viscous damping, arises from external sources that extract the kinetic energy from the system, such as drag through an atmosphere.

Consider the driven harmonic oscillator with a general force (Eq. 2.15). The damping can be introduced as another force term proportional to velocity, $F_d = -\gamma m(\dot{z} - \dot{z}_g)$, where γ is the damping coefficient, the minus sign implies a resistive force to the motion. The force is proportional to the difference between the velocities, as the damping is often with respect to a component or medium that is stationary in the lab frame. This leads to the EOM

$$\ddot{z} + \gamma(\dot{z} - \dot{z}_g) + {\omega_0}^2(z - z_g) = 0, \tag{2.34}$$

which yields the transfer function

$$T(\Omega) = \frac{{\omega_0}^2 + i\gamma\Omega}{{\omega_0}^2 + i\gamma\Omega - \Omega^2}. \tag{2.35}$$

The presence of viscous damping means that the position of the resonance is no longer preserved and, in the presence of strong damping, can induce a substantial shift. However, in the case of sufficiently weak damping ($\gamma \ll \omega_0$), the resonance remains very close to the natural resonance at ω_0 and the maximum response is given by the newly defined $Q = \omega_0/\gamma$. It is important to note that in Eq. 2.35 the numerator now contains an explicit frequency dependence. This term is the consequence of the damping taking place relative to the lab frame, which alters the high frequency response to fall off as $\sim \Omega^{-1}$.

Fig. 2.6 shows a system in the presence of no damping, structural damping and viscous damping. Both damped cases have been tuned for the same Q. The clear disadvantage of viscous damping over structural damping is its detrimental impact on high-frequency isolation performance. This is the consequence of the external source of the dissipative effect, which damps the suspension with respect to the moving lab frame rather than an underlying inertial frame.

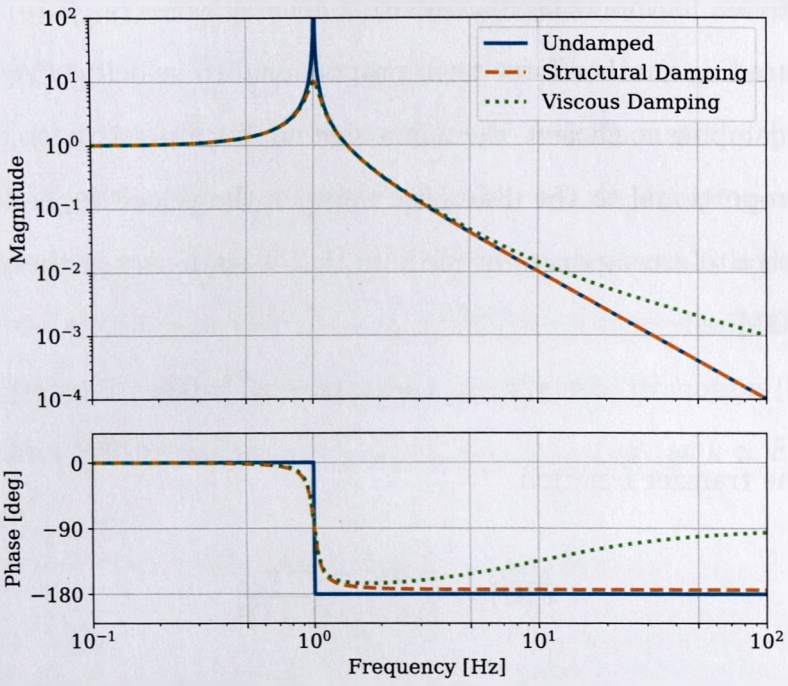


FIGURE 2.6: Bode plot of the transfer function of a generic single-stage isolation system with a pole frequency at 1 Hz. The system is shown in the case of no damping, structural damping, and viscous damping. The quality factor of the two damped examples is set at 10. This figure serves to illustrate two key effects of the damping system. Firstly, the presence of either damping source ensures a finite response of the system at the resonance with a smooth and continuous phase transition between the two discontinuous levels of the undamped response. Secondly, the two damping strategies show differing behaviour at frequencies well above the resonance. Whilst the structural damping approaches the undamped system in both magnitude and phase behaviour, the viscous damping curve follows a shallower gradient and a correspondingly lower phase lag of only 90°.

Viscous damping can be introduced to the system in a controlled manner using, for example, eddy current damping in the presence of metal components. Eddy currents can be induced by placing magnets in close proximity to a counterpart metal component without physical contact between them. One of these components is placed on the suspension/oscillator and the other on a support structure. In the case of a harmonic oscillator, relative motion between these two parts will result in damping in proportion

to the oscillator's relative velocity [46]. As the production of eddy currents relies on a relative difference between velocities, this damping manifests in the way described by the viscous damping formalism. Devices using eddy current damping have been considered for high performance suspensions of GW detectors [47].

It may seem more desirable to consider structural damping in preference to the viscous type. Unfortunately, there are few options for generating an effect of this kind. Most naturally, one could make use of the intrinsic mechanical losses within the chosen material but the mechanical loss in many materials is very small; the loss angle of fused silica—a high-performance material for suspension fibres—is in the range 1×10^{-8} to 1×10^{-7} [48, 49], which makes structural damping largely ineffectual in reducing the Q of the resonances to a manageable level. In another sense, it is fortunate that the loss angle is so low, particularly due its effect on thermal noise (Section 2.3). This is the first hint of the tension between competing requirements in the experimental design, which will be addressed in Chapter 4. In practice, there is little choice in the type of damping that is used, with controllable tunable damping falling typically into the viscous category.

The drawbacks of introducing strong viscous damping can be mitigated in a number of ways. Firstly, damping can be introduced in a limited and controlled manner using a feedback system. This falls under active control and is discussed in more detail in the next section. An alternative is through the use of an auxiliary suspension, such as that used in *Advanced LIGO* [50], which acts as the support structure for the actuation and damping components. This suspension filters the seismic noise through multiple isolation stages so that it is not directly injected into the relative motion between the test mass and support structure. Therefore, the lost factor of Ω in isolation at high frequencies can be partially recovered.

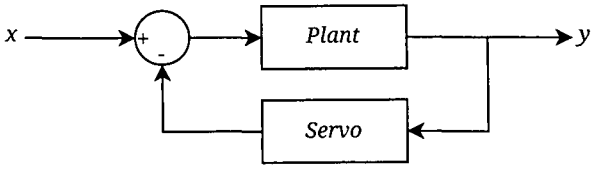


FIGURE 2.7: Diagram of an elementary feedback system. In the context of seismic isolation, the plant refers to the passive suspension and the servo refers to an electronic controller that accepts input from a displacement sensor and outputs to a force actuator on the suspended mass.

2.2.4 Active Isolation

As the name suggests, passive isolation systems work through their mere presence alone. There is no need for measurement or control to benefit from the noise-isolating capability. However, in certain applications, the performance of passive isolation systems may not be enough, or the situation is such that passive isolation is not appropriate or practical. In such cases, one can turn to the alternative approach of implementing an active isolation system instead (e.g. [51]).

In the most general terms, an active system allows for the compensation of an external disturbance through measurement and subsequent control of the system, usually through some form of feedback/feedforward scheme. The active system can thus be broadly separated into two key areas, those of sensing and actuation.

Displacement sensing can be achieved through a number of commercial, as well as custom, technologies. The sensing on and around the suspension is part of a so-called *local control* scheme [41] used in all contemporary GW detectors. *Advanced LIGO* utilises a multi-stage active isolation system [52] in addition to the multi-stage passive isolation system [42] and thus a variety of sensing techniques can be found on the chamber and suspension of a single test mass. For the initial stages of isolation, the suspension uses capacitive position sensors, optical levers [53] as well as a range of commercial sensors [52], such as the Nanometric Trillium 240 seismometer, Sercel L-4C,

and Geotech GS-13 geophones. Further custom sensing solutions also exist or have been proposed, such as the *beam rotation sensor* [54] and the 6D inertial seismometer [55,56]. The quadruple suspension itself is sensed by a number of *BOSEM* optical shadow sensors [50], with a range of interferometric sensors suggested as upgrades, such as the *HoQI* [57] and the *SmarAct* picosensor that is presented in Chapter 7. *Advanced Virgo* further uses the *linear variable displacement transducer* [58] based on magnetic readout. In terms of actuation, the landscape is generally more sparse with most techniques using some form of coil-magnet actuation such as that used in the *BOSEM* [50] and on the *Advanced Virgo* suspensions [59].

There are plenty of thorough reviews of control and feedback theory that speak of its application in general terms, such as [8]. However, discussing this theory in its general application at length is unnecessary and thus I will confine the analysis to specific examples as required, such as that discussed in 3.4. Instead, consider one example here due to its relevance to the damping techniques explored in Section 2.2.3.

Consider a very simple passive system controlled via a basic feedback loop as shown in Fig. 2.7. The passive system (in this case a simple spring) can be identified as the *plant*, following common control terminology. The displacement of the test mass attached to the spring is read out by some notionally ideal sensor with respect to the ground. This is then passed through a *servo* with a transfer function given by $G(\Omega)$, which controls an actuator that applies an external force on the test mass. The system can be modelled with reference to the spring EOM, driven by both the ground vibrations as well as an external force applied by the actuator (effectively combining Eq. 2.14 and 2.15). By transferring to the frequency domain immediately, this yields

$$-\Omega^2 Z + {\omega_0}^2(Z - Z_g) = G(\Omega)(Z - Z_g), \tag{2.36}$$

which in turn results in the transfer function

$$T(\Omega) = \frac{{\omega_0}^2 - G(\Omega)}{{\omega_0}^2 - \Omega^2 - G(\Omega)}. \tag{2.37}$$

Now, consider a couple of suggestions for the servo transfer function, G. Setting $G(\Omega) = -i\gamma\Omega$ immediately recovers Eq. 2.35, showing how the active system can be used to generate viscous damping, complete with its detrimental effects at high frequency. Instead, consider a modification to the servo where it is rolled off above the resonant frequency of the plant by a low pass filter, $G_{\mathrm{LP}}(\Omega)$, such that the new transfer function is given by $G(\Omega) = -i\gamma\Omega G_{\mathrm{LP}}(\Omega)$. Considering an idealised low-pass filter such that $G_{\mathrm{LP}}(\Omega < \omega_c) \approx 1$ and $G_{\mathrm{LP}}(\Omega > \omega_c) \approx 0$, where ω_c is the cut-off frequency and $\omega_c > \omega_0$, close to the resonance the transfer function becomes identical to the damped case. However, at high frequencies (above both ω_0 and ω_c), the transfer function recovers the same form as the undamped system. In that sense, it is possible to obtain the best of both systems: the damping of the resonance and the higher level of isolation at high frequencies. A comparison between the performances of the active and passive damping methods is shown in Fig. 2.8. This fairly rudimentary example may serve as motivation for the potential strengths of active systems that often justify their added complexity and cost.

2.3 Thermal Noise

It is a familiar concept that a thermodynamic system at a finite temperature will generally experience fluctuations on a microscopic scale, through randomised motion due to the distributed thermal energy. However, these thermally induced fluctuations lead to an ultimate uncertainty in even the macroscopic parameters of the system. Consider,

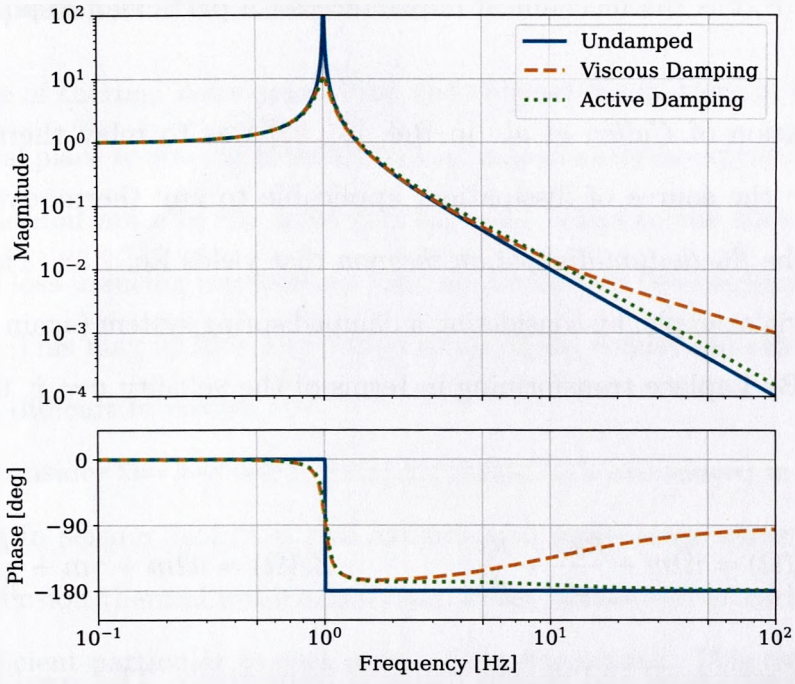


FIGURE 2.8: Bode plot of the transfer functions for a passive system that is undamped, experiencing viscous damping, and actively damped through feedback. The passive curves correspond to the system in Fig. 2.6. The active curve corresponds to Eq. 2.37 with an applied viscous-damping-type servo transfer function and cut-off at high frequency with a simple, single-pole low-pass filter at 5 Hz. This (by no means optimal) system serves to illustrate the key effect that the high-frequency isolation scaling of the undamped system can be recovered with active control.

as an example, the voltage fluctuations induced by Johnson–Nyquist noise [60], which relates the macroscopic quantity (voltage) to the energy dissipative property of the electric circuit—the resistance. Due to the pioneering work of *H. B. Callen et al.* [61,62], it is possible to extend the analysis to the general case and, in particular, derive the effect of thermal fluctuations in linear, mechanical systems. In a harmonic oscillator, the thermal noise consequently manifests as a fluctuating force that drives the system (as in Eq. 2.15) and has a PSD given by [63]

$$S_F(\Omega) = 4k_B T \mathfrak{Re}[Z(\Omega)], \tag{2.38}$$

where $Z(\Omega) = F/\dot{x}$ is the mechanical impedance for a particular displacement of the system, x.

The innovation of *Callen et al.* in Ref. [61, 62] was to relay thermally induced fluctuations to the source of dissipation, applicable to any thermodynamic system, giving rise to the *fluctuation-dissipation theorem* that yields Eq. 2.38. This relationship can be shown via example, by considering a damped spring system (again with reference to Eq. 2.15). By Laplace transforming in terms of the velocity $v = \dot{x}$, the impedance becomes

$$Z_{\mathrm{s}}(\Omega) = i\Omega m + \frac{k_0\phi}{\omega} + \frac{k_0}{i\Omega}, \qquad Z_{\mathrm{v}}(\Omega) = i\Omega m + \gamma m + \frac{k}{i\Omega}, \tag{2.39}$$

which apply to structural and viscous damping respectively. This results in a thermal force PSD given by

$$S_{F,\mathrm{s}}(\Omega) = \frac{4k_BTk_0\phi}{\omega}, \qquad S_{F,\mathrm{v}}(\Omega) = 4k_BT\gamma m. \tag{2.40}$$

The round-trip fractional loss in energy is given by $2\pi\phi$ or $2\pi\gamma/\omega_0$ [64], which shows that the loss angle, ϕ, and damping coefficient, γ, can be identified with the source of energy loss within the system. Together with Eq. 2.40, this makes explicit the direct link between the dissipative component of the system and the induced thermal motion.

There are numerous ways that thermal fluctuations can couple into the displacement measurement of an interferometer, from directly induced motion to optical effects. However, I will limit detailed discussions of thermal noise to the subset which is relevant to the work presented in the following chapters. The three sources of thermal noise that will be explored in more detail are the suspension thermal noise, the substrate Brownian noise, and the coating Brownian noise.

2.3.1 Suspension Thermal Noise

This source of thermal noise arises from the thermal fluctuations of the suspension. This is a natural place to start as it builds on Lagrangian analysis explored in Section 2.2. Suspension thermal noise in the most general sense refers to the thermal noise that arises from all loss-inducing mechanisms that are further up the suspension chain from the test mass. This may include any bolted connections, welds, and other sources that are, generally, difficult to model.

However, consider the loss angle/damping framework introduced in Section 2.2 to deal with loss in seismic isolation and extend it to multi-stage systems. In such a case, the suspension thermal noise arises from losses introduced at each stage, with a damping coefficient particular to each stage of the suspension. This treatment allows for an analytic solution to be obtained if the *effective* loss angle for each stage is known. This quantity is a combination of all the losses mentioned above and is generally a function of material properties, and aspects of the geometry, such as the details of the shape of the connection between the wires/fibres and the stage mass [65].

In the treatment of the multi-stage suspension, it was shown that each additional stage further enhances the high-frequency suppression of seismic noise (Eq. 2.31). Therefore, it is to be expected that the thermal noise should experience some form of filtering as it propagates down the suspension chain. Unfortunately, the level of complexity increases very rapidly with additional stages, which makes analysing a many-stage system rather complex and involved. However, key insights can be obtained by merely considering a two-stage system with damping at both stages.

Consider a two-stage spring-based system with two bodies of mass m_1 and m_2 suspended by springs with spring constant k_1 and k_2. Each mass experiences viscous damping of its velocity with coefficients γ_1 and γ_2, respectively. The masses are arranged

vertically with m_2 suspended from m_1 and the quantity of interest is the vertical position of m_2. Even such a system, which was analysed in Ref. [63], yields a rather cumbersome equation for the mechanical impedance (Eq 29 in Ref. [63]). In the high-frequency limit above both system resonances, the leading order term for each damping source is given by

$$\Re\mathfrak{e}[Z(\Omega)]\big|_{\gamma_1=0} = \gamma_2 m_2, \qquad \Re\mathfrak{e}[Z(\Omega)]\big|_{\gamma_2=0} = \frac{{k_2}^2\gamma_1}{m_1\Omega^4}. \tag{2.41}$$

The most notable observation from Eq. 2.41 is the dominance of the loss introduced at the lower mass, which maintains the same form for the force PSD as derived for the single-stage isolation in Eq. 2.40. In this sense, the additional isolation stages do not provide much benefit from the perspective of thermal noise reduction if the final stage is very lossy. In contrast, the thermal noise that couples into the lower mass motion via loss at the upper mass does experience significant suppression, with a frequency dependence proportional to Ω^{-4}. Consider that, in comparison, one additional suspension stage adds only a factor of Ω^{-2} in the suppression of seismic noise.

As shown in Eq. 2.41, the thermal noise that manifests at the test mass will be significantly more susceptible to losses that occur physically closer (i.e. lower down the suspension chain). This underscores the importance of keeping lossy components of the system as far removed from the test mass and suggests that the optimal strategy to reduce thermal noise should prioritise the closest stages to the test mass. As an example of this principle in action, consider contemporary, low-noise suspension designs, such as the *Advanced LIGO* quadruple pendulum suspension [42]. In this case, the first three stages are suspended by metal wires, whereas the final stage consists of a much more advanced monolithic construction, using precisely welded fused-silica fibres.

Although the oscillator treatment in Eq. 2.38 and 2.41 can be readily analogised to a pendulum-based suspension, it is worthwhile to elucidate further, as there are certain

effects that the spring-based treatment does not capture. First, consider that gravity provides some (often most) of the restoring force in the pendulum-based application. This is highly relevant as gravity is (as far as is understood) a lossless store of energy. As such, the calculation of thermal noise from a pendulum stage follows from the result in Eq. 2.29 and results in a mechanical impedance given by

$$Z(\Omega) = i\Omega m + \frac{k_g\phi_g}{\omega} + \frac{k_w\phi_w}{\omega} + \frac{k_g}{i\Omega} + \frac{k_w}{i\Omega}, \tag{2.42}$$

where setting $\phi_g = 0$, invokes the lossless gravity condition. From this, it is possible to obtain the PSD of the thermal force [66],

$$S_F(\Omega) = \frac{4k_BTk_0D\phi_w}{\omega}. \tag{2.43}$$

Written in this form, Eq. 2.43 can be readily compared to Eq.2.40, where the total spring constant $k_0 = k_g + k_w$ and $D = k_w/k_0$ is the so-called *dilution factor* [63, 66], which represents the 'dilution' of the round-trip energy loss due to only a part of the total energy being stored in the lossy anelastic potential. A fuller treatment of the pendulum-based thermal noise includes the contribution from the so-called violin modes of the pendulum [67–70], which represent the standing-wave resonances that form in the transverse direction along the suspension wires.

Naturally, one noise reduction strategy is to make the dilution factor as small as possible. As, in most situations, $k_g \gg k_w$, the dilution factor can be approximated to $D \approx k_w/k_g$. To this end, it is possible to consider non-circular geometries. For example, fused silica ribbons were investigated as an option for the early *Advanced LIGO* suspensions [71]. These allow for lower stiffness than circular wires in the beam direction (for ribbons oriented with the thin dimension along the beam axis), whilst

maintaining a high enough cross-sectional area to maintain safe stress levels. The trade-off is a stiffer suspension in the perpendicular direction, which means this strategy is more appropriate for an application where there is a defined direction along which thermal motion must be minimised (such as the beam axis for an optically probed test mass) and motion in the transverse direction is less important.

It is also appropriate to consider reducing the loss angle of the suspending wires directly. The low suspension thermal noise of *Advanced LIGO* is ensured by a unique set of fused silica fibres, which have a low loss angle (e.g. $\phi_{\text{tot}} \sim 10^{-9}$ for the fundamental violin mode [65]) compared to the more conventional steel wires used on the upper stages (e.g. $\phi_{\text{tot}} \sim 10^{-2}$ for the fundamental violin mode [66]).

2.3.2 Substrate Brownian Noise

So far, the discussion has not had to deviate significantly from a point-mass model of the pendulum-based suspensions. However, it now becomes necessary to consider the test mass as an extended body and understand the consequences. The propagation of thermally induced stresses in the test mass generates strains, which alter the internal structure of the test mass and deform the surface. As an extended body, the test mass can support a sequence of internal modes that arise due to the resonant amplification of these thermally induced stress waves throughout the material. One approach is to consider the vibrational modes of the mirror as individual harmonic oscillators [72]. However, exploration of this approach will be postponed until Chapter 6, as it is directly related to the work presented therein. Instead, consider the pioneering work of *Yu. Levin* [73]—the so-called direct approach—which is much more tractable in the frequency band of conventional GW measurement.

The insightful contribution in Ref. [73] was to express the mechanical impedance in the fluctuation-dissipation theorem (Eq. 2.38) in terms of the energy dissipated upon hypothetically applying an oscillatory pressure on the test mass surface. Crucially, the pressure must have an identical spatial profile to the intensity distribution of the light field used to perform the displacement measurement. This can be thought of as the complementary concept to that used in Ref. [72], whereby the displacement of the vibrating test mass surface 'seen' by the light field is given by the integrated displacement of the surface, weighted by the spatial profile of the light beam. In simple terms, for a Gaussian beam profile, the effective displacement measured by the beam is most sensitive to fluctuations of the mirror's surface at the centre of the beam, where the intensity is highest. A conceptually similar statement that arises from the direct approach is that the total energy loss induced by the applied pressure (and therefore the displacement noise measured) is weighted towards losses localised closest to the centre of the light beam on the test mass surface.

The thermal noise PSD in displacement is thus given by [73]

$$S_x(\Omega) = \frac{8k_BT}{\Omega^2}\frac{W_{\mathrm{diss}}}{{F_0}^2}, \tag{2.44}$$

where W_{diss} is the time-averaged power dissipated in the test mass and F_0 is the amplitude of the applied force. Through a set of substitutions, it can be shown (see Appendix B.2) that W_{diss} in a spring-based suspension in the high-frequency approximation is given by $W_{\mathrm{diss}} = (k_0\phi{F_0}^2)/(2m^2\Omega^3)$. Substituting this relation into Eq. 2.44 yields the same result that was arrived at via a direct application of the fluctuation-dissipation theorem that led to Eq. 2.40. Therefore, it is quite easy to motivate the validity of this approach through the example of a relatively straightforward and well-known system.

Although this simple consistency check shows that the direct approach can recover

the same results as previous thermal noise calculations [63,66], it is much more powerful. The more involved derivation in Ref. [73] shows that this approach generalises beyond systems that can be reduced to point-mass suspensions. Notably, this approach was applied to the case of a displacement measurement of a test mass surface, subject to deformation due to thermal excitation of the substrate's internal modes [73–75].

For the conventional GW detector, the internal resonances of the substrate are above and outside of the detection band. Therefore, it is possible to derive the substrate thermal noise using the direct approach by applying static pressure to the test mass surface. The original approach was applied in the derivation of thermal noise for an infinite-half-plane substrate and later works [74,75] have developed this approach further to more closely approximate the finite extent of real test masses. Keeping to the more tractable infinite half-plane solution [74]

$$S_x(\Omega) = \frac{4k_B T(1-\sigma^2)\phi_{\text{subs}}}{\omega\sqrt{\pi}Y_0 w_0}, \tag{2.45}$$

where σ is the Poisson ratio, ϕ_{subs} is the loss angle of the substrate, and w_0 is the beam radius[9]. A finite-mirror correction can be applied to this result to improve the accuracy of the noise model. The finite-mirror model is well approximated by Eq. 2.45 for smaller beam radii and the approximation worsens as the beam radius increases. However, this correction is rather cumbersome to insert here and applies a correction on the order of a factor of unity. For example, the thermal noise for a beam radius that is approximately half the radius of a mirror of equal thickness has a discrepancy of less than 50% [75] compared to the infinite half-plane result.

The primary limitation of the treatment that led to Eq. 2.45 lies in its static

[9]The beam radius is defined here as the distance at which the electric field amplitude has fallen by a factor of $1/e$ from maximum. The alternate definition of the beam radius as the distance, r_0, at which the intensity has fallen by $1/e$ is also sometimes used and simply relates as $w_0 = \sqrt{2}r_0$.

pressure assumption. This immediately undermines its validity at frequencies close to or above any internal resonances of the substrate. Although the direct approach is valid across all frequencies, solving the general system analytically becomes a significant challenge and alternative approaches may be considered instead (see Section 6.5). Nonetheless, obtaining an analytic equation for the substrate thermal noise that is valid over frequencies in the conventional detection band (well below substrate resonances) is, unequivocally, very useful. To underscore its usefulness, Eq. 2.45 will be directly applied in the construction of the noise budget in Section 4.3.

The treatment of substrate thermal noise considers so-called homogeneous thermal fluctuations in the substrate. These are a result of the bulk losses within the material. Beyond this consideration, there are distortions of the light field generated by inhomogeneous and optical effects, although these will not be discussed here. For further reading, particular considerations should be given to the so-called thermoelastic [75] and thermo-optic [76] noise.

2.3.3 Coating Brownian Noise

The optical coatings that are applied to the test masses in precision interferometric experiments need to be highly customisable (covering the full range between highly reflective and anti-reflective) and exhibit low optical losses. To achieve this, several techniques for depositing very thin layers of dielectric material have been developed and extensively used within precision interferometry generally and within GW instrumentation in particular.

The total coating thickness (often on the order of microns to tens of microns [77]) makes for a vanishingly small constituent of the total test mass volume. However, as the direct approach has shown, the location of the coatings on the surface of the

substrate means that their thermal noise contribution is disproportionate to their size, particularly in reflection, where the incoming electric field does not significantly penetrate into the substrate material. Furthermore, due to the substantially larger loss angle of conventional coatings ($\phi \sim 10^{-4}$ for silica/tantala (SiO_2/Ta_2O_5) based coatings [77] compared to $\phi \sim 10^{-9}$ for bulk fused silica [64]), the coating thermal noise dominates over the substrate contribution in all contemporary GW detectors.

For simplicity, consider a single layer coating of thickness d. The analysis of coating thermal noise follows as an extension of the direct approach used to obtain Eq. 2.45 and is given by [64]

$$S_x(\Omega) = \frac{8k_BT(1-\sigma)(1-2\sigma)d\phi_{\text{coat}}}{\omega\pi Y_0 {w_0}^2}, \tag{2.46}$$

where σ remains the Poisson ratio of the substrate and ϕ_{coat} is the loss angle of the coating material. With more complicated arrangements such as multi-layer coatings and anisotropic loss, the individual loss contribution of each subsequent layer and in all directions must be calculated separately and combined to yield the total coating thermal noise. However, this is merely an exercise in adding complexity through repeated application of the same principles that arrive at Eq. 2.46. Therefore, a more detailed derivation can be deferred until the specific application in a particular experimental configuration.

Due to the dominant nature of coating thermal noise over the substrate contribution, much discussion within the literature and research activity within the community is focused on the different strategies for reducing it. From Eq. 2.46 it is clear that a reduction in the noise can come from changes in the substrate (for example increasing the Young's modulus) or from improvements in the coating itself. To the latter point, much of the focus has been on reducing the loss of the existing coatings through better deposition and processing techniques, such as heat treatment of the coating

post-deposition [78].

A further avenue of research is in alternative coating materials with higher refractive indices (to reduce the coating thickness) or intrinsically lower loss, for example, the use of silicon-based coatings [79]. These alternatives may have advantages in the contemporary GW detector configurations but some solutions are better suited for future cryogenic operation and some may have a significant impact on aspects of the whole interferometer design. For example, the silicon coatings investigated in Ref. [79] were investigated at a laser wavelength of 2 µm which requires a change in both the laser and substrate to be used.

Finally, the optimisation of the thermal noise properties may come at the expense of the coating's optical properties, such as optical losses. As is evident from the variety of different techniques and material choices, their impact on thermal and optical properties, and their impact on the entire interferometer infrastructure, there is much to be considered when selecting the optimal coating solution for a particular experiment.

2.4 Laser Noise

The analysis of interferometers undertaken in Chapter 1 assumes that the input electric field is fully coherent up to the fundamental limitations imposed by, for example, quantum fluctuations. However, in a real laser system, the output field experiences fluctuations in both phase and amplitude due to classical effects, such as temperature and voltage/current fluctuations. These fluctuations in the field have rather different effects on the interferometer's readout and can be suppressed through different means.

Generally, these noises can be readily suppressed with existing, well-known techniques and do not pose a fundamental limitation to the sensitivity of contemporary

interferometric devices. Therefore, these noises are not considered fundamental and are often left out of noise budgets. In a sense, these noises are suppressed only as much as is necessary to achieve the required sensitivity level. If the need arises to reduce these noises further, it is generally possible to do so without a significant on the infrastructure.

2.4.1 Laser Frequency Noise

The fluctuation in the phase of the output field of the laser can be more conveniently analysed as a fluctuation on the laser's frequency or, analogously, in the laser's wavelength. The laser frequency noise can be most purely characterised by the spectral density of the frequency fluctuations in units of $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$. This noise is a specific property of the particular laser in question and is dependent on the technology used, the quality of the construction, and the stability of the control systems integrated into the laser package.

The laser frequency noise couples into the displacement readout of interferometric devices due to the symmetry between frequency and length in the interferometer response, as elucidated at multiple points in Section 1.2. The phase accumulated in a Fabry-Pérot interferometer is given simply by $\phi = 2\omega_0 L/c$ from Eq. 1.7, which results in the total derivative

$$\mathrm{d}\phi = \frac{\partial\phi}{\partial\omega}\mathrm{d}\omega + \frac{\partial\phi}{\partial L}\mathrm{d}L = \frac{2L}{c}\mathrm{d}\omega + \frac{2\omega}{c}\mathrm{d}L. \tag{2.47}$$

In the case of small perturbations, Eq. 2.47 can be used to approximate a small change in the cavity length as equivalent to a small change in laser frequency via the relation

$$\frac{\delta\omega}{\omega_0} = \frac{\delta L}{L_0}, \tag{2.48}$$

where ω_0 is the nominal (central) laser frequency and L_0 is the nominal (macroscopic) cavity length.

Due to this interchange between length and frequency, it is often important to reduce the laser frequency noise to below the sensitivity requirements of the given displacement measurement. Aside from the technological solution of constructing a better laser, there are a number of ways to reduce the laser frequency noise for a given setup.

One solution, which has already been discussed in Section 1.2.5, is the PDH locking scheme. This technique maintains the resonant condition between cavity and the electric field through a feedback to the cavity position or laser frequency. What matters to the PDH scheme is the round-trip phase accumulation inside the cavity, in other words, the suppression of $\mathrm{d}\phi$ in Eq. 2.47. As the phase change is a combination of both frequency and length fluctuations, the feedback loop will suppress both. However, the issue with this scheme is that the control signal to the locking servo is often used to infer the cavity displacement. If the frequency noise stabilisation is in loop with the displacement sensing, then the frequency noise is not isolated from the readout at all. As a result, there are two ways to proceed in suppressing the laser frequency noise in way that achieves a real increase to the SNR: pre-stabilising the laser frequency to an external reference, or reading out the displacement with reference to an out-of-loop local oscillator.

With regard to the first strategy, this is possible by obtaining a frequency reference that is more stable than the laser's own frequency control loop. One way to achieve this is with a very rigid auxiliary cavity. The laser can be stabilised with another PDH signal to this cavity first, with the main cavity then benefiting from the independently stabilised laser frequency. Another way to achieve the frequency pre-stabilisation is by locking the laser frequency to a frequency comb [80] or an atomic transition line [10].

Considering the latter strategy, the laser frequency is stabilised to the main cavity, meaning that an independent local oscillator is then required to extract the displacement fluctuations. In such a case, it is possible to, once again, utilise a frequency comb, assuming its own noise characteristic is sufficient for the application. Another technique to achieve this, which is implemented in Chapter 4, is to use another cavity locked to a second and independent laser. In this configuration, the two cavity signals can be interfered and the combined cavities' displacement can be inferred from the beat spectrum. The two cavities effectively become each other's frequency reference.

2.4.2 Laser Intensity Noise

The second laser noise to consider, the laser intensity noise, poses an issue in interferometric system due to the fact that the interferometer acts as a phase-to-intensity converter. Therefore, although the displacement measured by an interferometer is generally imparted upon the phase of the light, this information is converted into an intensity fluctuation, which then becomes indistinguishable from fluctuations of the laser intensity at the input to the device.

In this case I am considering the classical effect. As such, this noise should not be mistaken for the shot noise, which also leads to an intensity fluctuation as measured by a photodetector. Crucially, the QSN scales with the square-root of the measured power and is sourced by vacuum fluctuations that enter at every lossy component of the interferometric system. Conversely, the laser intensity noise is simply a fluctuation generated at the laser due to, for example, imperfect stabilisation of the laser's driver. Therefore, passive optical components do not affect the SNR of the laser intensity noise, as they will attenuate and amplify both the signal and noise equally. Due to this

simplified behaviour, a particular laser is often characterised in terms of its relative intensity noise (RIN), as this is the conserved quantity.

One way that an interferometric measurement can be affected by the laser intensity noise is by directly obscuring the interference-caused intensity fluctuations at the output. Consider a Michelson interferometer operated at the mid-point between the bright and dark fringes. As mentioned in Section 1.2.1, the response to small fluctuations is given by kP_{in} W m^{-1}. For an example displacement sensitivity of $1\,\text{pm}/\sqrt{\text{Hz}}$ and a laser wavelength of 1550 nm, this results in a RIN requirement of $8 \times 10^{-6}\,\text{Hz}^{-0.5}$, this is within the reasonable performance level of existing commercial laser systems. However, it does suggest that laser intensity noise is a key noise source to consider in novel, sub-picometre interferometric displacement sensors and is, indeed, discussed in some detail in Chapter 7.

Another issue with intensity fluctuations arises in light-test-mass, suspended interferometers due to the classical radiation pressure effect. In this effect, the fluctuating intensity excites the dynamics of the test mass and thus generates a true displacement rather than simply affecting the readout. This should, again, not be confused with the QRPN, which is a consequence of the quantum-level fluctuations of light, even though the mechanism of driving the cavity mirror with the photon's momentum is identical in both cases. To understand this effect, consider the radiative force on a fully reflective mirror, given by $F = 2P/c$. For frequencies above the mirror's suspension resonance, in the free-mass regime, this results in a displacement noise of

$$S_x(\Omega) = \left(\frac{2}{mc\Omega^2}\right)^2 S_P(\Omega). \tag{2.49}$$

This can have a significant effect on high-sensitivity measurements using light mirrors. For example, a 10 g mirror operating at an SQL frequency of 100 Hz can tolerate a

maximum power fluctuation of $3 \times 10^{-7}\,\mathrm{W}/\sqrt{\mathrm{Hz}}$. This, again, places somewhat stringent requirements on the performance of the laser or constraints on the maximum pump power allowed. Whilst this longitudinal-driving effect is worth considering, the, possibly more important, effects to consider are those concerned with the driving of the test masses' other degrees of freedom. These can result in instabilities, such as the so-called Sigg-Sidles instability [81], which can lead to a failure of the locking scheme, thus preventing any measurements from taking place altogether.

2.5 Summary

The continued success of precision interferometry is due to the efforts made to suppress or otherwise mitigate the noise sources that limit the sensitivity of these devices. In this chapter, I outlined the three most notable types of noise present in interferometric sensors, particularly contemporary GW detectors, which are the quantum, seismic, and thermal noise.

The quantum noise PSD in a conventional Fabry-Pérot interferometer is given using an intuitive, approximate derivation as well using the full quantum-mechanical approach. The quantum noise poses a fundamental limit to the operation of suspended interferometers. However, it can be manipulated in ways that allow for improvements in sensitivity, which is the topic of Section 3.1.3.

The seismic and thermal noises severely limit the sensitivity of contemporary GW detectors in different parts of the frequency spectrum. They are formidable obstacles that require dedicated and advanced solutions to be overcome. In this chapter, I devoted much time to discussing the many and varied strategies for seismic isolation and laid out the framework for analysing the isolation systems. The sections on thermal noise

present the approaches for deriving the thermal noise from different components of the system based on decades of pioneering work. I also outline some of the considerations that are made when selecting optimal materials and designing optimal geometries.

Chapter 3

Macroscopic Quantum State: Motivation and Suspension

The quantum nature of light generates fluctuations in the readout of interferometric measurements of test mass displacement. This poses a formidable obstacle to further improvement in precision measurement. Whilst the quantum fluctuations are observed in contemporary detectors, operation at the level of the SQL has not yet been achieved. On the one hand, a device that operates in the quantum-noise-dominated regime would provide useful insight into the nature of quantum noise and serve as a test-bed for SQL-surpassing techniques. In another sense, a quantum-limited system would be a valuable resource for exploring unanswered questions in fundamental physics. In this experiment, we consider (a) test mass entanglement via the QRPN [30], (b) semi-classical gravity manifest through test mass self-interaction [82], and (c) quantum gravity effects seen through the entanglement of adjacent optical cavities mediated via the gravitational interaction [83].

This chapter presents the design of a table-top experiment capable of reaching quantum-noise-limited sensitivity at the level of the SQL. The topic of quantum noise measurement and suppression is discussed in a contemporary context and the motivation

and strategy of the experiment are laid out. The remainder of the chapter is then devoted to the construction and performance of the larger infrastructure and the seismic isolation, with a particular focus on the internal suspension used to isolate the test masses from ground motion and cryocooler vibrations.

The discussion is continued in the next chapter, which focuses on the aspects of cryogenic performance and thermal noise. That chapter (and thus the whole discussion of this experiment) concludes with a presentation of the current status of the experiment as well as the key initial results.

The work presented here is my own and mostly consists of previously unpublished content. There is some overlap between the content in this chapter and the conference proceedings paper *J. Smetana, D. Martynov*, "Towards the Standard Quantum Limit in a Table-Top Interferometer" [84] of which I was the principal author.

3.1 Outline

Before considering the designs of the novel experiment, it is worthwhile to discuss the extent to which the quantum fluctuations have been measured so far, why the SQL is so elusive, and what strategies will be employed to reach it and, eventually, surpass it.

3.1.1 Contemporary Sensitivity

Reaching quantum-noise limited sensitivity within a part of the detection window is not such a challenge for contemporary gravitational wave detectors. For example, the *Advanced LIGO* [13] and *Advanced Virgo* [14] detectors are both constrained by QSN over a broad frequency band starting at frequencies above a few hundred Hz. Neither

detector is, however, operating at a level of sensitivity constrained by the QRPN, with the latest *Advanced LIGO* observing run showing the dominance of technical noises in the low-frequency band [17]. Neither detector is more than an order of magnitude away from measuring sensitivity at the SQL but in both cases, the detector is fundamentally prevented from reaching this level of sensitivity by the coating Brownian noise [13, 14]. The shot noise limit has also been reached in experiments outside of the field of GW detection in table-top experiments [85, 86], such as atom interferometry [87].

The radiation pressure noise is elusive due to its dominance at low frequencies, which brings it below classical noise sources that, generally, increase towards lower frequencies. This can, of course, be improved upon with enhancements in low-frequency suppression, particularly of seismic noise, which is indeed a topic of serious enquiry [88]. However, it is possible to increase the radiation pressure component through an appropriate selection of the parameters in Eq. 2.6 (e.g. increasing the incident radiative power on the test mass) and by performing the measurement over a higher frequency band, where other noises don not pose such an obstacle. This measurement was performed by *Cripe et al.* using a cantilever resonator in the 2 kHz to 100 kHz band [89].

3.1.2 SQL in Future Devices

Although the SQL is rather elusive, future GW detectors will continue to improve on the current sensitivity and will operate at near-SQL levels, which provides ample motivation for understanding and testing such systems on a smaller scale. Results from these tests will be very helpful to future detectors such as the Einstein Telescope which is designed to be broadly quantum-noise limited [90]. The growing adoption of cryogenic operation, such as in the current *KAGRA* [91] and future *LIGO Voyager* [92] detectors,

will lead to a suppression of thermal noise that will, in turn, make quantum-noise suppressing techniques even more essential.

The intrinsic link between the SQL and macroscopic quantum mechanics has naturally led to proposals for sites that can operate at this level of sensitivity. One such example is the *AEI* 10 m prototype facility [93]. This detector consists of a Michelson interferometer that will measure the quantum-noise limited displacement of 100 g test masses and allow for the exploration of phenomena ranging from macroscopic test-mass entanglement [94] to tests of quantum gravity [95].

3.1.3 Beating the SQL

Beyond the measurement of the SQL and the resulting quantum mechanical phenomena, there is a concerted effort to surpass the SQL in preparation for the stricter sensitivity requirements of the next generation of GW detectors. Although the SQL has not been measured directly, many of these techniques provide suppression over a broad range of frequencies such that the improvement can be observed in other parts of the spectrum. The limitations on sensitivity imposed by the quantum noise might seem to be fundamentally unavoidable. However, this is only true for a 'conventional' measurement like that used to derive the quantum noise spectrum in Eq. 2.7. The key insight is that typical measurements probe observables that do not commute with one another or with themselves at different times [25]. Therefore, improvements can be made by altering the measurement technique to perform measurements of only observables that minimally perturb the system, first proposed by *V. B. Braginsky* and termed *quantum nondemolition* [96].

Another significant leap in the suppression of quantum noise came in the form of the so-called *squeezed state* of light [25, 28, 29, 97]. This technique stems from the ability

to manipulate the noise content over different phases of light, which can be analysed in terms of the two quadratures introduced in Section 2.1.2. In the coherent state, these are identified as the amplitude and phase quadrature. In simple terms, the introduction of squeezing involves the reduction of noise in the quadrature that contains the desired signal at the expense of a corresponding increase in noise (anti-squeezing) in the other quadrature. This preserves the constraints of the uncertainty relation between the two observables, whilst effectively transferring the quantum noise onto a channel that does not participate in the measurement. This technique has been experimentally verified in table-top applications [98, 99] and also implemented within the *Advanced LIGO* and *Advanced Virgo* detectors [100, 101].

A proposal for achieving the squeezing effect internally (i.e. without the addition of a dedicated squeezer cavity) follows from the signal recycling already implemented in contemporary GW detectors [102]. Via careful tuning of the signal recycling cavity, the sensitivity curve can be reshaped, which arises from the correlated nature of the shot noise and radiation pressure noise [103, 104]. At the expense of worse sensitivity at low frequencies, it is possible to surpass the SQL within a particular frequency band.

The solutions proposed thus far comprise innovative ways of reducing the quantum noise within the same displacement measuring framework. This leaves room for much broader speculation, such as completely rethinking the observable of the system that is being probed. The test mass momentum (in the free-mass approximation) commutes with itself in time, thus circumventing the additional noise imposed by the conventional displacement measurement. This led to the theoretical development of the speed meter interferometer [105, 106]. Building upon these initial explorations, speed meter designs have been proposed for future generations of GW detectors [107–109]. Prototype facilities such as the one being built at the *University of Glasgow* may verify this

concept experimentally [110].

A final, novel approach to quantum noise suppression is the use of a phase-insensitive quantum amplifier that leads to signal enhancement without a corresponding increase in the quantum noise. An implementation of such a device that uses a filter cavity augmented by a mechanical oscillator forms the basis of Chapter 5 and a detailed discussion can be found therein.

3.1.4 Experimental Proposal

The creation of a quantum-noise limited system is evidently a challenging task that presents many obstacles. However, from the practical development of GW detector technologies to fundamental physics research, the scientific case for creating such a system was laid out in the previous section. In this chapter, we present a novel design for a table-top interferometer capable of operating at a quantum-limited sensitivity over a frequency band containing the SQL.

The overarching design requirements we lay out are for (i) an interferometer that consists of gram-scale mirrors, (ii) a high-finesse optical cavity, (iii) cryogenically cooled operation within the confines of a cryostat, (iv) a multi-stage suspension for seismic isolation, (v) quantum-noise limited operation in the 'free mass' regime and (vi) an SQL frequency of approximately 100 Hz.

The mass constraint in (i) represents a compromise between two competing aims. On the one hand, we wish to create a truly macroscopic system that consists of large optics akin to contemporary GW test masses. Furthermore, light suspended masses are difficult to operate as the classical radiation pressure induces angular instabilities [81], to which the lighter masses are more susceptible. However, due to the inverse scaling of the SQL boundary with mass (Eq. 2.9), the SQL sensitivity can be reached more

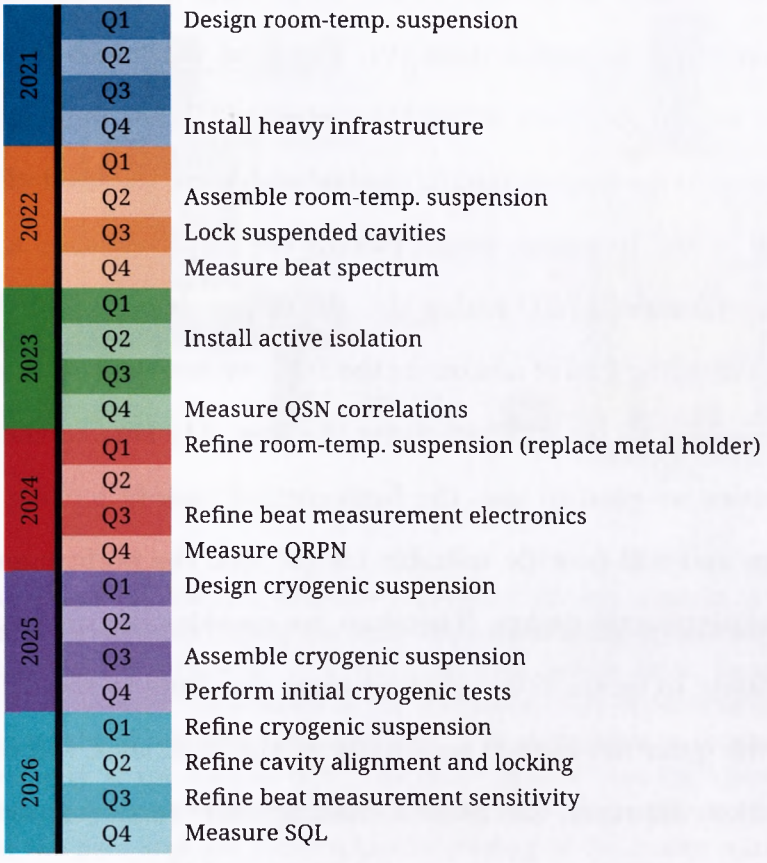


FIGURE 3.1: Timeline of experimental milestones. The timeline shows notable past events at the point they were achieved and an estimate of future milestones according to our predicted experimental progress.

easily with lighter test masses. In this sense, we depart from the strategy followed in GW detectors where the quantum noise should be minimised. Ultimately, we chose conventionally sized ($\sim$10 g) laboratory optics as a compromise and for ease of handling within a table-top environment. Requirement (ii) is particularly important for reaching a high level of quantum back-action by amplifying the carrier and vacuum fields within the cavity and enlarging the radiation pressure on the suspended mirrors.

Requirements (iii) and (iv) are the necessary practical steps that must be taken to suppress the two primary noise sources in the system—the thermal noise and seismic

noise respectively. These steps will enable us to achieve the primary goals of the experiment, represented by requirement (v). The final, and much looser, requirement (vi) connects our work more firmly within the context of GW instrumentation. Although creating a quantum-noise-limited interferometric system at the level of the SQL would be a novel result in any frequency range, placing the SQL frequency approximately in accordance with *Advanced LIGO* makes the connection clearer.

Due to the challenging goal of measuring the SQL, we have laid out a set of milestones, (1)–(4), that will have to be achieved along the way. (1) Due to the high finesse of the optical cavities we plan to use, the fundamental task of locking the cavity will pose a challenge and will provide valuable insight into the performance of the other aspects of the experimental design. Therefore, we consider the successful and reliable locking of the cavity to be the first milestone along the route to the SQL measurement. (2) Measuring the quantum-limited sensitivity at the SQL level will certainly require cryogenic operation. However, this poses a challenge over room temperature operation due to the stricter material requirements, deformation of the components upon cooling, and other such considerations. It is, therefore, valuable to consider an intermediate result that is achievable at room temperature. Due to the inverse scaling of the QRPN PSD with laser power (Eq. 2.6), we may deliberately amplify the radiation pressure noise to dominant levels once the cavity is locked and operational. The measurement of the quantum back-action effect within the free-mass regime would constitute a novel result. (3) We will face many challenges in the eventual transition to cryogenic operation, including a possible redesign of the suspension scheme, particularly in the final stage, where we will consider a quasi-monolithic suspension. Therefore, successful locking and long-term stability under cryogenic conditions would be another essential result. (4) Finally, a broadband measurement of the quantum noise at the level of the SQL is the

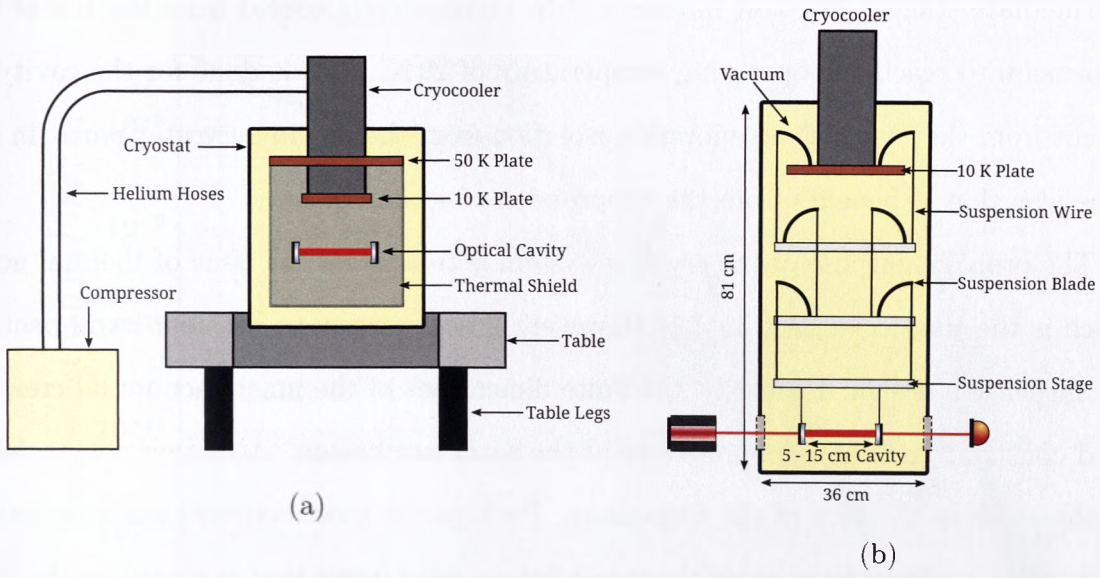


FIGURE 3.2: Panel (a) shows a simplified sketch of the key aspects of the system. The cryostat is shown placed atop the table connected to the grund with rigid legs. The cryocooler head is connected to a compressor via helium lines. Panel (b) shows the schematic of the inner components of the cryostat system in relation to the optical system (thermal shield removed for simplicity). The suspension is shown in its place isolating the cavity at the bottom of the vacuum chamber from the vibrations of the ground and the cryocooler that couple in through the 10 K plate at the top. This plate acts as the heat sink for the conductive cooling of the cavity mirrors via the suspension wires. The cavity is pumped and read out via a laser and photodetectors, respectively, placed in-air.

ultimate goal of the experiment and comprises the last milestone. A more granular breakdown of notable events and milestones is shown on the timeline in Fig. 3.1.

3.1.5 Experimental Schematic

The outline schematic of the experiment can be found in Fig. 3.2, which highlights the key aspects of the design, this being the cryostat system, the internal suspension, and the cavity. This diagram may serve as a point of reference for later detailed discussions of the design. The cavity will be formed by the test masses located at the bottom of the suspension to maximally benefit from the seismic filtering of the

intermediate stages. The test masses will be conductively cooled from the top of the suspension to reach the operating temperature of 20 K. This is done for the cavity to benefit from the suspension's vibration isolation from the moving cryostat parts, in the same way that it benefits from the suppression of seismic noise.

The primary purpose of the cryostat system is to address the issue of thermal noise, which is the subject of Section 4.1. However, it is necessary to consider its impact on the suspension design. Primarily, the finite dimensions of the inner vacuum lid create a hard constraint on the maximum size of the total suspension. A further constraint is on the attachment point of the suspension. Perhaps the most conventional attachment mechanism would be to suspend the stages from a rigid frame that is placed on the floor of the vacuum chamber. In this case, this is not desirable as (a) we would not create good thermal contact between the cold plate and the suspension, and (b) we would create a path for conductive heat transfer between the relatively hotter cryostat walls and the suspension. As a result, our design is to suspend from the cold plate directly.

3.2 Seismic Isolation Overview

To understand the requirements for our isolation design, we must first consider the ambient seismic motion. Fig. 3.3 shows the horizontal displacement amplitude spectral density (ASD) of a rigidly connected table top. The optical cavity must be isolated from this motion via multiple stages of seismic filtering to reduce the residual seismic noise to a level that is sufficiently below the quantum noise. We may use Eq. 2.9 to approximate the level of sensitivity required to reach the SQL. With a mirror mass of 10 g, we obtain a spectral density of $4.6 \times 10^{-19}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at the SQL frequency of 100 Hz. By comparing this to the spectral density at 100 Hz in Fig. 3.3, we can determine that

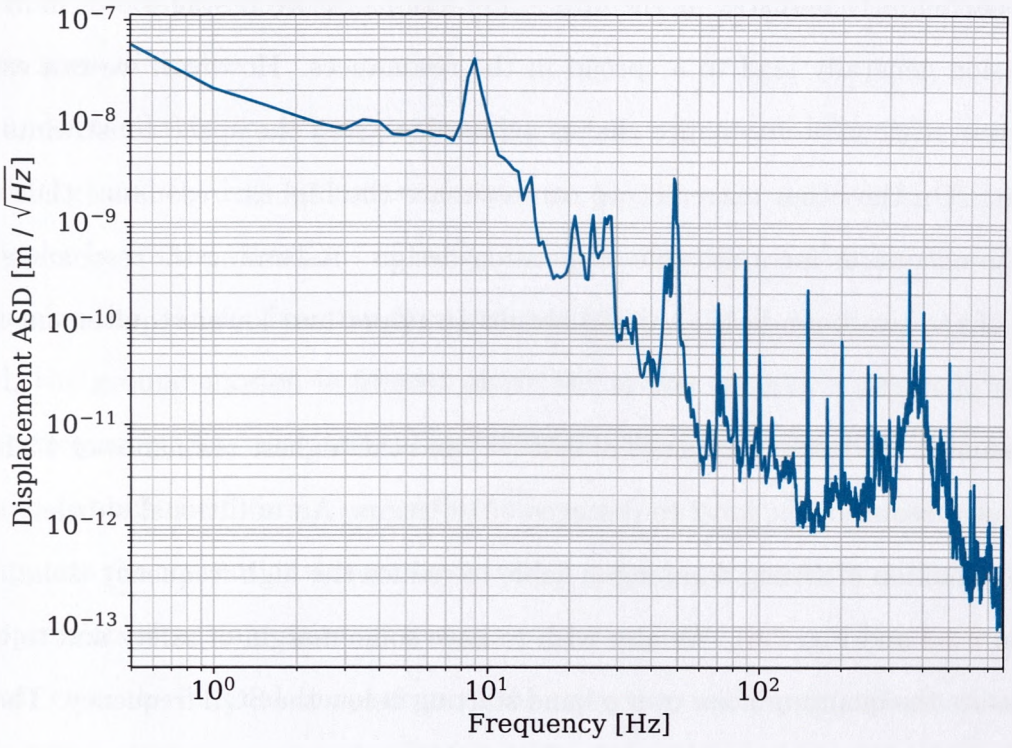


FIGURE 3.3: Horizontal, seismically-induced displacement spectrum of a table top connected with rigid legs. Although this motion is not the most direct measurement of the ground, it serves the purpose of later analysis better, such as the performance analysis of the seismically isolating suspension, whose input is better represented by the motion of the table over the motion of the floor.

a suppression factor of at least 10^8 is required.

We can use this result to motivate the number of isolation stages that are needed to achieve a sufficient reduction in seismic noise. Using Eq. 2.31 in the high-frequency limit, and further assuming that all resonances occur at the same frequency, we obtain

$$N \approx \frac{\log(T(\omega_{\mathrm{SQL}}))}{2(\log(\omega_0) - \log(\omega_{\mathrm{SQL}}))}, \tag{3.1}$$

where $T(\omega_{\mathrm{SQL}})$ is the desired value of the transfer function at the SQL frequency, in other words the required level of suppression. The approximation of identical resonances

is not particularly realistic, as the cross-interactions between stages of a multi-stage suspension generally lead to a spread in the resonances. However, we can estimate the lowest reasonable resonance that is achievable given the space constraints of the cryostat. On the other extreme, we can estimate the highest resonance that we are likely to encounter for a sensible suspension design. Assuming all resonances lie at either of the two boundaries, we can obtain an upper and lower requirement for the number of stages.

Assuming the lowest resonance of 0.55 Hz[1] and the highest resonance of 4.5 Hz[2], we arrive at a requirement for 2 to 3 stages of isolation. An additional subtlety is that the introduction of strong damping is liable to reduce the high-frequency isolation (see Section 2.2.3 and Fig. 2.6). We also wish to have some margin of safety and to be able to measure the quantum noise over a band starting below the SQL frequency. Therefore, to ensure a sufficient margin, we opt for four stages of isolation.

3.3 Optical Table Isolation

Before we consider a dedicated seismic isolation solution, we may first consider filtering via the optical table. It is common to use floated tables to provide some level of seismic isolation, which provides some stability even to the in-air optical components. This strategy generates a soft connection between the table surface and the table legs, which is indeed useful for the isolation of motion that couples to the table via the table legs. However, another significant source of vibrations to consider is the cryocooler head that connects to the tabletop from above. This direct connection to the tabletop

[1]The lowest resonance is given approximately by a pendulum of length 81 cm, which is the longest pendulum that can be suspended in the vacuum chamber (see Fig. 3.2).

[2]Given our plan to use standard 1 in mirrors, the shortest suspension we can create is approximately 0.5 in in length.

not only bypasses the leg isolation but will in fact make the overall task of vibration isolation more difficult.

To understand this, consider a single-spring system with a mass given by the tabletop ($m_t = 600\,\mathrm{kg}$) and the legs acting as springs with some spring constant, k_l. For motion that couples in via the vibrations of the ground, we follow the approach of Section 2.2.1 to determine the transfer function from ground motion to table motion. As previously derived, the ground motion is filtered above the spring resonance ($T_x \approx {f_0}^2/f^2$ for resonant frequency, f_0), with unity response below. Thus, it is preferable to reduce the resonance to shift the filtering property of the spring to the lowest possible frequency. This supports the conventional requirements for a softer leg connection, which prioritises the suppression of seismic noise.

In our case, the internal components of the cryopump undergo significant vibrations as well. Although the cryopump assembly is isolated from direct contact with the rest of the cryostat, a connection remains between the two parts via a bellows-like component that transfers the vibrations from the cryopump to the vacuum chamber. The bellows connection has some unknown force transfer and does not necessarily behave as an ideal spring. However, we can still make a general argument about the effects of a stiff connection between optical table and ground.

The optical table will experience a force with some, as yet unknown, spectral density which is applied to it by the cryopump. The response of the table to this generic force is given by

$$T_F = \frac{1}{k_l - m_t\Omega^2}. \tag{3.2}$$

At high frequencies, the transfer function scales as the free-mass response ($1/(m_t\Omega^2)$), which is independent of the leg stiffness. At low frequencies, however, the response scales as $1/k_l$, meaning that the stiffer the legs are, the lower the response of the table

top to an input driving force from above. The optimal filtering thus works in the opposite sense to filtering of motion from the ground, as it is ideal to raise the stiffness of the legs and the resonant frequency to bring the system into the stiffness-limited regime over the broadest frequency region.

With the two competing constraints in mind, we consider the isolation of cryocooler vibrations to be more important at the tabletop stage. We can design an internal suspension that is theoretically capable of isolating the ground vibrations within the vacuum chamber. As such, it is not entirely necessary to implement additional pre-isolation in the form of a floated table top. In contrast, we anticipate the cryocooler vibrations to be substantially larger than the ground motion and thus providing stiffer resistance to these vibrations has been deemed more important.

3.4 Cryostat Frame

The next stage in the isolation chain to consider is the frame that connects the cryostat to the optical table. The cryostat windows that give access to the cavity are located at the bottom of the cryostat's vacuum lid, as shown in Fig. 3.2. Therefore, it is necessary to raise the cryostat to place the windows above table-top level, in line with the in-air optics. Via the same argument used to justify a stiff table-leg connection in the previous section, it is optimal for this frame to be as stiff as possible as well. Our design was made to conform closely to a tetrahedron in shape to benefit from the high rigidity of this geometry[3].

Fig. 3.4 shows the schematic of the cryostat frame system with a more detailed set of technical drawings shown in Appendix D.1. The tetrahedral shape is achieved through

[3]The intrinsically rigid nature of the tetrahedral geometry leads to its ubiquitous use in engineering for load-bearing components such as trusses.

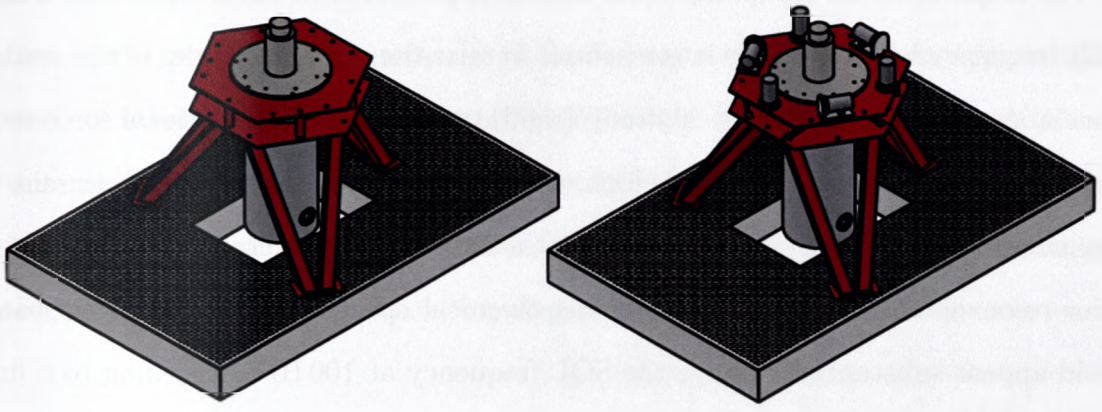


FIGURE 3.4: Computer-aided design (CAD) drawing of the cryostat support frame in two different configurations: passive (left), and actively controlled (right). The frame is shown in the context of the surrounding infrastructure, such as the table and the supported cryostat. The frame itself is colourised in red to highlight the components that belong to the frame assembly. The frame is designed to allow for easy transition between the two modes of operation, whereby the pillars connecting the two hexagonal support plates can be replaced with piezo actuators in situ.

the meeting of the frame legs at three points on the tabletop defining an equilateral triangle. The planes between each set of legs form the remaining triangular faces of the tetrahedron. In order to allow room for a support plate at the top of the frame, the frame legs cannot meet at a point to fully complete the tetrahedral geometry and, as such, the actual shape is truncated from above. The frame's height is defined by the height of the vacuum lid, which places the cryostat attachment plate 90 cm above the tabletop. The leg separation is constrained by the width of the table, which is 150 cm. The frame height is fixed, which then defines the ideal frame size to achieve equilateral side faces. Fortunately, the table is sufficiently wide to accommodate the necessary spread of the frame legs. As such the effective dihedral angle of the tetrahedron formed by the frame is within one degree of the 70.5° dihedral angle of a regular tetrahedron. This seemingly small mismatch does not account for the ambiguity in defining the correct plane from which to measure in the more complicated frame design.

The majority of the cryopump RMS motion is predicted to occur within the 1 Hz to 10 Hz frequency band and thus it is essential to raise the resonant modes of the cryostat frame above this band to ensure a strong (rigid) coupling of the vibrational force to the table and, via the already rigid table legs, to the ground. It is also preferable to raise the resonances well above the detection band of the suspended interferometer to prevent the frame resonances from appearing in the displacement spectrum. Ideally, the resonances would appear substantially above the SQL frequency at 100 Hz. According to a finite element analysis (FEA) of the cryostat frame design, the lowest frame resonance occurs at approximately 50 Hz. This achieves the stricter requirement of being above 10 Hz but is within the detection band. However, as determined by a finite element model, the lowest resonances are dominated by the azimuthally distributed modes of the hexagonal top plates. These modes appear to couple very weakly to inner cryostat degrees of freedom and thus may not appear in the cavity spectrum. It is yet to be seen whether these modes will be present in the final analysis.

The cryostat frame has been designed with two modes of operation in mind. The LHS design in Fig. 3.4 shows the passive frame with the upper hexagonal plate (directly bolted to the cryostat) connected via rigid, removable pillars to the lower hexagonal plate (welded to the frame legs). In this configuration, the frame's only use is to rigidly support the cryostat above the table. In the RHS design, we propose to actively control the cryostat for additional vibration isolation. We achieve this by replacing the pillars connecting the upper and lower hexagonal plates with 6 piezoelectric actuators, which are controlled through feedback from commercial sensors, such as the Sercel L-4C geophone, mounted atop the upper support plate. This isolation will be akin to the control, in feedback, of a Stewart platform [111]. Such a process has been implemented in the context of cryogenic suspensions [112] and the testing of inertial

seismic sensors [113].

Due to the potential for operation of the frame in the rigid or actively stabilised scheme, we label the cryostat frame as Stage 0 of the seismic isolation stack. Should the internal suspension (Stage 1–4) prove to be insufficient for the suppression of environmental vibrations, the potential for an active Stage 0 upgrade serves as a contingency plan in case more isolation stages are needed.

3.5 Internal Suspension

In the initial set-up of a passive Stage 0, the cryostat is rigidly connected to the ground, which means that the cryostat cold plate will experience vibrations with a PSD that closely resembles that of the ground motion. Therefore, the first stage of the seismic isolation begins at the internal suspension of the cavity (sketched out in Fig. 3.2.

The most essential suppression is that of horizontal motion along the cavity axis, which directly couples to the displacement readout of the suspended cavity. With this in mind, we base our suspension design on the filtering capability of the pendulum, which was discussed in detail in Section 2.2.1. The requirements are for the softest suspension of practical design that will fit within the confines of the cryostat's vacuum envelope.

3.5.1 Horizontal Performance—Separate Suspension

To a first-order approximation, the suspension is designed such that most degrees of freedom of the input motion will couple directly to the corresponding degree of freedom of the test mass and no other. Aside from the notable horizontal-pitch coupling,

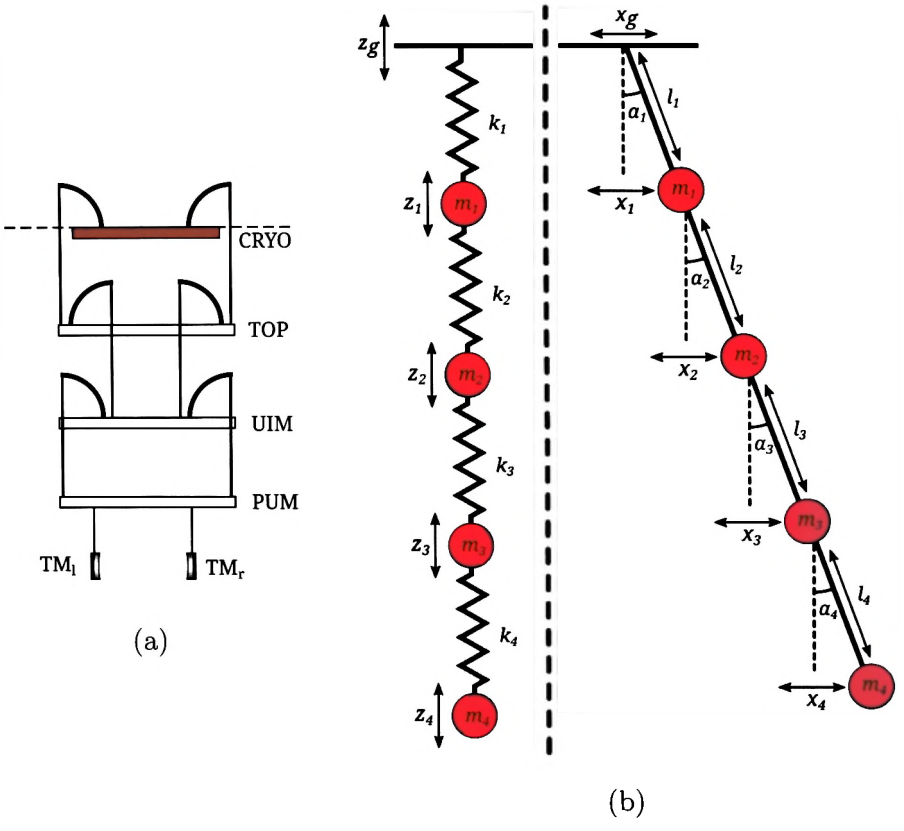


FIGURE 3.5: Panel (a) shows a sketch representation of the suspension, highlighting the individual stages and key components such as the suspension wires and blades. The sketch defines the naming convention used to refer to the individual stages, that is to say, the cryostat cold plate (CRYO), the top stage mass (TOP), the upper intermediate mass (UIM), the penultimate mass (PUM) and the two test masses (TM). Panel (b) shows how the suspension design is abstracted into simplified models for Lagrangian analysis. These models are further separated into the model used for the analysis of vertical input motion on the left and the model used for horizontal input motion on the right. The figure also defines the parameters and coordinates used in the relevant equations.

other cross-couplings are either non-linear or will couple through mismatches and misalignment of the real suspension. We initially consider only the key direct couplings, with the troublesome cross-couplings addressed in a later section.

As motivated at the start of Section 3.2, and shown in Fig. 3.5a, we begin by considering a four-stage pendulum-based suspension to achieve the necessary level of

horizontal seismic isolation. In the simplified treatment, we may consider the system as consisting of four point-mass pendulums chained together, as shown on the RHS of Fig. 3.5b. In the final analysis, we will discuss the additional effects of an extended body. The system in Fig. 3.5b can be analysed through the Lagrangian formalism laid out in Section 2.2.1. As discussed in that section, we may also decouple the consideration of horizontal-horizontal coupling from vertical-vertical coupling and construct two separate Lagrangians, which is shown by the two separate representations of the same system in Fig. 3.5b. We can define the inertial-frame coordinates of each mass through the generalised coordinates, α_i, which results in

$$x_n = x_g + \sum_{i=1}^{n} l_i \sin(\alpha_i), \tag{3.3}$$

$$y_n = y_0 - \sum_{i=1}^{n} l_i \cos(\alpha_i), \tag{3.4}$$

for all $n \in \{1, 2, 3, 4\}$.

Given this definition of the coordinate system, we can construct the Lagrangian for this system from the kinetic energy term,

$$T_{\text{pm}} = \sum_{i=1}^{4} \left(\frac{1}{2} m_i \dot{x}_i^2 + \frac{1}{2} m_i \dot{y}_i^2 \right), \tag{3.5}$$

and the potential energy term,

$$V_{\text{pm}} = \sum_{i=1}^{4} m_i g y_i. \tag{3.6}$$

The subscript 'pm' refers to the quantities in the point-mass approximation. The resulting Lagrangian yields the following set of EL equations

$$l_1 m_{1,4}\ddot{\alpha}_1 + l_2 m_{2,4}\ddot{\alpha}_2 + l_3 m_{3,4}\ddot{\alpha}_3 + l_4 m_4 \ddot{\alpha}_4 + m_{1,4} g \alpha_1 + m_{1,4}\ddot{x}_g = 0, \tag{3.7}$$

$$l_1 m_{2,4}\ddot{\alpha}_1 + l_2 m_{2,4}\ddot{\alpha}_2 + l_3 m_{3,4}\ddot{\alpha}_3 + l_4 m_4 \ddot{\alpha}_4 + m_{2,4} g \alpha_2 + m_{2,4}\ddot{x}_g = 0, \tag{3.8}$$

$$l_1 m_{3,4}\ddot{\alpha}_1 + l_2 m_{3,4}\ddot{\alpha}_2 + l_3 m_{3,4}\ddot{\alpha}_3 + l_4 m_4 \ddot{\alpha}_4 + m_{3,4} g \alpha_3 + m_{3,4}\ddot{x}_g = 0, \tag{3.9}$$

$$l_1 m_4 \ddot{\alpha}_1 + l_2 m_4 \ddot{\alpha}_2 + l_3 m_4 \ddot{\alpha}_3 + l_4 m_4 \ddot{\alpha}_4 + m_4 g \alpha_4 + m_4 \ddot{x}_g = 0, \tag{3.10}$$

where we define

$$m_{k,l} = \sum_{i=k}^{l} m_i \tag{3.11}$$

for convenience.

Through a sequence of algebraic eliminations, we can decouple the second order derivatives from each other to construct a set of four differential equations, one for each of the second derivatives of the state variables,

$$l_1 m_1 \ddot{\alpha}_1 = -m_{1,4} g \alpha_1 + m_{2,4} g \alpha_2 - m_1 \ddot{x}_g, \tag{3.12}$$

$$l_2 m_2 \ddot{\alpha}_2 = m_2 g \left(1 + \frac{m_{2,4}}{m_1}\right) \alpha_1 - m_{2,4} g \left(1 + \frac{m_2}{m_1}\right) \alpha_2 + m_{3,4} g \alpha_3, \tag{3.13}$$

$$l_3 m_3 \ddot{\alpha}_3 = m_3 g \left(1 + \frac{m_{3,4}}{m_2}\right) \alpha_2 - m_{3,4} g \left(1 + \frac{m_3}{m_2}\right) \alpha_3 + m_4 g \alpha_4, \tag{3.14}$$

$$l_4 m_4 \ddot{\alpha}_4 = m_4 g \left(1 + \frac{m_4}{m_3}\right) \alpha_3 - m_4 g \left(1 + \frac{m_4}{m_3}\right) \alpha_4. \tag{3.15}$$

Due to the cumbersome nature of solving such a system in the analytic transfer function formalism, we will instead construct a state space model in the form laid out in Eq. 2.16 and 2.17. The coupling of the ground displacement in Eq. 3.12 through its second derivative is rather inconvenient, as it is preferable to couple the ground

displacement directly. We can apply a change in our internal state variables in the form

$$\alpha'_n = \begin{cases} \alpha_1 + \frac{x_g}{l_1}, & n = 1 \\ \alpha_n, & \text{otherwise} \end{cases} \tag{3.16}$$

which replaces the ground acceleration with ground displacement in the equations of motion. Due to the second order nature of Eq. 3.12–3.15, we must introduce a set of intermediate variables, corresponding to the states' angular velocities, given by $\zeta'_i = \dot{\alpha}'_i$. The state variables are stored in the state vector, $\boldsymbol{x}_h$, such that

$$\boldsymbol{x}_h = \begin{pmatrix} \alpha'_1 & \alpha'_2 & \alpha'_3 & \alpha'_4 & \zeta'_1 & \zeta'_2 & \zeta'_3 & \zeta'_4 \end{pmatrix}^T. \tag{3.17}$$

The vector of external inputs simplifies to $\boldsymbol{uh} = x_g$ for this single input system. From Eq. 3.12–3.15, we can define the state space matrices, $\boldsymbol{A}_h$ and $\boldsymbol{B}_h$, as

$$\boldsymbol{A}_h = \left(\begin{array}{c|c} \mathbf{0} & \boldsymbol{I} \\ \hline \boldsymbol{M}_h & \mathbf{0} \end{array}\right), \tag{3.18}$$

and

$$\boldsymbol{B}_h = \begin{pmatrix} 0 & 0 & 0 & 0 & \frac{m_{1,4}g}{m_1 {l_1}^2} & -\frac{m_{1,4}g}{m_1 l_1 l_2} & 0 & 0 \end{pmatrix}^T, \tag{3.19}$$

where $\boldsymbol{I}$ and $\mathbf{0}$ is the 4×4 identity matrix and matrix of zeros, respectively, and

$$\boldsymbol{M}_h = \begin{pmatrix} -\frac{m_{1,4}g}{m_1 l_1} & \frac{m_{2,4}g}{m_1 l_1} & 0 & 0 \\ \frac{m_{1,4}g}{m_1 l_2} & -\frac{m_{1,2}m_{2,4}g}{m_1 m_2 l_2} & \frac{m_{3,4}g}{m_2 l_2} & 0 \\ 0 & \frac{m_{2,4}g}{m_2 l_3} & -\frac{m_{2,3}m_{3,4}g}{m_2 m_3 l_3} & \frac{m_4 g}{m_3 l_3} \\ 0 & 0 & \frac{m_{3,4}g}{m_3 l_4} & -\frac{m_{3,4}g}{m_3 l_4} \end{pmatrix} \tag{3.20}$$

is the matrix constructed from the four differential equations. The output vector can, in general, contain the inertial coordinates of each mass, but to simplify the analysis to a single-input-single-output system, we only consider the horizontal displacement of the test mass, such that $\boldsymbol{y}_h = x_4 = x_g + l_1\alpha_1 + l_2\alpha_2 + l_3\alpha_3 + l_4\alpha_4$ using a linear expansion of Eq. 3.3. We must restate this equation in terms of the new state variables (Eq. 3.16), which yields the remaining state space matrices,

$$\boldsymbol{C}_h = \begin{pmatrix} l_1 & l_2 & l_3 & l_4 & 0 & 0 & 0 & 0 \end{pmatrix}, \tag{3.21}$$

$$\boldsymbol{D}_h = 0. \tag{3.22}$$

In the point-mass approximation, the only parameters that can be tuned are the lengths and masses of each pendulum stage. These parameters must be optimised such that we achieve the largest suppression above the resonances. This is generally achieved by lowering all resonant frequencies of the system as much as possible (see Eq. 2.31). However, it is not possible to bring each resonance to its lowest value simultaneously because, in general, changing a parameter such that it lowers one resonance will simultaneously lead to an increase in another resonance. We can discuss this problem in a more tractable way by considering a two-stage pendulum suspension instead of the significantly more complicated four-stage system under investigation. Fig. 3.6 shows the change in the two resonances of such a system by varying the mass ratio (Fig. 3.6a) and length ratio (Fig. 3.6b) of the two stages.

From these results, some general inferences can be made. In Fig. 3.6a it is clear that the lower mass should not be substantially heavier than the upper mass, else the upper resonance increases too much. This can be argued intuitively; if the upper mass is negligible, the system begins to act like a single pendulum, with the upper resonance

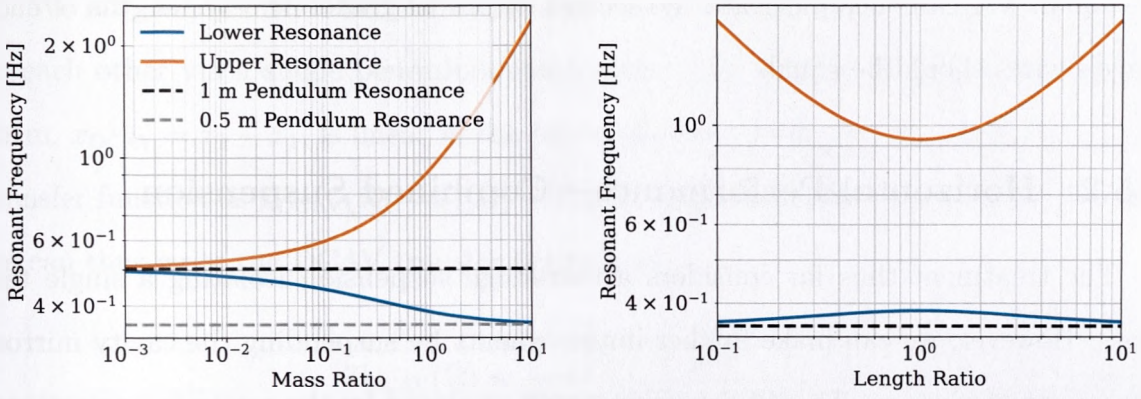


FIGURE 3.6: Changes in the resonant frequencies of a two-stage pendulum-based suspension. The legend is common to both plots. The dashed lines represent the resonances of a single pendulum of the respective length. Panel (a) shows the effect of changing the ratio of the lower mass to the upper mass. Panel (b) shows the effect of changing the ratio of the lower suspension length to the upper suspension length.

diverging towards infinity. In contrast, if the upper stage increases in mass relative to the lower stage, the influence of the lower mass becomes negligible, and thus the upper-mass resonance asymptotes towards that of a single-stage pendulum. The lower resonance varies by a mere factor of $\sqrt{2}$ between its extremal values. It may seem ideal to progressively decrease the mass of subsequent stages. However, the coupling strength of motion from the lower mass to the upper mass decreases as the upper mass becomes relatively heavier, which causes an issue with damping efficiency (see discussion in Section 3.5.5). From these arguments we determine the first empirical guideline: the stages should all be of approximately equal mass.

In Fig. 3.6b, we have a much more straightforward result. Given that we are limited by the ~70 cm internal height of the cryostat's vacuum chamber, we place a hard constraint on the total height of the suspension. In such a case, the lower resonance is largely invariant and remains close to the resonance of a single 70 cm long pendulum. The upper resonance is very clearly optimised when the lengths of the two pendulums

are equal. We can thus motivate the second empirical guideline: the lengths of each stage's wires should be equal.

3.5.2 Horizontal Performance—Combined Suspension

The treatment thus far considers a four-stage suspension isolating a single test mass. However, we can make further improvements by suspending the cavity mirrors from a common stage. The displacement signal produced by the cavity is proportional to the cavity length, that is to say, a differential measurement (a property that is discussed in Section 1.2). This degree of freedom, referred to as differential cavity motion (DCAV), is denoted by x_{DCAV}. This means that we can, in theory, tolerate high levels of common motion, which will not couple to the cavity readout. One way to achieve this straightforwardly is with a fixed cavity akin to a conventional tabletop experiment. We cannot adopt this approach directly, as this would defeat one of the primary requirements, which is to provide an effective 'free mass' for the radiation pressure interaction in the detection band. This means that the test masses must be individually suspended through at least one stage. As a consequence, we consider the design already shown in Fig. 3.5a, which consists of three intermediate, common stages with the two test masses (the final stage) individually suspended from the same penultimate stage. Another way to achieve an analogous effect is by utilising a suspension point interferometer [114], which could be used to lock the penultimate masses of two completely separate test mass suspensions. This would be particularly useful if the test masses were sufficiently separated that the two test penultimate stages could not be connected with a single bar. In our case, the suspension is sufficiently compact that a shared penultimate stage is a simpler solution.

We may analyse the effect of a coherent drive of two test masses suspended near to each other via a single pendulum stage each. The relative displacement between them, $x_{\mathrm{DCAV}} = x_2 - x_1$, is linear in the two quantities. Thus, we may write the DCAV transfer function as $T_{\mathrm{DCAV}}(\Omega) = T_2(\Omega) - T_1(\Omega)$. Using the solution derived in Eq. 2.26, we can thus write the DCAV transfer function as

$$T_{\mathrm{DCAV}}(\Omega) = \frac{(l_2 - l_1)g\Omega^2}{(g - l_1\Omega^2)(g - l_2\Omega^2)}. \tag{3.23}$$

The two key results here are that (1) the isolation at high frequencies remains proportional to Ω^{-2} as in the case of a single mass, and (2) the common mode rejection leads to a low-frequency dependence proportional to Ω^2. In this case, there is suppression of the input motion away from the resonances in both directions in frequency. The transfer function gain is given by the difference between the two pendulum lengths, which leads to the intuitive result that if the two suspensions are identical, then all the motion is common and the transfer function is zero.

For the suggested four-stage suspension, the updated Lagrangian analysis can largely retain the same approach as before, except that Eq. 3.3 and 3.4 must now contain a further coordinate for the second test mass. This addition then cascades through the analysis to add a further set of terms to the Lagrangian and ultimately increases the dimensions of the state space matrices by a further pair of state variables. In the case of identical test-mass stages, the DCAV transfer function reduces to zero, as motivated by Eq. 3.23. Although it is our plan to design the test masses and their suspensions to be identical, in reality, the test mass stages cannot be perfectly identical due to manufacturing tolerances, material inhomogeneity and human error in assembly. Therefore, the transfer function remains finite with only limited common mode suppression. With yet more degrees of freedom, an analytic derivation of the

transfer functions becomes intractably complicated. Therefore, all further extensions to the system in Section 3.5.1 are solved computationally.

3.5.3 Vertical Performance

Vertical ground motion directly couples into the common vertical motion of the cavity mirrors. The response of the cavity to vertical translation can be related to the concept of beam jitter, which manifests due to a misalignment or mismatch between the input beam and the cavity mode. The cavity is conventionally tuned to be resonant for the fundamental Gaussian mode within its mode basis. However, the electric field's spatial profile for a generic beam can be represented in terms of a summation over a series of multiple higher order modes, which are in themselves independent solutions to the electric-field wave equation [115]. The interaction between the input beam and a misaligned cavity can be analysed as the scattering of the input fundamental mode into higher order modes [116, 117]. The resonant condition is not met simultaneously for all spatial modes and thus a cavity resonantly tuned for the fundamental Gaussian mode will reject the newly scattered higher order mode components of the electric field. This will lead to a reduction in the circulating cavity power, which is shown explicitly in Eq. 5 of [118] to be a second order effect in the tilt of the cavity mirrors. A vertical displacement of curved cavity mirrors can be equated to a mirror tilt.

These fluctuations will have an effect on the stability and sensitivity of the cavity, which is yet to be seen. However, due to the second order nature of this effect, we anticipate that its suppression will not be as crucial to the success of the experiment as the suppression of the direct, linearly coupled noise sources. In spite of this argument, the resulting fluctuations in the circulating power are a clearly undesirable effect, which

would benefit from some level of suppression. We may achieve this by reducing the vertical-vertical coupling through the suspension.

In the idealised model, the vertical ground motion will only affect the common vertical motion of the cavity mirrors. However, in the real system, we anticipate a level of vertical-horizontal coupling that will occur through some non-ideal mechanism or asymmetry created during manufacturing, assembly and alignment. We anticipate that the dominant source of cross-coupling will arise from the imperfect angular positioning of the cavity mirrors, which will shift the cavity axis away from the horizontal. In such a case, the DCAV degree of freedom will be given by a linear combination of the test mass horizontal and vertical motion. As we expect the angular shift of the cavity axis to be small, we may consider the cavity as purely horizontal, with the vertical-horizontal coupling given by the appropriate geometric constant multiplied by the vertical-vertical transfer function. Although this geometric pre-factor is significantly less than unity, an issue arises if there is no direct vertical isolation, which is the case for a pure pendulum-based suspension. This means that at high frequencies, the horizontal-horizontal coupled motion is strongly suppressed and the vertical-horizontal cross-coupled motion begins to dominate the residual seismic noise. To reduce the cross-coupled motion, we must suppress the vertical-vertical coupled motion directly, which provides yet more motivation for a dedicated vertical suspension.

To improve the suppression of vertical motion, we consider the installation of blade springs, which will be used as a buffer between a particular stage and the launch point of the wires to the next stage. Such springs are commonly used for vertical isolation in the high performance suspensions of GW detectors [42, 119]. The blades we consider are thin sheets of uniform thickness with a triangular profile. These blades, when arranged to lie flat in the horizontal plane, bend in a circular arc and act as cantilever springs.

For small displacements, the spring's performance can be abstracted into the linear treatment laid out in Section 2.2.1, which only requires knowledge of the effective spring constant, k. Using standard beam theory, the spring constant of a cantilever can be derived to be [120, 121]

$$k_b = \frac{Yah^3}{4\alpha {l_b}^3}, \tag{3.24}$$

where Y is the blade material's Young's modulus, a is the base width, h is the blade thickness, l_b is the blade length, and α is a geometric factor, which differs depending on the blade profile [122]. This coefficient has a value of 1.5 for a triangular blade. The reduction of the blade stiffness usually leads to a very thin blade in relation to its other dimensions. The high performance blade springs on GW detectors are manufactured from high-strength materials operated close to the yield limit. Thus, when designing such a spring, it is essential to compute the stress in the blade. The peak stress is given by [122]

$$\sigma = \frac{6Pl_b}{ah^2}, \tag{3.25}$$

where P is the load applied to the blade tip ($P = mg$ for m equal to the total mass suspended from the blade). The unique property of a triangular blade is that the stress along the blade remains constant along its length. This is due to the fact that the cross-sectional area increases linearly with distance from the blade tip, as does the applied moment, resulting in a constant stress.

To maximise the level of vertical isolation, we want to reduce the spring constant of each blade. The spring constant is a function of the blade's geometry only and thus the design can be optimised in isolation. The design of the blade spring departs slightly from the idealised triangular design and follows a very similar strategy to the blades deployed on contemporary GW detectors and is shown in Fig. 3.7. The practical details

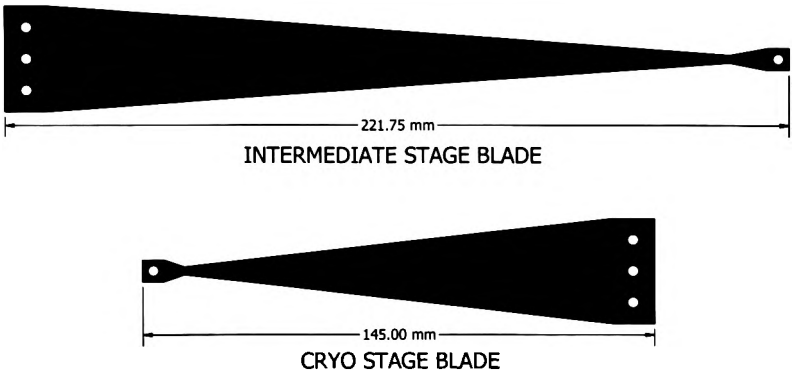


FIGURE 3.7: Top view of the two different profiles of blade spring. The intermediate stage blades are designed to span almost the full diameter of the cryostat vacuum chamber, whilst the CRYO stage blades are designed to be arranged in a triangular configuration. Further details of the design are found in Appendix C.1.

of the design are discussed in Appendix C.1.

Although the blade design could be optimised without reference to the suspension, we now turn to the modelling of the vertical isolation performance, which requires embedding the optimised blades within the broader suspension design. Firstly, we must consider how many vertical isolation stages are required. We once again refer to the *Advanced LIGO* suspension design strategy, which features three vertical stages in contrast to the four horizontal stages. This is a good strategy for us, as we should avoid placing the blades on the final stage if possible. As outlined in Section 2.3.1, the thermal noise is more sensitive to losses located closer to the test masses in the suspension chain. It is thus not desirable to have potentially lossy blades on the final stage, particularly as the vertical restoring force is purely elastic and thus does not benefit from the dilution of loss in the same way as the pendulum. The omission of a stage is compensated for by the smaller cross-coupling coefficient, which has an effect of reducing the coupled vertical motion at 100 Hz approximately akin to an extra isolation stage. We, therefore, arrange the vertical isolation on each of the uppermost three

stages and leave the test mass stage without a further vertical stage.

Due to the finite stiffness of the wires themselves, they also provide vertical isolation above their resonances. It is highly preferable to operate the wires below their yield limit, allowing us to derive the vertical spring constant of the wires in the elastic regime (see Appendix C.2) as $k_{w,z} = YA_w/l$. The wires alone are generally too stiff to act as effective vertical isolation, hence the need for dedicated spring stages. If we neglect the mass of the blades, wires and the blade-wire clamping mechanism, we can approximate that the blade and wire form a combined spring with a vertical constant given by $1/k_v = 1/k_b + 1/k_{w,z} \approx 1/k_b$. Only in the test mass stage, which contains no blades, is the wire stiffness of importance.

The Lagrangian analysis of the three-stage vertical isolation follows an identical procedure to that laid out in Section 3.5.1. It is, however, significantly easier to derive given the nature of the system. The vertical isolation can be represented by the model on the LHS of Fig. 3.5b with the top three springs corresponding to the blade spring stages and the final spring given only by the final stage wire stiffness. The Lagrangian of the system is given by

$$L_v = \sum_{i=1}^{4} \frac{1}{2} m_i \dot{z}_i^2 - \sum_{i=1}^{4} \frac{1}{2} k_i (z_i - z_{i-1})^2, \tag{3.26}$$

where we define $z_0 = z_g$ as the input vertical ground motion. The resulting equations of motion are much simpler than in the horizontal case and no further change of variables is necessary. Therefore we may define the internal state vector as

$$\boldsymbol{x}_v = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 & v_1 & v_2 & v_3 & v_4 \end{pmatrix}^T, \tag{3.27}$$

for $v_i = \dot{z}_i$. We may also define $\boldsymbol{u}_v = z_g$ and $\boldsymbol{y}_v = z_3$. This leads to the resulting set of state space matrices,

$$\boldsymbol{A}_v = \left(\begin{array}{c|c} \mathbf{0} & \boldsymbol{I} \\ \hline \boldsymbol{M}_v & \mathbf{0} \end{array}\right), \tag{3.28}$$

$$\boldsymbol{M}_v = \begin{pmatrix} -\frac{k_1+k_2}{m_1} & \frac{k_2}{m_1} & 0 & 0 \\ \frac{k_2}{m_2} & -\frac{k_2+k_3}{m_2} & \frac{k_3}{m_2} & 0 \\ 0 & \frac{k_3}{m_3} & -\frac{k_3+k_4}{m_3} & \frac{k_4}{m_3} \\ 0 & 0 & \frac{k_4}{m_4} & -\frac{k_4}{m_4} \end{pmatrix}, \tag{3.29}$$

$$\boldsymbol{B}_v = \begin{pmatrix} 0 & 0 & 0 & 0 & \frac{k_1}{m_1} & 0 & 0 & 0 \end{pmatrix}^T, \tag{3.30}$$

$$\boldsymbol{C}_v = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}, \tag{3.31}$$

$$\boldsymbol{D}_v = 0. \tag{3.32}$$

The vertical-vertical transfer function that is generated using this model is implemented in Section 3.6.1 to estimate the vertical-horizontal cross-coupling.

3.5.4 Ancillary Degrees of Freedom

So far, we have considered the linear degrees of freedom. However, it is worthwhile to briefly consider some further cross-couplings and the angular degrees of freedom as well. The first consideration is an amendment to the existing model of the horizontal isolation stages. The point-mass model is effective at determining the horizontal-horizontal transfer function. However, certain couplings can only be modelled by considering the stages as extended bodies.

A particular effect to consider is horizontal-pitch coupling, that is to say the transfer function from x_g to Ry. This arises if the attachment point of test masses is above their centre of mass, which induces a torque around the centre of mass. The extended Lagrangian analysis involves adding a further degree of freedom to the test mass suspension to allow for pivoting around the lower test mass attachment point. This, of course, increases the complexity of the process even further, but the procedure follows the same logical progression described in Section 3.5.1. The key additional effect of an extended body on the Lagrangian is an additional term in the kinetic energy,

$$T_{\mathrm{eb}} = T_{\mathrm{pm}} + \sum_{i=1}^{K} \frac{1}{2} I_i \dot{\beta}_i^2, \tag{3.33}$$

for each stage capable of rotation around its centre of mass, where I_i is the moment of inertia around the rotation axis passing through the centre of mass and β_i is the pitch angle of mass i (this should not be confused with the angle α_i, which is the inclination from vertical of each suspension wire and retains a meaning consistent with its use in Section 3.5.1). Depending on the stiffness of the wire at the lower attachment point to a given stage, this may have a negligible effect on the horizontal-horizontal transfer function (soft connection) or may lead to a shift in the resonant frequency (rigid connection). The latter effect is caused by a higher effective inertia to linear displacement, as a linear centre of mass displacement will be accompanied by a rotation as well.

Depending on the strength of the horizontal-tilt coupling (growing with separation between the attachment point and centre of mass), the tilt resonances may couple significantly into the horizontal-horizontal transfer function and thus appear in the cavity readout. In terms of the direct effect of mirror tilt on the cavity response, this

generates fluctuations in the position and angle of the cavity axis, which leads to the same second order beam jitter discussed in Section 3.5.3 and thus the same arguments apply.

There is also the potential to induce pitch motion in the intermediate masses through direct pitch-pitch coupling. The finite stiffness of the wire material and, more importantly, the softening introduced by the blade springs, means that relative tilts can occur between the intermediate stages, as the stages are not rigidly confined to the plane defined by their DC position. With motion in this degree of freedom, and generally for all ancillary degrees of freedom, the motion should not directly couple into the DCAV degree of freedom. It is, however, likely that the resulting resonances may show up in the cavity readout through some non-deal couplings and, in any case, fluctuations in these degrees of freedom will affect the cavity circulating power through beam jitter.

In general, we do not consider these degrees of freedom as thoroughly as the more directly coupled ones. However, it is at the least beneficial to be aware of the likely location of the resonances and consider tuning these in advance to avoid any resonant peaks appearing within the detection band. A similar state space model that was produced for the horizontal and vertical degrees of freedom can be constructed from the Lagrangian that governs the energy in each ancillary degree of freedom. Such a model is used in Section 3.6.1 to estimate the locations of the upper stage pitch and yaw resonances.

3.5.5 Damping

A final consideration is the damping strategy. The suspension in-vacuum will experience negligible damping, with much of the natural damping arising from loss (see

Section 2.2.3) in the suspension materials. This is likely to lead to a rather large Q for most resonances, particularly for the pendulum modes that will further benefit from the dilution of loss due to energy stored in the gravitational potential.

The motion around resonances often contributes significantly to the RMS motion. In the case of weak damping, the RMS motion may become excessively high due, as the RMS contribution from a particular resonance is scales positively with Q. High RMS motion in the DCAV degree of freedom means that there is a stricter requirement on the cavity locking servo gain to maintain operation. This is generally undesirable because (a) the amplitude of the motion may be so large as to bring the servo out of range, (b) the servo bandwidth may not be sufficiently high to catch the error signal to acquire lock, and (c) operating the corresponding actuators at a higher range may lead to a larger injection of actuation noise. Maintaining low RMS motion in the ancillary degrees of freedom is also important, as high RMS motion will lead to a high RMS fluctuation in the circulating cavity power. If allowed to fluctuate too much, even this second-order effect may severely degrade controllability.

Therefore, we must consider a dedicated damping solution to reduce the RMS motion generated by all suspension resonances. An active system has benefits over a passive one, as elaborated in Section 2.2.4. However, the added challenge of designing and operating cryogenic and vacuum-compatible active components leads us to consider conventional, passive strategies. Our strategy is derived from the approach demonstrated within the GW community (see Section 2.2.3), using a set of eddy-current dampers to reduce the Q of all resonances.

A key aspect to settle is the location of the damping. Consider a generic two-stage suspension for simplicity, which we wish to passively damp. Eddy currents lead to the damping of relative motion. For the two-stage configuration, we can consider damping

the relative velocities between (a) the ground and the upper mass, (b) the upper mass and the lower mass, and (c) the ground and the lower mass.

Fig. 3.8 shows the effect of the three different damping locations on the transfer function. The clearest result from Fig. 3.8 is that damping should occur between two consecutive isolation stages, in this case, either configuration (a) or (b). Configuration (c) couples ground velocity directly into the motion of the lower mass, effectively bypassing the isolation provided by the upper stage entirely. Although the damping in both configurations (a) and (b) exhibits the same low- and high-frequency behaviour, the efficiency at which each resonance is damped differs between the two: type (a) appears to damp the lower resonance more effectively than type (b) and vice-versa for the upper resonance.

In our suspension design, the highest resonance can be identified as the pendulum resonance of the final test mass stages, which are significantly lighter and shorter than the intermediate stages. Due to this difference, the motion of the test masses couples weakly into the motion of the penultimate stage, that is to say, the test masses respond strongly to drive from the penultimate mass above, but the penultimate mass responds weakly to drive from the test masses below. In this case, it will be rather more challenging to dampen the test mass resonances via damping introduced at the upper stages. However, we do not wish to insert damping directly at the test masses, as damping introduces loss, which is in contravention of the suspension thermal noise requirements that losses should be minimised at the final stage. As a compromise, we have designed a damping solution between Stage 2 and Stage 3 (UIM and PUM), that is to say, one stage up the suspension chain from the test masses.

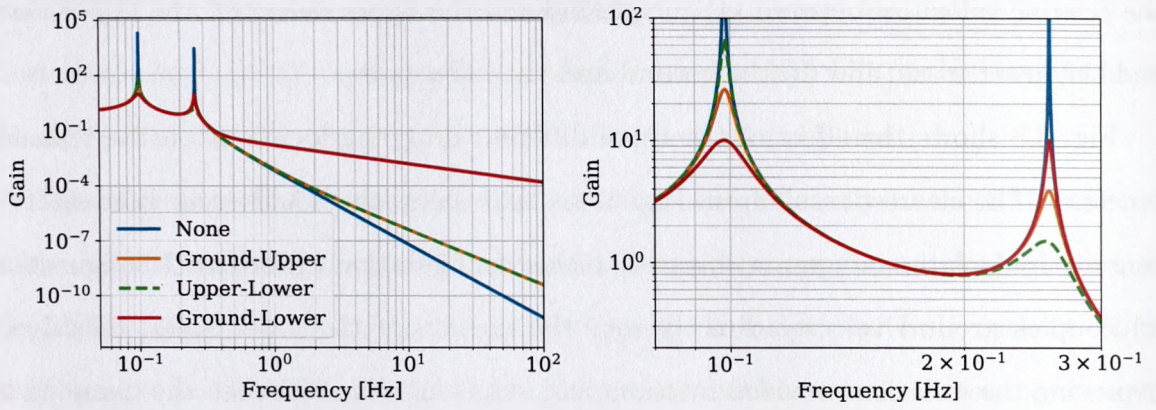


FIGURE 3.8: Comparison of differently located damping implemented on the same suspension. The ground-upper, upper-lower and ground-lower plots refer to configurations (a)–(c) respectively. The RHS plot is a zoomed view around the resonances shown in the LHS plot.

3.6 Practical Design

Using the models, constraints and guidelines discussed in the previous sections, we were able to design a suspension that should meet our requirements. The overall design is shown in Fig. 3.9. A more detailed view of each stage and major component, including the assembly procedure is shown in Appendix D.2. This design is for a room-temperature prototype capable of reaching the milestone of measuring the quantum back-action. There are multiple challenges in designing a suspension that will work well in cryogenic conditions, some of which include accounting for the contraction of the components during the cooling process and maintaining cavity alignment at the operating point when the alignment was performed at room temperature. These considerations will be left for a future upgrade, however, once we gather the necessary experience from operating the prototype design.

It is worthwhile to qualitatively comment on some of the aspects of the design in Fig. 3.9. Due to space constraints on the suspension stages, we opted for pre-curved

blades. These blades are curved when unloaded with the correspondingly negative curvature profile to that experienced by flat blades under load. When the appropriate load is applied to the curved blades, the curvature generated by the bending of the blades compensates for the blades' natural curvature, leading to blades lying flat under load. The clearest benefits of such a design choice are that (a) the stages are more compact without having to accommodate the height of curved blades, and (b) we avoid significant horizontal-pitch coupling by placing the suspension point of the wires closer to the centre of mass of each blade's host stage.

The ring-shaped stage masses may also have a slightly unconventional appearance, particularly in comparison to the bulky stage masses on contemporary GW detectors. The purpose of these is two-fold. To optimise the design of the blade springs, we must maximise their length. To achieve this, the blades must span close to the full diameter of the inner volume of the vacuum chamber. This means the stages must be appropriately wide to carry such blades. Secondly, we wish to match the mass of the intermediate stages to the test mass stages as closely as possible for efficient damping as mentioned in Section 3.5.5. On account of the very low mass of the mirrors, we aim to minimise the mass of each stage whilst keeping the outer diameter of each stage as large as the space permits. These constraints are satisfied by concentrating the mass in a ring-like structure as shown in Fig. 3.9.

The final aspect of the design to comment on is the test mass assembly, that is to say the lower test mass clamping, the mirror and its holder. The test mass assembly in the prototype design is designed to be easy to remove and disassemble. We will make use of a custom, metal optic holder so that no permanent changes need to be made to the mirror itself and the mirror can be secured in the holder with a removable screw. This design decision was made for convenience in the initial phase of the experiment, while

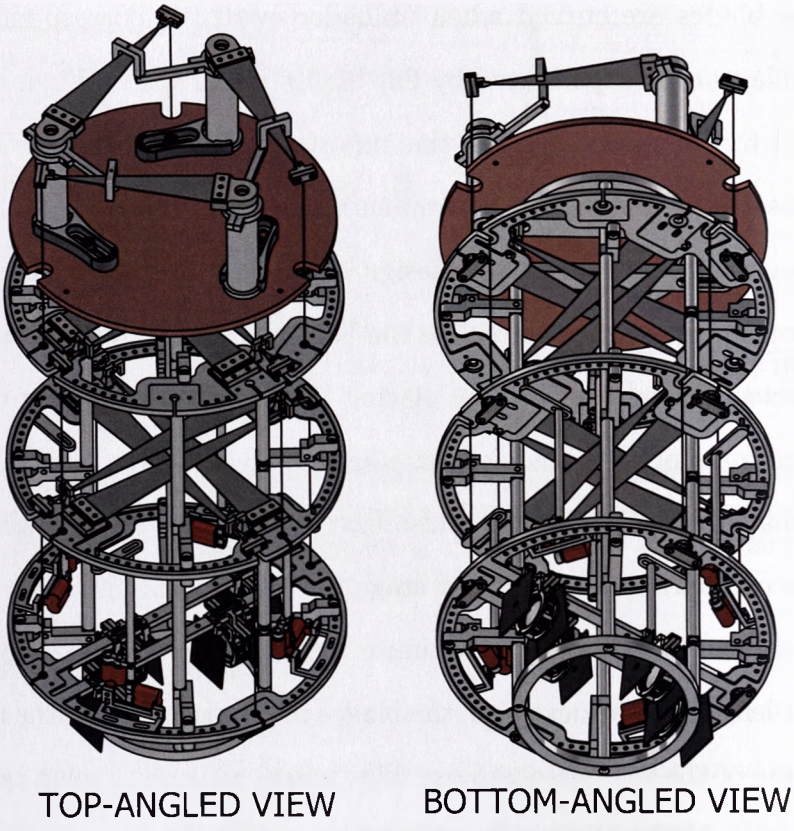


FIGURE 3.9: Design drawing of the room-temperature prototype suspension from a top- and bottom-angled viewpoint. This design incorporates all of the aspects discussed in this chapter, with four stages of horizontal isolation and three stages of vertical isolation. The horizontal stages consist of the three ring-like, suspended masses and the final, test mass stage for each cavity mirror. A more detailed breakdown of the suspension design including the assembly procedure of key components is shown in Appendix D.2.

we gather experience with cavity locking and balancing of the suspension. Redesigns of the test mass suspension mechanism will be necessary in the future, particularly due to thermal noise considerations. The topic of thermal noise reduction in the future test mass designs is revisited in Section 4.1.4.

3.6.1 Design Parameters and Transfer Functions

The prototype design's performance can be estimated using state space models that are in essence slight extensions of the models laid out earlier in this chapter. Table 3.1 lists the nominal values of the parameters used in these models. The estimated transfer functions for the corresponding set of parameters are shown in Fig. 3.10.

The transfer functions shown in Fig. 3.10 represent the primary means of coupling seismic noise into the DCAV readout channel. The horizontal-horizontal transfer function was derived from an extension of the model in Section 3.5.1, which also includes the common-mode suppression effect discussed in Section 3.5.2. The vertical-horizontal cross-coupling was estimated from the vertical-vertical transfer function derived from the model in Section 3.5.3. The dominant source of this cross-coupling in our case is expected to be due to a tilt of the cavity mode axis away from the horizontal. The cross-coupling is modelled as a combination of the vertical-vertical transfer function, multiplied by the cross-coupling coefficient (a geometric pre-factor given approximately by the cavity axis tilt angle), and with additional common-mode suppression included similarly to that derived for the horizontal coupling.

The cross-coupling mechanism of the ancillary degrees of freedom was not investigated, as we do not expect them to dominate across any broad region of the seismic noise spectrum. However, we model the location of the resonances of the tilt and yaw degrees of freedom in Fig. 3.10 as it is possible that we may see the sharp resonant features rise above the broadband spectra of the more strongly coupled degrees of freedom.

TABLE 3.1: Parameters of the prototype suspension

Parameter	Nominal Value
TOP suspension wire number, n_1	3
TOP suspension wire length, L_1	190 mm
TOP suspension wire radius*, r_1	200 µm
TOP mass, m_1	475 g
TOP pitch moment of inertia, I_1	3.04×10^{-3} kg m^2
TOP vertical spring constant, k_1	91.3 N m^{-1}
UIM suspension wire number, n_2	4
UIM suspension wire length, L_2	190 mm
UIM suspension wire radius*, r_2	200 µm
UIM mass, m_2	423 g
UIM pitch moment of inertia, I_2	2.53×10^{-3} kg m^2
UIM vertical spring constant, k_2	17.0 N m^{-1}
PUM suspension wire number, n_3	4
PUM suspension wire length, L_3	190 mm
PUM suspension wire radius*, r_3	200 µm
PUM mass, m_3	586 g
PUM pitch moment of inertia, I_3	3.52×10^{-3} kg m^2
PUM vertical spring constant, k_3	7.87 N m^{-1}
Number of test masses, n_{TM}	4
TM suspension wire number, n_4	2
TM suspension wire length, L_4	20 mm
TM suspension wire radius*, r_4	100 µm
TM mass, m_4	12 g
TM pitch moment of inertia, I_4	5.76×10^{-7} kg m^2
TM vertical spring constant, k_4	3.14×10^{6} N m^{-1}

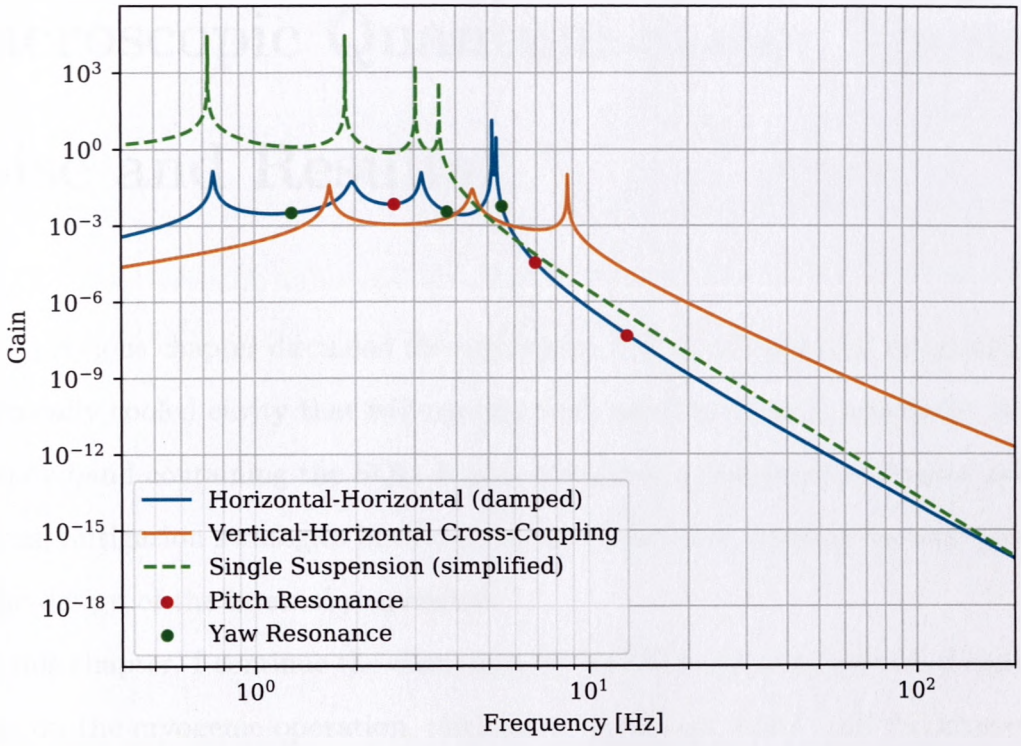


FIGURE 3.10: Key transfer functions into the DCAV readout channel. The single-suspension transfer function shows the result of the simplified model in Section 3.5.1 with no damping or elastic restoring force present. The horizontal-horizontal transfer function shows all effects combined. The ancillary degrees of freedom have the locations of their resonances represented.

Chapter 4

Macroscopic Quantum State: Thermal Noise and Results

The previous chapter discussed the motivation behind the proposal for a suspended, cryogenically cooled cavity that will operate with quantum-limited sensitivity over the frequency band containing the SQL. It also contained a discussion of seismic isolation concerns, mitigation strategies both outside and inside the cryostat vacuum envelope, and the design of the internal suspension.

In this chapter, I continue the discussion of the table-top experimental design, with a focus on the cryogenic operation, mitigation of thermal noise, and the impact that cryogenic operation has and will continue to have on the suspension design. I conclude the chapter, and the discussion of the experiment, with the status of the experiment to date and present notable results obtained thus far.

4.1 Thermal Considerations

The cryostat was mentioned in the previous chapter due to its central role in the constraints placed upon the suspension design. However, the cryostat's main purpose

is to enable cryogenic operation of the cavity. As outlined in Section 3.1.1, thermal noise (particularly the coating thermal noise) poses a significant obstacle to measuring broadband quantum noise over the SQL frequency. This naturally motivates our extensive efforts toward reducing it. In a general sense, all sources of thermal noise regardless of origin are amplified for larger amounts of stored thermal energy. This is one of the conclusions drawn from the *fluctuation-dissipation theorem* and is reflected in the direct proportionality of the noise PSD with temperature in Eq. 2.38. Thus cooling the system is the most straightforward means of reducing thermal noise.

However, many of the thermal and mechanical properties of the materials have themselves a dependence on temperature. For example, the thermal noise is generally related to some loss mechanism within the system, which is often a strongly temperature-dependent quantity. It is thus essential to address the key properties of our selected materials and their performance at different temperatures. It is preferable to select materials whose properties improve with lower temperature or, at the very least, avoid materials whose properties significantly worsen.

4.1.1 Cooling Efficiency

The first consideration is reaching the temperature operating point of the test masses. This allows for a discussion of the cooling strategy in a broad sense but also has significant consequences on the material choices in several areas of the suspension design.

The cooling rate, as well as the final temperature, will be influenced by the two primary heat exchange channels, these being radiative and conductive cooling. The suspension will radiatively exchange heat with the walls of the cryostat lid, which becomes a significant mechanism in a vacuum chamber with no exchange gas and very

little thermal contact to facilitate conductive cooling. The radiative power from a body of temperature, T, and surface area, A_s, is given by the well-known Stefan-Boltzmann law:

$$P_{\text{rad}} = \epsilon \sigma A_s T^4, \tag{4.1}$$

where σ is the Stefan-Boltzmann constant, and ϵ is the emissivity of the material. By approximating the surface of the suspension to be around $0.5\,\text{m}^2$, it can be irradiated with heating power on the order of 10 W, should the suspension be fully exposed to the room-temperature black-body radiation of the cryostat walls. Such a large amount of energy would overwhelm the cooling power of the cryocooler (see Fig. 4.1) and the suspension would never reach cryogenic operation. To remedy this, we make use of a two-stage cryocooler system, with a 50 K heat shield that can absorb the room temperature radiation from the outer vacuum lid, leaving a residual radiative heat load of around 10 mW on the enclosed suspension. The inner heat shield is cooled by the cryocooler's first stage, which can tolerate a much higher heat load. Therefore, the inner shield will cool significantly faster, meaning that during the initial cooling process, there is a net radiative outflow of heat from the suspension. This is a very effective cooling channel initially, but drops in efficiency rapidly due to the fourth-order scaling of Eq. 4.1 with temperature, which means we cannot rely on this cooling process below around 100 K.

The second mechanism to consider is conductive cooling. This is the only means of cooling the suspension below 50 K when the suspension reaches thermal equilibrium with the inner heat shield. The maximum rate of cooling of the suspension components depends on the maximum rate of heat extraction at the cryocooler's second stage. The estimated cooling power of the 10 K stage over temperature is given by Fig. 4.1, which shows the inferred rate of heat extraction over a single cool-down cycle of a sample plate

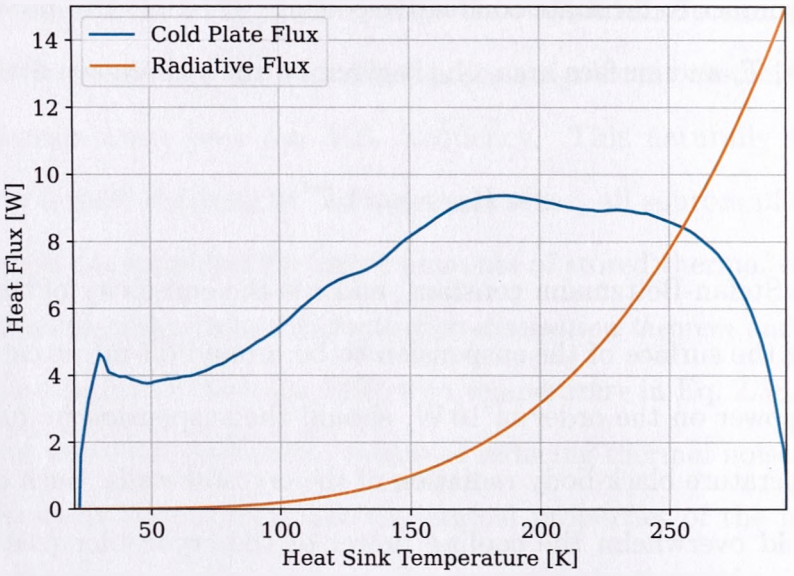


FIGURE 4.1: Heat extracted at the 10 K plate inferred from temperature measurements of the cold plate and sample plate connected via copper rods. The heat extraction power does not change significantly over with temperature except at the two extremes of the temperature range. The temperature is swept over from high to low, which means that in the initial stage the cryocooler pump was not operating at maximum power. Therefore, the drop in heat load at high temperature is due to transient behaviour. The low-temperature fluctuations represent noise due to insufficient data resolution in this region. According to the manufacturer specification, the heat extraction should begin to decline at 20 K, reaching zero at around 10 K.

connected to the cold plate by four cylindrical rods. From these results, we can clearly see that, even for a component whose cooling is not restricted by thermal contact with the cold plate, radiative cooling is a significant, if not dominant, cooling channel at high temperatures.

The real efficiency of conductive cooling is, however, limited by the heat flux across the whole suspension and can vary significantly according to the thermal conductivity of the material and the heat-flow cross-section. Determining the rate of conductive cooling precisely is not a trivial matter for a suspension design that does not follow simple geometry. However, the heat flow rate will be limited most critically by the

small cross-sectional area of the suspension wires, which restrict the heat flow between the stage masses. We can thus approximate that the thermal masses will always be internally in a state of thermal equilibrium and the heat flow is driven by a thermal gradient along the wire length. For wires of constant cross-section, A_w, and length, L_w, the heating power is given by:

$$P_{\text{con}} = \frac{\kappa(T) A_w \Delta T}{L_w}, \tag{4.2}$$

where κ is the thermal conductivity of the wire material and ΔT is the temperature difference between the two ends of the wire. The result in Eq. 4.2 assumes a constant thermal gradient in the wire, which relies on a constant κ over the temperature intervals. However, the thermal conductivity of materials can vary significantly over a small temperature range at cryogenic temperatures, which leads to a non-linear thermal gradient across the wire length and thus only an approximate nature to Eq. 4.2.

4.1.2 Upper Suspension

The design of the upper suspension stages is largely driven by the seismic isolation requirements, meaning that we must often compromise on some of the thermal considerations, which often suggest the opposite approach. Whether in the room-temperature or cryogenic phase, the suspension parameters are fixed to the values given in Table 3.1. However, there remains some freedom to make design choices regarding thermal considerations as well, particularly concerning the materials to be used for each component. The choices are generally driven by the efficiency of the suspension cooling, as the upper suspension stages are unlikely to contribute significantly to the suspension thermal noise. The three main areas of interest are the material composition of (a) the stages

masses, (b) the vertically isolating blade springs, and (c) the suspension wires.

The stage masses will be manufactured from aluminium for both modes of operation, which is chosen for its sufficient strength, durability and low density. The mass of the intermediate stages must be kept very low in the interests of effective test mass damping (discussed in Section 3.5.5) for which the low density of aluminium is highly beneficial. The current stages are formed of a standard aluminium alloy for ease of manufacture and handling. In the future upgrade, we consider replacing the stage material with high-purity aluminium. Aluminium with a very low density of impurities exhibits some excellent thermal properties at cryogenic temperatures, particularly, high thermal conductivity (can reach around $10^4\,\mathrm{W\,m^{-1}\,K^{-1}}$ at 20 K [123]). This will be very helpful for the reduction of the cooling time. A significant challenge for this future adoption is the low strength and hardness of the material (up to 100 times weaker than the strongest aluminium alloys). This makes it both challenging to machine and use in a larger assembly without damaging the components. As a result, prior to the potential implementation of high-purity aluminium in the design, we will first gain experience operating with the more robust alloy.

For vertical blade springs in the initial investigation, we will use 316L grade stainless steel. This material does not exhibit particularly high yield strength (typically below 300 MPa [124]) and thus we cannot be particularly ambitious with reducing the stiffness of the blades. This grade of steel is sufficiently strong for initial investigations and is easier to machine. However, we will eventually adopt alternative materials that have a better level of performance. For the room temperature design, we may consider grades of maraging steel with high yield strength, such as 440C stainless steel with yield stress as high as 1900 MPa upon correct heat treatment [125]. The issue with continuing use of stainless steel into the cryogenic phase of the experiment is in their

poor thermal conductivity (around $2\,\mathrm{W\,m^{-1}\,K^{-1}}$ [126] with slight variation depending on the grade), which would cause significant reduction in the rate of conductive cooling of the suspension. Thus, in the cryogenic phase, we will look to materials with better cryogenic thermal properties, such as beryllium copper, with thermal conductivity of around $10\,\mathrm{W\,m^{-1}\,K^{-1}}$ [126], or tungsten, with thermal conductivity approaching $10^4\,\mathrm{W\,m^{-1}\,K^{-1}}$ at 20 K [127].

The upper stage wires are not likely to pose much of a problem with the widely available 316 stainless steel wires being both sufficiently strong and sufficiently thin to not affect the suspension performance greatly. The initial assembly presented in Section 4.4 uses particularly thick wires, which led to the significant shift in the pendulum resonances seen in Fig. 3.10. However, the wire thickness can be substantially reduced whilst maintaining a sufficient safety factor for operation. The wire material will have to be changed for cryogenic operation due to the same issue of low thermal conductivity of stainless steel. In this case, we may consider tungsten wires or beryllium copper ribbons as an alternative.

Even with the choice of highly conductive materials, the restriction of heat flow up the suspension chain may lead to an unreasonably long cooling time. If we consider tungsten wires operating at a factor of two below their yield stress, using Eq. 4.2, we can estimate the maximum heat flux between the PUM and the UIM to be around 2 mW[1]. We may relax the seismic isolation requirements and allow for thicker wires to increase the cooling efficiency. With an increase of the wire thickness to give an equal distribution of energy between the gravitational potential and elastic potential, the maximum heat flux rises to 0.8 W but comes at a cost of a factor of 2 in seismic isolation at 100 Hz in an equivalent single-stage suspension. The thinnest wires are

[1] assuming the UIM is at 10 K and taking the peak heat flux in the PUM temperature range 20 K to 100 K

insufficient to even extract the heat load from the black-body radiation of the inner shield. Whilst the thicker wires should be sufficient to maintain the test masses at the operating temperature, they are still likely too thin to facilitate a reasonable cool-down time. As such, we may consider implementing some form of retractable connection between the stage masses and the cryocooler that can aid in the initial cooling procedure and be removed after the correct temperature is reached.

4.1.3 Substrate Choice

The careful choice of substrate material is an essential part of our low thermal noise strategy. Fused silica is a natural choice due to its small loss angle of 10^{-7} at room temperature [128]. Indeed, this material choice is appropriate for the room-temperature milestones of our experiment and is a very common choice for similar experiments in the field. We will make use of standard-sized table-top optics for the initial fused silica substrates of one-inch diameter and one-quarter-inch thickness.

However, the mechanical loss of fused silica degrades significantly at cryogenic temperatures, particularly over a broad peak that rises to 10^{-3} close to our 20 K operating point [129]. The larger loss angle degrades sensitivity due to the substantial increase in the substrate thermal noise (see Eq. 2.45). As such, we cannot continue the use of this material into the cryogenic phase.

The *KAGRA* [15] detector is designed for cryogenic operation and makes use of sapphire substrates, whilst *ET-LF* [130] considers crystalline silicon substrates. Both of these substrates are a reasonable choice with a mechanical loss of around 10^{-8} at 20 K [123]. Given the choice of the two suitable candidates, we will consider silicon in the low-noise, cryogenic phase of the experiment, due to its other properties of interest. Particularly, silicon possesses two roots in its coefficient of thermal expansion

around 20 K and 120 K [131], which leads to a substantial reduction in the substrate thermoelastic noise [75]. From a cooling perspective, the switch to silicon is also highly beneficial, with the thermal conductivity of silicon exceeding $10^3\,\mathrm{W\,m^{-1}\,K^{-1}}$ in the 20 K to 50 K range compared to the rather low thermal conductivity of silica [123], which declines from around $1\,\mathrm{W\,m^{-1}\,K^{-1}}$ at room temperature down to $0.1\,\mathrm{W\,m^{-1}\,K^{-1}}$ at 20 K.

4.1.4 Test Mass Suspension

The substrates will be held inside a metal holder via a screw and the final-stage suspension wires will be clamped to the metal holder. This process allows for easy interchange with standard off-the-shelf mirrors and the ability to re-suspend and dismantle the entire test mass assembly. This is beneficial in the early prototyping and commissioning phase as we may make changes to the design based on gathered experience. This design also avoids making permanent alterations to the mirrors or bonding components directly to the substrate. The initial test mass assembly (shown in detail in Appendix D.2) shares some similarities with the design of the optic holder in the small optic suspensions of *Advanced LIGO* [132].

A significant limitation of this attachment strategy is the large number of bolted connections close to the test mass. This is predicted to contribute significantly to the suspension thermal noise and possibly severely limit the sensitivity in the detection band. As an intermediate step, we plan to loop the suspension wires around the substrate in a sling-like configuration akin to the *Initial LIGO* test mass suspensions [133]. In this case, we anticipate that a level of customisation of the substrate shape will be necessary to implement the standoffs needed to minimise the contact area between the wire sling and the substrate.

The final upgrade to the suspension strategy will coincide with the switch to silicon substrates. At this point, the test mass assembly will need to exhibit significantly lower thermal noise than the intermediate designs. Therefore, it is at this stage that we will design a monolithic test mass suspension with silicon fibres replacing the metal wires and direct bonding applied to the substrate. The exact details of the future designs have not been finalised at the date of writing but are certain to share similarities in design with the monolithic fibre suspension of the *Advanced LIGO* test masses [65].

4.1.5 Coating Choice

The coating thermal noise is, in some sense, the most difficult challenge to overcome, given its limitation on the sensitivity of conventional GW detectors. In the initial stages of the experiment, we plan to use commercial Ta_2O_5/SiO_2 coatings, which are widely used and well understood within the GW community. With a loss angle of 1×10^{-4} to 4×10^{-4} [77], they are sufficient for the room-temperature milestones of the experiment. However, these coatings lose performance at 20 K, with a loss angle that rises up to 9×10^{-4} [78].

In the switch to silicon substrates, we have the option to change the coating material as well. In this case, we may consider several options that show promise at the cryogenic operating temperature. These options include amorphous coatings, such as a-Si[2]/SiO_2 (loss angles of 2×10^{-5} to 5×10^{-5} [79] and 1×10^{-5} to 4×10^{-5} [134], respectively), a-Si/SiN (loss angle of 1×10^{-5} for SiN [135]), or crystalline coatings, such as GaAs/AlGaAs (loss angle of 2.5×10^{-5} to 5×10^{-5} [136, 137]). The exact choice of coating will largely depend on the progress made in coating development over the near future. A particular point of concern that will have to be resolved first is

[2]amorphous silicon

the comparatively large absorption in the current state-of-the-art coatings. We may ultimately continue to employ Ta_2O_5/SiO_2 coatings in the long term to mitigate the high risk of committing to prototype coatings. However, in such a case it is uncertain whether it will be possible for us to reach the SQL level of sensitivity. Due to the level of uncertainty concerning the ongoing coating strategy, we leave this open to future investigations.

4.2 Laser and Cavity

The eventual transition to silicon substrates justified in Section 4.1.3 leads to an important change in our choice of laser wavelength. Silicon is often used in conjunction with laser light at 1550 nm, due to the low absorption of silicon at this wavelength, which is in contrast to its high absorption at the 1064 nm wavelength commonly used in contemporary GW detectors. The wavelength of 1550 nm is suitable for use with both silicon and fused silica substrates, hence we adopt its use from the beginning.

The high sensitivity requirements on the displacement measurement lead to a significant issue with the laser frequency noise. In order to suppress the laser frequency noise we lock the laser to the cavity using a PDH locking scheme as laid out in Section 1.2.5. However, this leads to the need, as discussed briefly in Section 2.4.1, for an external frequency reference in order to make an out-of-loop measurement of the displacement fluctuations and get a real-terms suppression of the laser frequency noise. We achieve thus by including a second cavity, that is as identical to the first and locked to its own laser in the same way. In this case, the transmitted signals experience phase fluctuations according to the DCAV motion of each cavity, free of laser frequency noise.

These signals can be interfered to generate a beat signal within the RF band. With

TABLE 4.1: Cavity parameters of the initial room-temperature design.

Parameter	Nominal Value
Wavelength	1550 nm
Transmissivity	8 ppm
Radius of Curvature	500 mm
Cavity Length	90 mm
Free Spectral Range	1.67 GHz
Finesse	3.93×10^5
Cavity g-factor	0.67

an appropriate shifting of the two laser frequencies by integer multiples of each cavity's FSR, it is possible to bring the beat note to 10 MHz to 100 MHz, and thus within the range of high-bandwidth electronics. Demodulating the signal around the beat note yields a measure of the combined DCAV motion, whose PSD is approximately double that of the individual cavities, given the cavities are near-identical and the noise is uncorrelated.

The cavities will be formed by two identical curved mirrors at a tunable separation. The initial design consists of identical fused silica mirrors with high reflectivity to achieve a very high finesse and thus a high build-up of circulating power. This will be of use in particular during the room-temperature measurement, where the high circulating power will generate a high radiation pressure on the substrates and thus amplify the QRPN. The parameters of the initial design are summarised in Table 4.1.

The switch to the cryogenic operation of silicon substrates will move us along the path toward operating at the quantum-noise-limited sensitivity at the level of the SQL. At this point, the substrate and coating may be significantly different from the coating used in the room-temperature prototype. Therefore, some of the parameter values given in Table 4.1 may change significantly in the future.

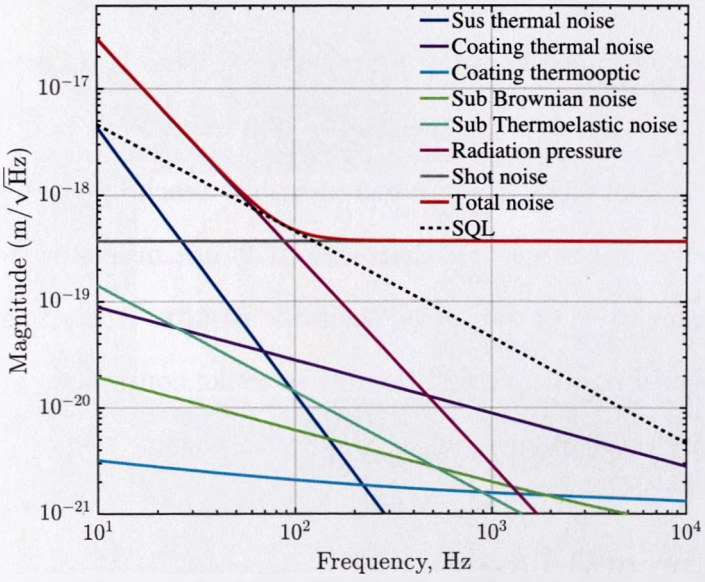


FIGURE 4.2: Noise budget for the final, cryogenic design capable of reaching the SQL and enabling investigations of macroscopic quantum mechanics.

4.3 Final Design Noise Budget

We take the optimistic case where we are able to use low-loss coatings at their advertised level of performance on silicon substrates. For operating in the low-noise cryogenic regime, and with the laser power tuned to place the SQL frequency at 100 Hz, we predict that we will be able to create a quantum limited system over the SQL region according to our requirements. The noise budget for this low-noise regime is shown in Fig. 4.2.

4.4 Status and Results—Room Temperature

The key design decisions, simulated performance, and parameters of the suspension and cavity were discussed in some detail over the course of Chapter 3 and the previous sections of this chapter. In this section, we explore the current status of the experiment

and present results of interest. As laid out in the timeline in Fig. 3.1, the initial stages of the experiment take place at room-temperature. This is necessary from an early commissioning perspective, as it is essential to gain experience in operating the cavity during more forgiving conditions before incrementally reaching the design sensitivity in the low-noise, cryogenic regime. Furthermore, it is not necessary to reach cryogenic operation in order for us to obtain novel scientific results. According to our predicted experimental goals in Fig. 3.1, we will be able to make continuous progress for several years into the future prior to engaging in the first cryogenic tests.

4.4.1 Experimental Layout

The nominal experimental design, as elucidated in Section 4.2, consists of two independent cavities suspended from the same PUM. The cavities are pumped by two separate lasers with identical control schemes. Fig. 4.3 shows the configuration of the in-air table-top components in relation to the two suspended cavities. At the moment, the cavities are operated completely independently with no beating of the transmitted fields. Thus, we cannot yet obtain the low-noise displacement spectrum from the combined cavity readout.

The locking electronics in Fig. 4.3 are abstracted into a single component. Therefore, Fig. 4.4 shows an expanded view of the electronic components of the control scheme. The individual components are further discussed in the following sections.

The experimental layout can be further appreciated with reference to Fig. 4.5, which shows the optical layout, locking electronics, and the cryostat with its support frame at the time the measurements in this section were taken. The suspension design presented in Fig. 3.9 can be seen fully assembled in Fig. 4.6.

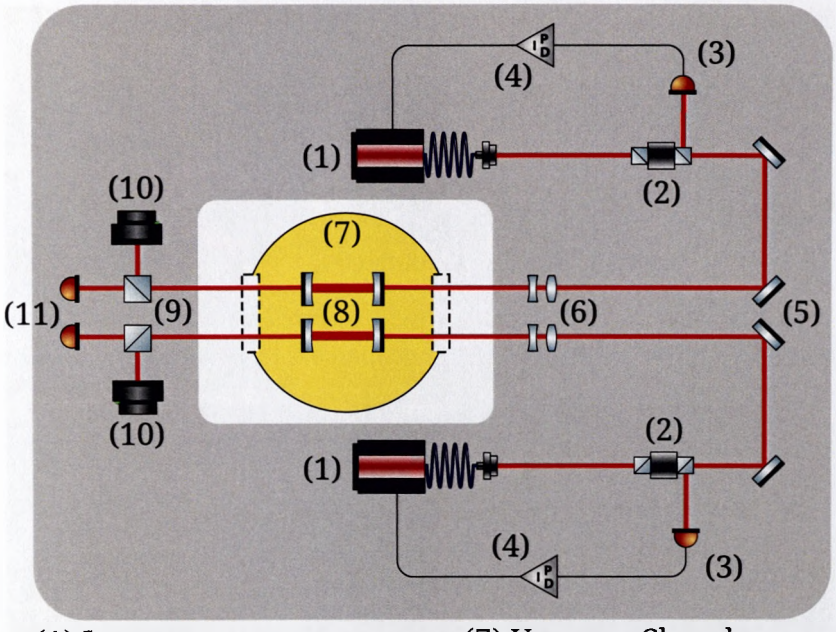


FIGURE 4.3: Optical layout of the experiment for the room-temperature prototype in the configuration discussed throughout Section 4.4. The two cavities are controlled independently with two separate but identical control systems. The cavity displacement extraction via the beating of the two cavities is not yet present in this layout.

4.4.2 Room-Temperature Noise Budget

The noise budget for the room-temperature prototype is shown in Fig. 4.7. It can be seen that with the expected level of performance of the suspension and mirrors, we should be capable of reaching a QRPN dominated level of sensitivity over a frequency band around 100 Hz. The radiation pressure effect is deliberately amplified, which raises the SQL above 1 kHz. To bring the SQL frequency to the target region around 100 Hz, it will be necessary to reduce the laser power and thus lower the quantum noise in this frequency band. Given the budget in Fig. 4.7, it can be seen that the suspension

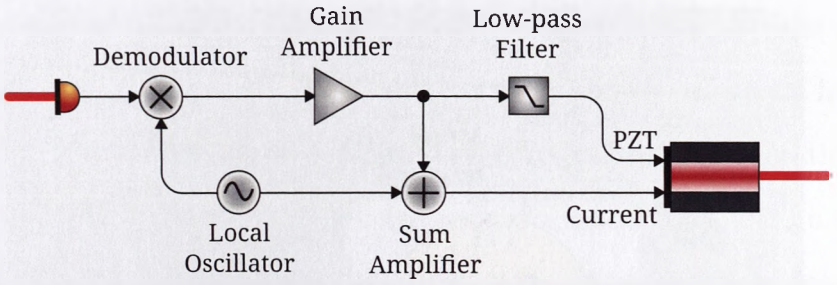


FIGURE 4.4: Detailed layout of the locking electronics, shown as (4) in Fig. 4.3. Discussion of the parameters of the components can be found in Section 4.4.5.

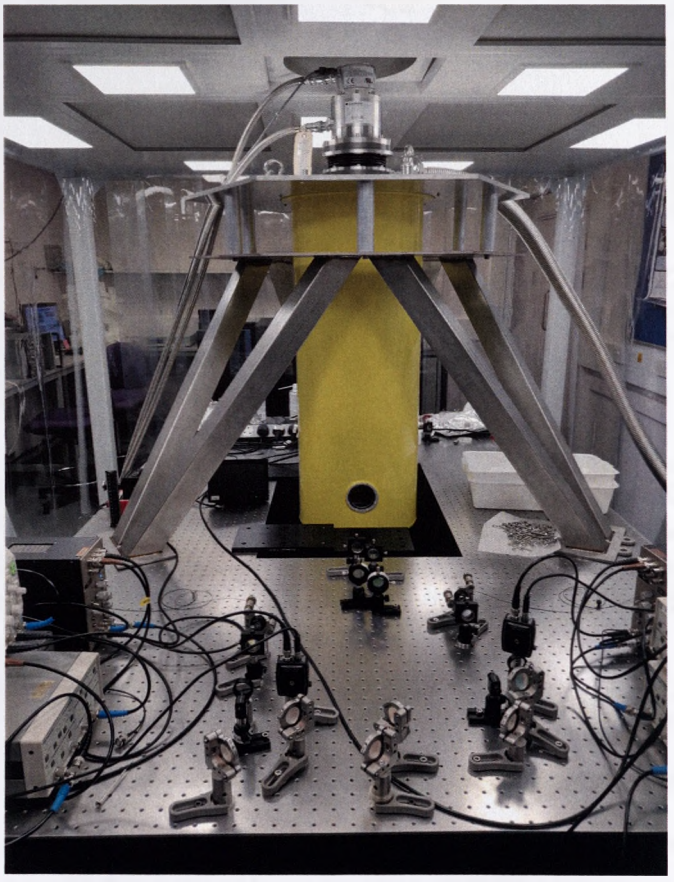


FIGURE 4.5: Photo of the experimental set-up at the time of the measurements laid out in Section 4.4. The photo shows the yellow cryostat on its tetrahedral support frame in the passive configuration, as shown in Fig. 3.4. The optical components in the lower portion of the photo correspond to the components laid out on the RHS of the vacuum chamber in Fig. 4.3. The electronics in both lower corners of the photo make up the locking servo laid out in Fig. 4.4.

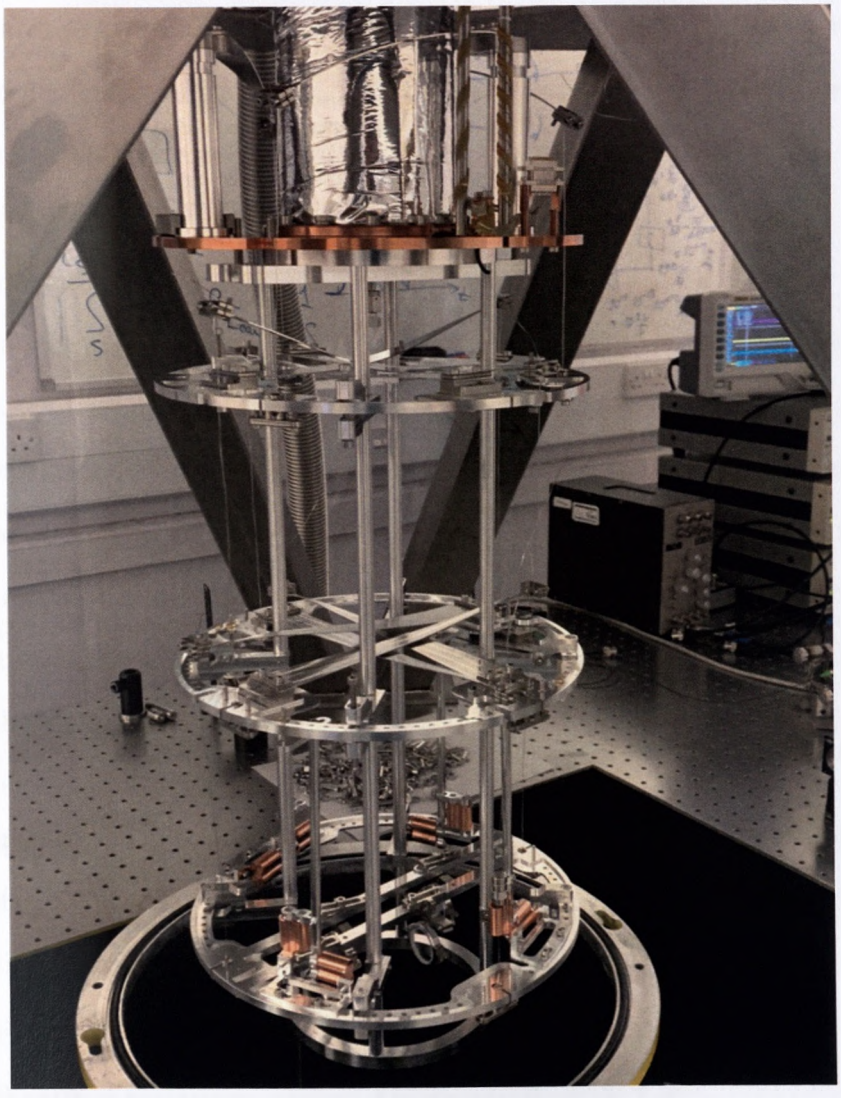


FIGURE 4.6: Photo of the suspension installed in its operating position within the vacuum chamber interior. The suspension is assembled as shown in Fig. 3.9, with the three ring-like masses clearly identifiable and one of the test masses visible just below the PUM. The blade springs are shown in the upwards pre-curved configuration under varying levels of load before fine-tuning the balance.

thermal noise and coating thermal noise will then begin to dominate. We can thus see directly, the necessity of finding a solution to the thermal noise problem that forms so much of the discussion of Section 4.1.

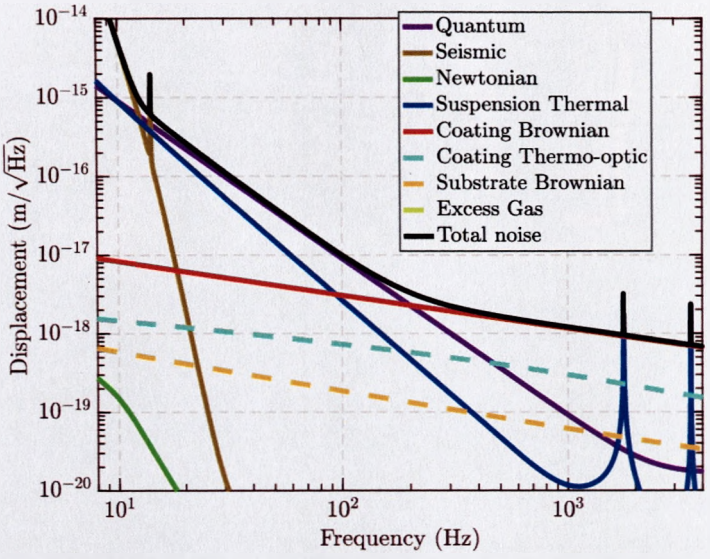


FIGURE 4.7: Noise budget for the room-temperature prototype of the experiment capable of operating at a QRPN-dominated level of sensitivity.

4.4.3 Laser Frequency—Actuation and Noise

The cavity resonance is maintained through a PDH locking scheme similar to that described in Section 1.2.5. We choose to avoid any in-vacuum actuation on the suspended cavity, leaving it effectively free-swinging in all degrees of freedom at all times. Motion in all degrees of freedom is filtered through the seismically isolating suspension and the eddy-current damping scheme lowers the Q of the suspension resonances to reduce their contribution to the RMS motion. The residual fluctuations in the cavity length are compensated for by actuation on the laser frequency.

The laser frequency of our lasers can be controlled through a combination of laser cavity PZT drive and laser current drive. The PZT can be driven with a high response from DC up to the piezo resonance, at which point the response declines. The current, on the other hand, can be driven to much higher frequencies but the input is high-passed, leading to very poor response at low frequencies. For this reason, we refer to these

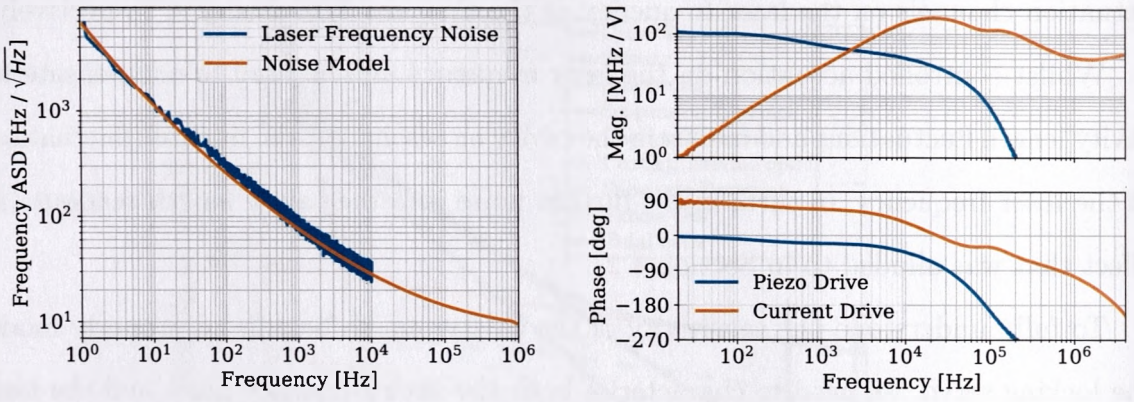
actuation channels on the laser frequency as the slow and fast channels, respectively.

Whilst controlled actuation on the laser frequency can be used to compensate for cavity length fluctuations and maintain the cavity on resonance, the intrinsic fluctuations of the laser frequency, in turn, inject further noise into the cavity length readout (an effect that was alluded to in Section 4.2).

To fully understand the sensitivity of the cavity and to be able to properly model the locking servo, we need to characterise both the laser frequency noise and the laser frequency actuation response. We measured both by beating the two near-identical lasers together, resulting in a beat-note at 15 MHz. We then implemented an electronic delay line to demodulate the beat-note signal. We make the assumption that both lasers have approximately equal noise spectra to infer the laser frequency noise of each laser from the demodulated beat-note ASD, which is shown in Fig. 4.8a. We then measured the response of the demodulated beat-note signal to a swept-sine input at the slow and fast channels. From this we computed the voltage-to-laser-frequency transfer functions for each channel. The results of the transfer function measurements are shown in Fig. 4.8b.

4.4.4 Cavity Stability and Alignment

The DC cavity alignment is achieved through relatively rough hand-tuning of the angle at the upper clamping hook of the final suspension stage. The alignment procedure is performed manually in-air, which means we have no means of compensating for any drifts that occur once the chamber is sealed. This is a potential concern for the later cryogenic operation, where the suspension wires will undergo significant shortening upon cooling down and may lead to DC distortions of the cavity alignment.



(a) (b)

FIGURE 4.8: Panel (a) shows the laser frequency noise ASD as measured by a beat measurement between the two near-identical lasers. The noise of a single laser is estimated by dividing the beat note ASD by a factor of $\sqrt{2}$, which is included in the plot. Panel (b) shows the laser response bode plot for the PZT (slow channel) and current (fast channel) drives.

For now, the tuning procedure is sufficient to allow for aligning the input beam with the cavity mode using external, in-air steering mirrors. The result is that the cavity mode does not rest at a purely horizontal orientation, but with some angular offset, which leads to the vertical-horizontal cross-coupling that was discussed in Section 3.5.3.

4.4.5 Lock Acquisition

The lock acquisition process is complicated by a combination of (a) the substantial laser frequency noise at high frequency, which demands a high-bandwidth servo; (b) the high RMS motion arising from the high Q of the final stage resonances, which strains the range of the servo; and (c) the high cavity finesse, which leads to a small cavity linewidth, and a large power buildup at resonance.

The first consideration we turn to is the servo response. The full cavity locking system consists of several components, which are shown in Fig. 4.3 and 4.4. The PDH

scheme begins with a 4 MHz modulation of the laser current, which modulates the laser frequency and, to some extent, the laser amplitude. The error signal is produced from the power measured at the PD in reflection (component (3) in Fig. 4.3), which is demodulated with the local oscillator inside the locking servo (see Fig. 4.4). The residual amplitude modulation generates an offset in the error signal, shifting the principal zero crossing away from the nominal position on resonance. This can be compensated for by an appropriate shift in the servo set-point and does not pose a significant problem. The signal passes through a gain amplifier and is then split to be passed to both control channels of the laser frequency. The slow channel (PZT) control signal is passed through a 10 Hz low-pass filter with a tunable gain, whilst the fast channel (current) has the control signal combined with the local oscillator, as the fast channel is responsible for high-frequency control and for generating the RF sidebands.

Using a model of the cavity response, together with the response of the slow and fast channels, we estimate the open loop transfer function shown in the Bode plot in Fig. 4.9. We are able to tune the unity-gain frequency to 320 kHz with a resulting phase margin of 22°. A key area of concern for the stability of the servo is at the cross-over point between the slow and fast channels. The ratio between the PZT and current drive is shown in the Bode plot in Fig. 4.10. The cross-over occurs at the 0 dB point and its position can be tuned by changing the gain on the PZT low-pass filter, which shifts the magnitude plot vertically and thus shifts the position of the 0 dB crossing. The phase remains unchanged, so we can determine the optimal location for the cross-over to achieve the largest phase margin. As shown on the plot, the phase margin is maximised at the indicated frequency of 20.5 kHz with a value of 18°. We thus tune the low-pass gain to place the cross-over point at this frequency, which is reflected in the open-loop transfer function in Fig. 4.9.

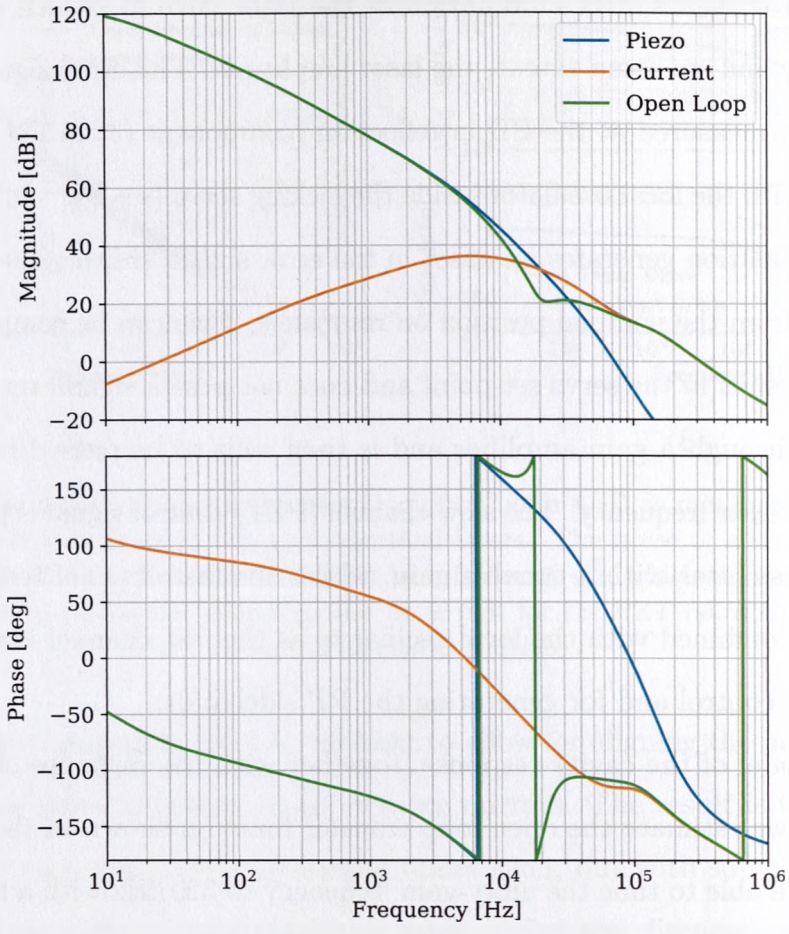


FIGURE 4.9: Bode plot of the open-loop transfer function constructed from the individual transfer functions of the PZT and current drives. The transfer functions are constructed from the cavity response, the particular responses of each channel (shown in Fig. 4.8b) and a further low-pass applied to the PZT.

The high cavity finesse also applies a constraint on the RMS velocity that can be tolerated to achieve lock. The cavity remains largely empty before acquiring lock and the electric field must build up over some finite time when the resonant condition is met. In the usual presentation of an error signal (see Fig. 1.3), it is assumed that the cavity is in the steady state. In the case of a conventional cavity sweep, it is assumed that the traversal over the cavity resonance is sufficiently slow such that the cavity has

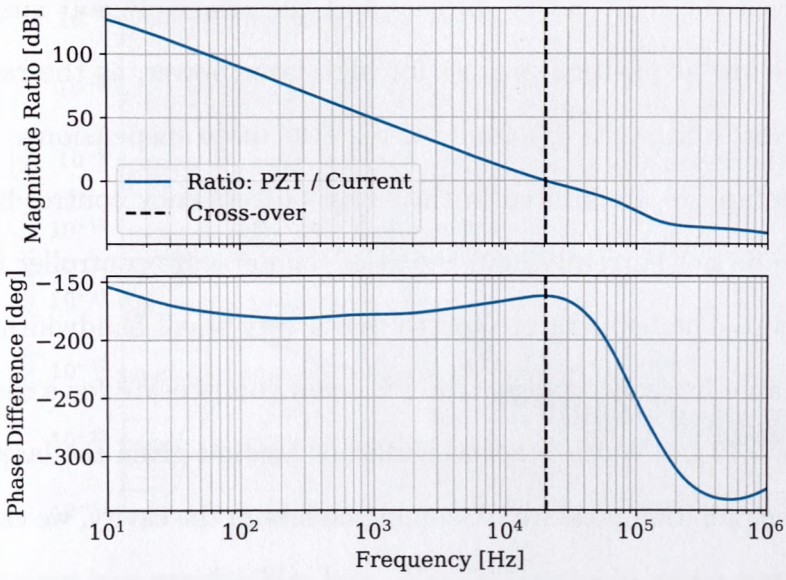


FIGURE 4.10: Bode plot of the ratio of the PZT and current transfer functions shown in Fig 4.9. The cross-over between the two transfer functions occurs at the 0 dB crossing and is indicated with the vertical dashed line. The corresponding location is indicated on the phase plot and determines the phase margin at the blend. The relative gain is optimised to place the cross-over frequency at the location of the maximum phase margin.

reached close enough to the steady state. The cavity build-up time related to the cavity linewidth, or similarly the cavity storage time, given by [138] $\tau_s \approx \mathcal{F}/f_{\mathrm{FSR}}$. For our cavity, this build-up time is on the order of 0.2 ms. The cavity motion traverses the error signal on the time scale given by $\tau_e \approx \lambda/(v_{\mathrm{RMS}}\mathcal{F})$. In order for the cavity field to build up sufficiently, $\tau_e \gg \tau_s$, which yields an upper limit on the cavity velocity,

$$v_{\mathrm{RMS}} \ll \frac{\lambda f_{\mathrm{FSR}}}{\mathcal{F}^2} \ll 17\,\mathrm{nm\,s^{-1}}. \tag{4.3}$$

This strict requirement on the velocity upper limit explains the need for strong and effective damping of the suspension resonances discussed in Section 3.5.5.

The large buildup in the cavity generates a significant amount of radiation pressure,

particularly given the light cavity mirrors and the relatively soft suspension. This complicates the low acquisition process for high laser power, as the rapid buildup of circulating power excites the dynamics of the test mass suspensions. In the current experimental setup, we are limited in the range of frequency control by the range of the PZT, as we do not currently have the laser temperature controller integrated into the feedback loop. The temperature control has a very small bandwidth but allows for tuning over a large frequency range. The PZT can compensate for a maximum cavity displacement of 0.19 µm, which is not a significant amount given the large displacement generated by the radiation pressure. Assuming no loss in the cavity, we can only tolerate a maximum input power of around 0.5 mW, and still achieve and maintain lock.

The cavity should be impedance matched with mirrors of equal transmissivity, and thus the steady-state reflected power on resonance should be zero. However, due to cavity losses, we see a non-zero level of reflected power, which allows us to estimate these cavity losses. We thus estimate losses of around 16 ppm, which results in a correction to the finesse to a new value of 1.3×10^5. We anticipate that further improvements are possible as the mode matching is not ideal, which contributes to the loss estimate as well.

4.4.6 Cavity Noise Spectrum

We can determine the spectral density of the cavity length fluctuations (the DCAV degree of freedom) from the spectrum of the control signal calibrated into units of metres. However, in its current status, the experiment operates at an insufficient sensitivity to observe quantum noise effects because of the dominant presence of laser frequency noise in this readout channel. This will be remedied in the future with the beating of the transmitted signal between the two cavities, at which point we will be able to measure

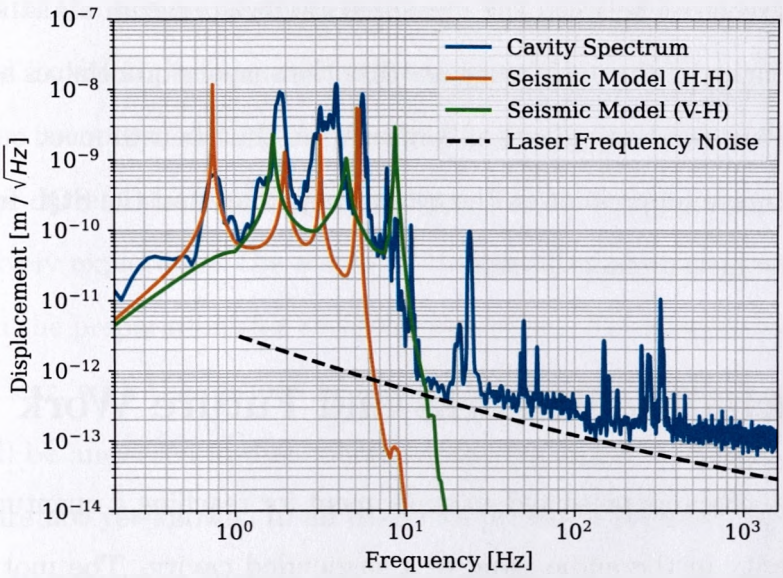


FIGURE 4.11: Cavity displacement ASD with principal noise sources modelled for comparison. The residual seismic noise spectrum is estimated from the table motion spectrum multiplied through by the directly coupled horizontal-horizontal (H-H) and the cross-coupled vertical-horizontal (V-H) transfer functions. The key result is the relatively close agreement on the slope of the pre- and post-resonance region of the seismic noise and the agreement on the location of many of the resonances. Towards the higher frequency region, the laser frequency noise begins to dominate over the cavity spectrum as expected.

the spectrum of their beat note in the absence of laser frequency noise (as discussed in Section 4.2).

Whilst many sources of true cavity displacement are obscured, the dominant nature of the seismic noise at low frequency means that we may observe the key features of the residual seismic noise coupling even in the cavity spectrum inferred from the control signal. We can use this to test the validity of the suspension model discussed in Chapter 3 by comparing the low-frequency seismic-noise-dominated region of the cavity spectrum to our seismic noise model. The model is generated by applying the transfer function generated in Fig. 3.10 to the seismic spectrum on the optical table, shown in Fig. 3.3. The result of this analysis is presented in Fig. 4.11 and shows the relatively

high level of agreement between the measured cavity spectrum and the seismic noise model. The accuracy of the model suggests that the suspension behaves according to our expectations. Thus, we have a level of confirmation that we will meet our requirements for the seismic noise suppression in the crucial region around the SQL frequency in the future beat note measurement as well.

4.5 Concluding Remarks and Future Work

The last two chapters presented an experiment for reaching a quantum-noise limited level of sensitivity in the audio band of a suspended cavity. The motivation for the experiment arises from the necessity to understand the quantum noise spectrum around the so-called standard quantum limit and to develop techniques for surpassing this sensitivity limit. This is driven primarily by the requirements of future generations of GW detectors, which will have increasingly stringent sensitivity requirements below the contemporary quantum noise level. The experiment is also embedded within the greater context of fundamental physics research on macroscopic quantum mechanics.

The key features of the experiment are concerned with the need to suppress the two dominant sources of noise in such high-performance, suspended, optical cavities, these being seismic noise and thermal noise. To the latter point, a cryostat is employed to cool the system down to 20 K, which has significant implications for the rest of the experimental design. Chapter 3 was principally concerned with the former, with an exploration of the notable departures from the norm that are brought about by the presence of the cryostat (particularly the focus on increasing the rigidity of many conventionally soft connections). The largest focus of Chapter 3, however, was on the analysis and modelling of the internal suspension, which is the most challenging and

complex aspect of the seismic isolation strategy. The chapter concludes with a model of the predicted isolation performance of the suspension.

This chapter is concerned with the thermal noise mitigation strategy, as well as a discussion of the current status of the experiment. The thermal noise considerations are not as conclusively explored as the seismic reduction strategy, principally because we are yet to begin the preparations for cryogenic operation. The avenues for thermal noise reduction are clear, with the options laid out in Section 4.1. However, which one is to be adopted will be answered in due course depending on the outcome of future work, whose results are not yet known. In all cases, we present a primary, high-risk strategy, which is likely to achieve or surpass our requirements, and an alternative, lower-risk mitigation strategy to deal with unforeseen difficulties, whose ability to achieve the requirements is, nonetheless, in doubt. The discussion concludes with the current status of the experiment and the initial results of the room-temperature prototype. In these sections, we lay out the approach to operating the cavity and compare the cavity displacement spectrum we obtained to our seismic noise model.

The experiment will naturally progress to the stage of low-noise operation, which will involve the beating together of both suspended cavities to suppress the laser frequency noise in the cavity displacement spectrum. Once we are able to extend the range of our laser frequency tuning, we will raise the laser input power and magnify the radiation pressure effect to measure the QRPN. This will achieve one of the significant milestones of the experiment.

In the more distant future, the route is less clear and will depend on external factors. However, in the near term, we expect the significant changes to be: (a) the replacement of current blade springs with high-strength maraging steel, which will lower the vertical suspension resonances and significantly reduce the cross-coupled seismic noise, (b)

the redesign of the test mass suspension stage with a departure away from a metal holder and bolted connections, and (c) explorations of the suspension performance upon initiating the cooling procedure.

Chapter 5

Optomechanical Filters for Interferometric Measurements

Quantum fluctuations of light are, in themselves, notable phenomena for study, as laid out in Chapter 3. However, much of contemporary GW instrumentation is concerned with strategies to suppress this noise source, as discussed in Section 3.1.3. One promising approach is in the use of optomechanical filter cavities that serve as quantum amplifiers. These additions to the contemporary detector design provide enhancements to the GW strain sensitivity and can be installed alongside complementary techniques such as squeezing without conflict. In this chapter we present theoretical work on such an implementation in the *Advanced LIGO* detector configuration. The novelty of the work consists of a dual approach through an analytic, frequency-domain derivation as well as a numerical, time-domain simulation. My focus will be on the time-domain simulation and the two key contributions that this approach brings, this being the verification of the results of the analytic approach with a fully independent technique, and providing insights to the behaviour of the system that is not as easy to calculate or interpret using the analytic technique.

The content of this chapter is a reformatted copy of the paper *X. Li, J. Smetana,*

A. S. Ubhi et al. "Enhancing interferometer sensitivity without sacrificing bandwidth and stability: Beyond single-mode and resolved-sideband approximation" [139]. I was considered the joint, primary coauthor with *X. Li*. *X. Li* was responsible for writing and/or editing Sections 5.1–5.3 and I was responsible for writing and/or editing Sections 5.4 and 5.5. Section 5.4 presents the outline and results from my work on the time domain simulation, which I constructed and analysed myself.

5.1 Introduction

The detection of GWs from a binary black hole merger in 2015 [16] opened a new window of astronomy observation. Binary black hole systems have so far been the most commonly measured GW sources [140, 141]. The demand for extracting richer properties of the ringdown stage [142–146], as well as other astrophysical processes that produce pronounced gravitational waves at high frequencies, e.g., the binary neutron star mergers [147–149] and core collapse supernovae [150–152], calls for the broadband and high-frequency sensitivity of GW detectors. For current advanced detectors and even future detectors including *Einstein Telescope* [153, 154] and Cosmic Explorer [155], the quantum shot noise limits the detector sensitivity from a few hundred Hz to kilo Hz [156]. Similarly, recently proposed detectors of axion-like-particles in the galactic halo suffer from the photon shot noise across their sensitivity bands [157–160].

In the canonical interferometer configuration [29], resonant arm cavities are used to increase the relative signal strength by effectively extending the length of the detector via repeated reflections of the optical field. However, the positive dispersion of the arm cavity makes the signal at higher frequencies no longer resonant. This leads to an inverse relationship between the peak sensitivity and bandwidth of the detector,

known as the Mizuno limit [31]. This can be traced back to the energetic quantum limit [24, 161], which is also called the quantum Cramer-Rao Bound (QCRB) [162], and is therefore limited by the quantum fluctuation of the intracavity light field [18].

One approach to broaden the bandwidth without sacrificing peak sensitivity is to attach a negative-dispersion optomechanical filter cavity to the arm cavity [163], which can compensate for the phase gained in the arm cavity and thus resulting in a *white light cavity* effect. However, such a scheme is dynamically unstable and thus an additional stabilizing controller must be implemented. We will call this unstable white light cavity scheme as *uWLC* for short. In the original proposal [163], in addition to the filter cavity, there are several auxiliary optics, either for impedance match with the input mirror of the arm cavity or for steering the field to the filter, which leads to a rather complex setup.

In a later study [164], it was found that converting the signal-recycling cavity (SRC) into the optomechanical filter can lead to bandwidth broadening with a much simpler optical layout, as illustrated in Fig. 5.1. The parameter regime considered [164], however, still leads to an unstable system. We recently realised that, when the optomechanical interaction strength is smaller than or equal to the coupling frequency between the arm cavity and the filter cavity, the system will be self-stablised [165]. More interestingly, the peak sensitivity is improved together with the bandwidth, not due to squeezing but a significant enhancement of the signal response. We shall call this stably operated white light cavity scheme as *sWLC* for short. So far, the sWLC scheme has only been analysed using the Hamiltonian in the single-mode and the resolved-sideband approximation [164, 165], which treats the arm cavity signal field, the mechanical oscillator, and the field in the filter cavity as single modes, separately. The stability issue of the system is based on the poles of the resulting input-output relation [165]. One

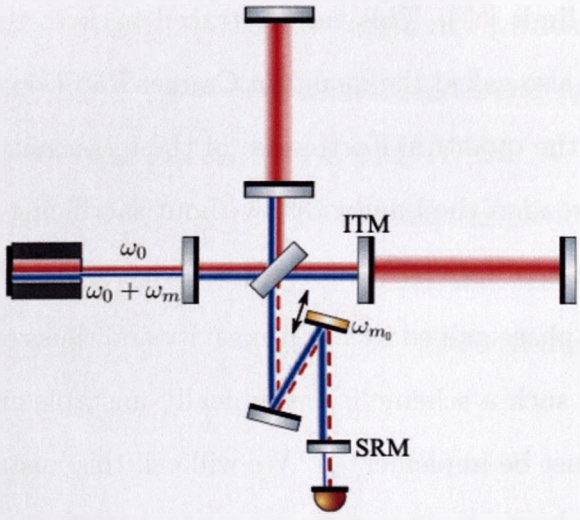


FIGURE 5.1: A schematics of the interferometer configuration with the signal-recycling cavity (SRC) as the optomechanical filter. Both the main arm cavities and the SRC resonant at ω_0, but they have different free spectral ranges due to the different cavity lengths. One of the mirrors in the SRC acts as the mechanical degree of freedom that resonant at ω_{m_0}. It couples to the signal sidebands around ω_0 via the radiation pressure due to the blue detuned SRC pumping field at $\omega_0 + \omega_m$. Note that $\omega_{m_0}/\omega_m \approx 1$, with slight difference caused by the optical spring effect due to the blue-detuned pumping (will be explained later near Eq. 5.17). ITM: input test mass, SRM: signal-recycling mirror.

natural question to ask is, whether the stability and sensitivity improvement remain valid when these approximations are removed by considering realistic parameters. Answering such a question defines the theme of this paper.

Note that in the low-frequency band (below 40 Hz), quantum backaction in the main interferometer will further restrict the sensitivity, because the radiation pressure force from the strong field would interact with the test mass. Throughout the paper, we will be focusing on the shot-noise-limited sensitivity, leaving the consideration and cancellation of back-action noise for future studies [165].

The outline of this paper goes as follows: in Section 5.2 we revisit the idealised Hamiltonian dynamics and stability requirement, and introduce the full analysis method

beyond the single-mode and resolved-sideband approximations. In Section 5.3, we solve the system dynamics in the frequency domain. We also analyse the stability using the Nyquist criterion and show the resulting sensitivity. In Section 5.4, we carry out a detailed time-domain simulation, and show the agreement with the frequency-domain analysis.

5.2 From single-mode approximation to full analysis

In this section, we will first recap the idealised Hamiltonian dynamics under the single-mode approximation as analysed in Refs. [164, 165] for the sWLC, and also revisit the stability requirement. We will then introduce the analysing framework considering the realistic setup, by abandoning the approximations applied in the previous treatment. The detailed analysis in the frequency and time domain will be carried in Section 5.3 and 5.4, respectively.

The idealised sWLC mode interaction in the rotating frame of frequency ω_0, as illustrated in Fig. 5.2, can be described by the following Hamiltonian:

$$\hat{H}_{\rm int} = i\,\hbar\omega_s(\hat{a}\,\hat{b}^\dagger - \hat{a}^\dagger\hat{b}) + i\,\hbar G(\hat{b}^\dagger\hat{c}^\dagger - \hat{b}\,\hat{c})\,, \tag{5.1}$$

with $\hat{a}\,,\hat{b}\,,\hat{c}$ being the quantum operators of the differential optical mode of the arm cavity, the SRC (i.e., the filter cavity) optical mode, and the mechanical mode, respectively. Here ω_s is the beam-splitter-type interaction strength between mode $\hat{a}$ and $\hat{b}$, and G describes the optomechanical interaction strength between mode $\hat{b}$ and mode $\hat{c}$.

Considering the GW strain signal input h with coupling strength α, as well as the coupling to the external bath $\hat{b}_{\rm in}$ with rate γ [166], we obtain the Heisenberg equations

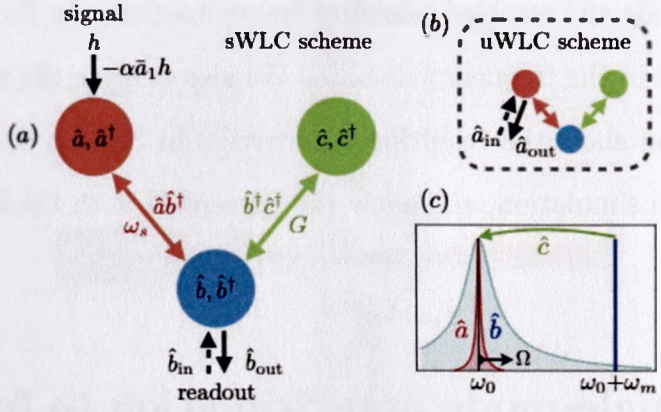


FIGURE 5.2: Idealised mode interaction structure of the optomechanical system illustrated in Fig. 5.1. **(a)** The sWLC scheme. Here $\hat{a}$ represents the differential mode of the arm cavity which couples to the GW strain h, $\hat{b}$ represents the SRC (filter) mode, and $\hat{c}$ represents the mechanical mode of the suspended mirror. The coupling between $\hat{a}$ and $\hat{b}$ ($\hat{b}$ and $\hat{c}$) is characterised by ω_s (G). Mode $\hat{b}$ is coupled to the external ingoing field $\hat{b}_{\rm in}$, which carries the vacuum noise, and the outgoing field $\hat{b}_{\rm out}$, which carries the signal and will be read out. **(b)** The uWLC scheme for comparison. In the intracavity readout scheme presented in Ref. [163], the external fields couple to the arm cavity mode $\hat{a}$, rather than the filter cavity mode $\hat{b}$. **(c)** The frequency-domain mode structure under single-mode (for both $\hat{a}$ and $\hat{b}$) and resolved sideband (for $\hat{b}$) approximations. The parametric interaction $\hat{b}^\dagger\hat{c}^\dagger$ is realized by the optomechanical coupling under blue-detuned pumping by the mechanical resonant frequency. Note that all mode operators are defined in the rotating frame of frequency ω_0, and Ω is the sideband frequency with respect to it.

of motion for the three modes:

$$\begin{aligned}
\dot{\hat{a}}(t) &= -\omega_s \hat{b}(t) + i\alpha h(t)\,, \\
\dot{\hat{c}}^\dagger(t) &= G\,\hat{b}(t)\,, \\
\dot{\hat{b}}(t) &= -\gamma\hat{b}(t) + \omega_s\hat{a}(t) + G\hat{c}^\dagger(t) + \sqrt{2\gamma}\,\hat{b}_{\rm in}(t)\,,
\end{aligned} \tag{5.2}$$

with outgoing field given by $\hat{b}_{\rm out}(t) = -\hat{b}_{\rm in}(t) + \sqrt{2\gamma}\,\hat{b}(t)$. The above equations can be solved in frequency domain via Fourier transform:

$$\hat{o}(\Omega) = \mathcal{F}[\hat{o}(t)] \equiv \int_{-\infty}^{\infty} {\rm d}t\, e^{i\Omega t}\hat{o}(t)\,, \tag{5.3}$$

where Ω is the sideband frequency in the rotating frame of frequency ω_0, and $\hat{o}$ represents $\hat{a}, \hat{b}, \hat{c}$ or $\hat{b}_{\mathrm{in,out}}$. The resulting solution for the outgoing field is:

$$\begin{aligned}\hat{b}_{\mathrm{out}}(\Omega) =& \frac{i\Omega(\gamma + i\Omega) - G^2 + \omega_s^2}{i\Omega(\gamma - i\Omega) + G^2 - \omega_s^2}\,\hat{b}_{\mathrm{in}}(\Omega) \\ &+ \frac{i\sqrt{2\gamma}\,\omega_s \alpha}{i\Omega(\gamma - i\Omega) + G^2 - \omega_s^2} h(\Omega)\,. \end{aligned} \tag{5.4}$$

Interestingly, regardless of the value of G and ω_s, the outgoing field is not squeezed, as the modulus of input-to-output transfer function remains equal to unity. The signal is contained entirely in the phase quadrature [27, 167] defined as $\hat{Y} = (\hat{b}_{\mathrm{out}} - \hat{b}_{\mathrm{out}}^{\dagger})/(\sqrt{2}\,i)$. The resulting signal-referred shot noise spectral density when measuring the phase quadrature is given by

$$S_{hh}(\Omega) = \frac{\Omega^2\gamma^2 + (G^2 - \omega_s^2 + \Omega^2)^2}{4\gamma\omega_s^2\alpha^2}\,. \tag{5.5}$$

The sensitivity given different relations between G and ω_s is plotted in Fig. 5.3. One interesting case is when $G = \omega_s$ where the noise spectral density vanishes at DC as the signal response diverges:

$$S_{hh}(\Omega)|_{G=\omega_s} = \frac{\Omega^2(\Omega^2 + \gamma^2)}{4\gamma\omega_s^2\alpha^2}\,. \tag{5.6}$$

The system stability, as analysed in Ref. [165], is determined by the poles of the transfer function in Eq. 5.4, i.e. the roots of

$$i\Omega(\gamma - i\Omega) + G^2 - \omega_s^2 = 0\,. \tag{5.7}$$

When the imaginary part of any root becomes positive, the system becomes unstable.

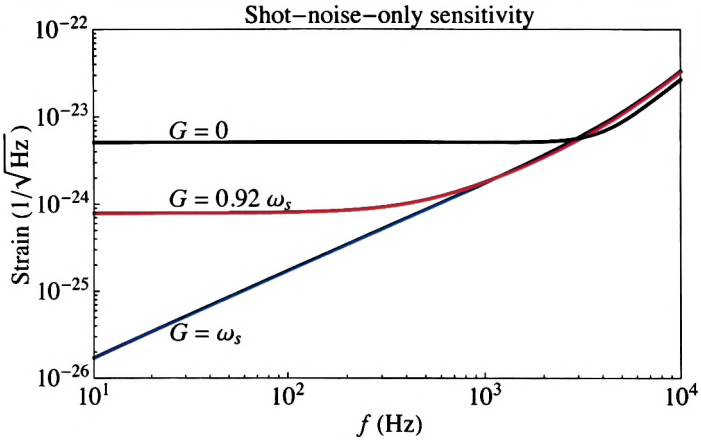


FIGURE 5.3: The signal-referred shot noise spectral density for different optomechanical interaction strengths G. Different curves are plotted based on Eq. 5.5 with the relation between G and ω_s shown near each line.

The locations of roots are determined by the relation between G and ω_s. It can be proven [165] that whenever G is equal to or less than ω_s, there is no unstable root. The critical point happens when $G = \omega_s$, and there is a pole at DC, which is consistent with the resulting signal response being infinite at DC, as can be seen in Eq. 5.4 and 5.6. To summarise, under the idealised Hamiltonian, the system is stable with $G \leq \omega_s$ and at the same time, the shot-noise-limited sensitivity will be improved. In the following content, we will show that these features will remain intact even after we relax the approximations applied in deriving the idealised Hamiltonian.

In terms of the physical parameters described in Table 5.1, the two approximations to produce the idealised Hamiltonian in Eq. 5.1 are: (i) the single-mode approximations, i.e. $\Omega L_{\text{arm}}/c \ll 1$ and $2\omega_{m_0} L_{\text{SRC}}/c \ll 1$, which treats the arm cavity and the filter cavity as single modes each described by an annihilation operators $\hat{a}$ and $\hat{b}$, and (ii) the resolved-sideband approximation, i.e. $\omega_{m_0} L_{\text{SRC}}/c \gg \gamma$, which treats the mechanical sidebands around ω_0 as a single operator $\hat{c}$, ignoring the higher sidebands around $\omega_0 + 2\omega_{m_0}$ that should be involved in the interaction between the blue detuned filter

TABLE 5.1: A list of parameters and nominal values

Parameters	Description	Value
$L_{\rm arm}$	arm cavity length	4 km
$P_{\rm arm}$	arm cavity power	800 kW
$T_{\rm ITM}$	ITM power transmissivity	0.5%
$L_{\rm SRC}$	SRC length	40 m
$T_{\rm SRM}$	SRM power transmissivity	0.02
P_b	filter cavity power	6.4 kW
λ	laser wavelength	1064 nm
m	oscillator mass	10 mg
$\omega_{m_0}/(2\pi)$	mechanical frequency	10^5 Hz
Q_m	mechanical quality factor	∞[1]

cavity and the mechanical oscillator. Note that all the parameters ω_s, G, α, γ used in Eq. 5.1–5.7 above are effective parameters, that can be approximately expressed in terms of the physical parameters. Under the single-mode and resolved-sideband approximations, the mode interaction strength ω_s and G can be expressed as:

$$\omega_s = \frac{c\sqrt{T_{\rm ITM}}}{2\sqrt{L_{\rm arm}L_{\rm SRC}}}, \quad G = \sqrt{\frac{8\pi P_b}{m\lambda\omega_{m_0}L_{\rm SRC}}}. \tag{5.8}$$

where P_b is the power of the filter cavity pumping field that impinges on the mechanical oscillator. Also, the signal coupling strength α and the SRC cavity half-bandwidth (i.e. the decay rate, or bath coupling rate) γ are defined as:

$$\alpha = \sqrt{\frac{P_{\rm arm}L_{\rm arm}\omega_0}{c\hbar}}, \quad \gamma = \frac{cT_{\rm SRM}}{4L_{\rm SRC}}. \tag{5.9}$$

For the full analysis, we adopt the approach in Ref. [29] by propagating the fields through the interferometer and taking into account their interactions with the mechanical degree of freedom via the radiation pressure. As all optical elements are axisymmetric,

for simplicity, the optical fields can be treated as 1D propagating ones along the optical axis. At each location, the optical field is represented as:

$$\hat{E}(t) = \hat{o}(t)e^{-i\omega_0 t} + \hat{o}^\dagger(t)e^{i\omega_0 t}\,, \tag{5.10}$$

where $\hat{o}(t)$ can represent $\hat{a}(t)$ and $\hat{b}(t)$, the slowly-varying field operators in the rotating frame of ω_0 for the light inside the arm cavity and filter cavity respectively. We can define sideband operators $\hat{a}(\omega)$ and $\hat{b}(\omega)$ via Fourier transform as in Eq. 5.3, but for a large sideband frequency ω up to the order of ω_{m_0}. Note that so long as $\omega \ll \omega_0$, the slowly-varying operators in time domain and sideband operators in frequency domain will be well-defined.

The relevant fields fully describing the scheme in Fig. 5.1 are illustrated in Fig. 5.4. We will focus on the differential mode of the two arm cavities, and the right panel of Fig. 5.4 shows the simplified representation considered here. We assume both the arm cavity and the SRC are tuned such that the cavity lengths are integer numbers times the wavelength of the carrier at frequency ω_0. Since we are looking at the linear dynamics, the equations of motion will only involve linear terms of the sideband operators. The Heisenberg equation of motion is formally identical to the classical Maxwell equation.

The field operators are described by the following set of equations:

$$\hat{a}_1(t) = \sqrt{T_{\rm ITM}}\,\hat{b}_3(t) + \sqrt{R_{\rm ITM}}\,\hat{a}_2(t)\,, \tag{5.11a}$$

$$\hat{a}_2(t) = \hat{a}_1(t - \tau_{\rm arm}) + 2ik_0 A_{\rm arm} L_{\rm arm} h(t)\,, \tag{5.11b}$$

$$\hat{b}_4(t) = \sqrt{T_{\rm ITM}}\,\hat{a}_2(t) - \sqrt{R_{\rm ITM}}\,\hat{b}_3(t)\,, \tag{5.11c}$$

$$\hat{b}_2(t) = \hat{b}_4(t - \tau_{\rm SRC}/2) + 2ik_b A_b e^{-i\omega_{m_0} t}\hat{x}(t)\,, \tag{5.11d}$$

$$\hat{b}_3(t) = \hat{b}_1(t - \tau_{\rm SRC}/2) + 2ik_b A_b e^{-i\omega_{m_0} t}\hat{x}(t)\,, \tag{5.11e}$$

$$\hat{b}_1(t) = \sqrt{T_{\rm SRM}}\,\hat{b}_{\rm in}(t) - \sqrt{R_{\rm SRM}}\,\hat{b}_2(t)\,, \tag{5.11f}$$

$$\hat{b}_{\rm out}(t) = \sqrt{T_{\rm SRM}}\,\hat{b}_2(t) + \sqrt{R_{\rm SRM}}\,\hat{b}_{\rm in}(t)\,. \tag{5.11g}$$

Here the round-trip delay times are defined as $\tau_{\rm arm} = 2L_{\rm arm}/c$ and $\tau_{\rm SRC} = 2L_{\rm SRC}/c$, for the arm cavity and SRC respectively. The two cavities are pumped with frequency ω_0 and $\omega_0 + \omega_{m_0}$, with wave vectors being $k_0 = \omega_0/c$ and $k_b = (\omega_0 + \omega_{m_0})/c \approx k_0$ ($\omega_0 \gg \omega_{m_0}$) respectively. The steady-state field amplitudes $A_{\rm arm}$ and A_b are given by $A_{\rm arm} = \sqrt{P_{\rm arm}/(2\hbar\omega_0)}$ and $A_b = \sqrt{P_b/2\hbar\omega_0}$. The mechanical motion is driven by the radiation pressure in the presence of the blue-detuned pump field:

$$\ddot{\hat{x}}(t) + \gamma_m \dot{\hat{x}}(t) + \omega_{m_0}^2 \hat{x}(t) = \frac{\hat{F}_{\rm rad}(t)}{m}\,, \tag{5.12}$$

where the radiation pressure $\hat{F}_{\rm rad}$ reads:

$$\hat{F}_{\rm rad}(t) = \frac{2\hbar\omega_0 A_{\rm c}^*}{c} e^{i\omega_{m_0} t}\left[\hat{b}_1(t) + \hat{b}_4(t)\right] + \text{h.c.} \tag{5.13}$$

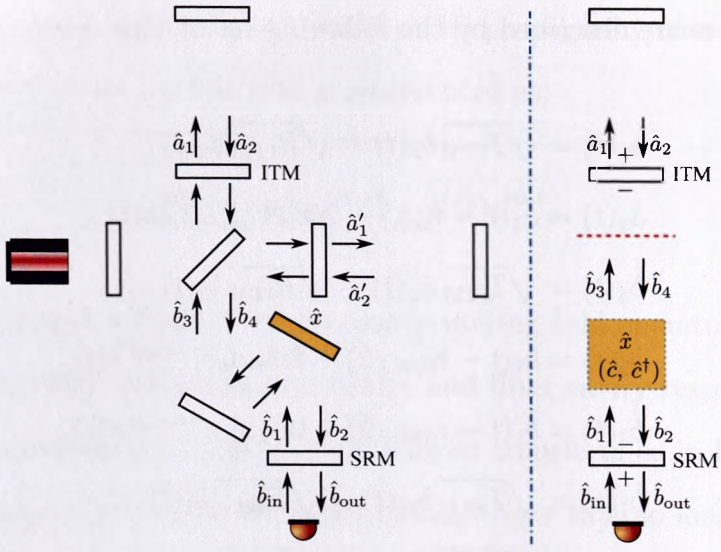


FIGURE 5.4: The fields involved in the complete analysis of the scheme shown in Fig. 5.1. The right figure is a simplified version of the left one when only looking at the differential mode of the two arms, which is adopted in the analysis in the main text. In the right figure, the sign convention for the mirror reflectivity is also shown. The arm cavity and the SRC length are both tuned to be integer numbers times the wavelength of the carrier at frequency ω_0.

The displacement operator $\hat{x}$ is related to the mechanical mode operator $\hat{c}$ by:

$$\hat{x}(t) = \sqrt{\frac{\hbar}{2m\omega_{m_0}}}\left[\hat{c}(t)e^{-i\omega_{m_0}t} + \hat{c}^\dagger(t)e^{i\omega_{m_0}t}\right] . \tag{5.14}$$

5.3 Frequency-domain analysis

In this section, we will solve the system dynamics in the frequency domain, and analyse the stability using the Nyquist criterion [168].

5.3.1 Formalism

In the frequency domain, Eq. 5.11a–5.11g can be converted into algebra equations in the matrix representation. For the slowly-varying field operators defined in Eq. 5.10, we apply the time-shifting relation in the Fourier transform of optical modes, i.e.

$$\mathcal{F}[\hat{o}(t-\tau)] = \hat{o}(\omega)e^{i\omega\tau} \text{ for } \hat{o} = \hat{a} \text{ and } \hat{b}\,, \tag{5.15}$$

to connect propagating light inside the cavities, where ω is the sideband frequency in the rotating frame of ω_0. The Fourier transform of the mechanical mode in the filter cavity reads:

$$\mathcal{F}[e^{-i\omega_{m_0}t}\hat{x}(t)] = \hat{x}(\omega - \omega_{m_0})\,. \tag{5.16}$$

With the ω_m blue-detuned pumping, the optomechanical interaction will couple the optical sidebands at frequency $\omega = \pm\Omega$ ($\Omega \ll \omega_{m_0}$) with the sidebands at $\omega = 2\omega_m \mp \Omega$. As mentioned in Fig. 5.1, the value of ω_m is slightly difference from the mechanical resonant frequency ω_{m_0}. The difference is caused by the optical spring effect,

$$\omega_m = \omega_{m_0} + \frac{P_b\omega_0}{2m\omega_{m_0}^2c^2\tau_b}. \tag{5.17}$$

when we consider the realistic scenario beyond the resolved-sideband approximation. As illustrated in Fig. 5.5 (a)–(b), when there is optical spring, the parametric interaction will be shifted by a small frequency, causing the demanded sidebands not to be correctly coupled. As our system works near a critical point of PT symmetry [165], the effect of a rather small optical spring effect is fairly important and needs to be carefully compensated.

As illustrated in Fig. 5.5 (c), the sidebands around ω_0 are named as the signal

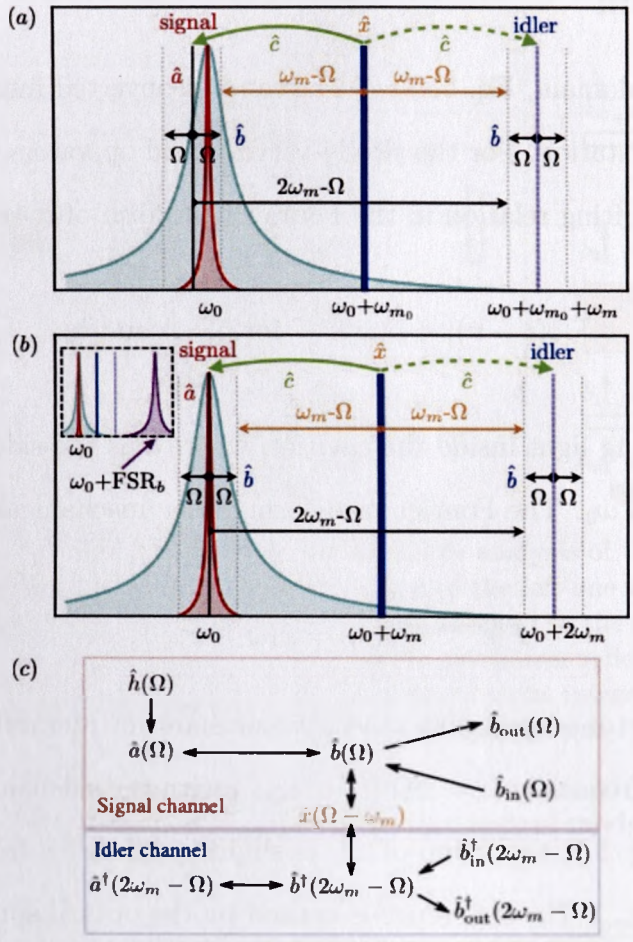


FIGURE 5.5: Coupling between the optical modes and the mechanical modes represented in frequency domain. **(a)** Without optical spring compensation, the parametric interaction will be shifted by a small frequency (not to scale) which will greatly ruin the designed filter response. **(b)** With optical spring compensation, the correct sidebands are coupled. It is achieved by making the value of blue-detuning to be slightly larger than the bare mechanical resonant frequency, as described in Eq. 5.17. Insert: the best scenario for resolved-sideband is to set $2\omega_m$ around half FSR$_b$. It will make ω_m larger than γ_b as much as possible and, at the same time, make the influence from the higher FSR as little as possible. **(c)** The full circulation loop within the signal channel (sidebands around ω_0, represented by the red box in the upper panel) and the idler channel (sidebands around $\omega_0 + 2\omega_m$, represented by the blue box in the lower panel). The idler channel is ignored in the idealised Hamiltonian analysis.

channel, and the ones around $\omega_0 + 2\omega_m$ are named as the idler channel. We will first treat the filter cavity as an effective mirror, obtain the transformation relation for signal

and idler channels, as well as mixing between the channels, and then combine it with the circulation loop of the arm cavity. As the filter cavity is pumped with frequency $\omega_0 + \omega_m$, for convenience, we temporarily use $\tilde{o}(\tilde{\omega})$ to represent the sidebands in the rotating frame of $\omega_0 + \omega_m$, where $\tilde{o}$ can be $\tilde{a}$ or $\tilde{b}$, and $\tilde{\omega}$ is in the order of magnitude of ω_{m_0}. The input-output relation of the filter cavity can be represented as follows:

$$\begin{bmatrix} \tilde{b}_{\text{out}}(\tilde{\omega}) \\ \tilde{b}^{\dagger}_{\text{out}}(-\tilde{\omega}) \\ \tilde{a}_1(\tilde{\omega}) \\ \tilde{a}^{\dagger}_1(-\tilde{\omega}) \end{bmatrix} = \tilde{\mathbf{M}}_{\text{filter}}(\tilde{\omega}) \begin{bmatrix} \tilde{b}_{\text{in}}(\tilde{\omega}) \\ \tilde{b}^{\dagger}_{\text{in}}(-\tilde{\omega}) \\ \tilde{a}_2(\tilde{\omega}) \\ \tilde{a}^{\dagger}_2(-\tilde{\omega}) \end{bmatrix}. \tag{5.18}$$

where $\tilde{\mathbf{M}}_{\text{filter}}(\tilde{\omega})$ is determined by the optomechanical interaction and the field circulation inside the filter cavity.

Before connecting the filter cavity to the main interferometer, let's comment on the required functionality of its transfer function in the full circulation loop. Notice that in the lossless case, the open-loop circulation in the arm cavity will result in a phase delay, without harming the amplitude. Thus, we expect the feedback gain provided by the signal channel reflection $\hat{a}_1(\Omega) \rightarrow \hat{a}_2(\Omega)$ to compensate that effect by providing the same amount of phase advance with an amplitude 1. For the phase advance, we require $4\omega_0 P_f/(m\omega_m c^2 T_{\text{ITM}}) = c/L_{\text{arm}}$ [163] and optical spring compensation to correctly amplify the demanded sidebands. For the physical parameters shown in Table 5.1, the optical spring leads to a shift of the mechanical frequency by around 77 Hz. In the actual simulation, this value will be compensated numerically to avoid any tiny discrepancy. The influence of optical spring in the signal-referred shot-noise and stability will be analysed in the following Sections 5.3.2 and 5.3.3. The other

requirement is that the amplitude of reflection to be 1, as we want the idler channel to mix with the signal channel as little as possible. The best scenario for that purpose is to set $2\omega_m$ around half FSR_b, as shown by the insert of Fig. 5.5 (b). It will make $\omega_m >> \gamma_b$ as much as possible and, at the same time, suppress the higher FSR from amplifying the idler channel. In addition, we want the mechanical loss to be as small as possible, as the filter cavity contains a parametric process and thus the mechanical loss will ruin the reflection amplitude.

By setting $\tilde{\omega} = \mp\omega_m + \Omega$, we can extract relevant matrix components to compose the transformation matrices for the optical fields outside the filter cavity:

$$\begin{bmatrix} \hat{b}_{\mathrm{out}}(\Omega) \\ \hat{b}^{\dagger}_{\mathrm{out}}(2\omega_m - \Omega) \\ \hat{a}_1(\Omega) \\ \hat{a}^{\dagger}_1(2\omega_m - \Omega) \end{bmatrix} = \tilde{\mathrm{M}}_{\mathrm{filter}}(-\omega_m + \Omega) \begin{bmatrix} \tilde{b}_{\mathrm{in}}(\Omega) \\ \tilde{b}^{\dagger}_{\mathrm{in}}(2\omega_m - \Omega) \\ \tilde{a}_2(\Omega) \\ \tilde{a}^{\dagger}_2(2\omega_m - \Omega) \end{bmatrix}, \tag{5.19a}$$

$$\begin{bmatrix} \hat{b}_{\mathrm{out}}(2\omega_m + \Omega) \\ \hat{b}^{\dagger}_{\mathrm{out}}(-\Omega) \\ \hat{a}_1(2\omega_m + \Omega) \\ \hat{a}^{\dagger}_1(-\Omega) \end{bmatrix} = \tilde{\mathrm{M}}_{\mathrm{filter}}(\omega_m + \Omega) \begin{bmatrix} \tilde{b}_{\mathrm{in}}(-\Omega) \\ \tilde{b}^{\dagger}_{\mathrm{in}}(2\omega_m - \Omega) \\ \tilde{a}_2(-\Omega) \\ \tilde{a}^{\dagger}_2(2\omega_m - \Omega) \end{bmatrix}. \tag{5.19b}$$

For simplicity, we define $\tilde{\mathrm{M}}_{\mp} \equiv \tilde{\mathrm{M}}_{\mathrm{filter}}(\mp\omega_m + \Omega)$ and thus have:

$$\mathbf{R}_{aa}(\Omega) \equiv \begin{bmatrix} \tilde{\mathrm{M}}_{-}^{3,3}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{3,4} \\ 0, & \tilde{\mathrm{M}}_{+}^{4,4}, & \tilde{\mathrm{M}}_{+}^{4,3}, & 0 \\ 0, & \tilde{\mathrm{M}}_{+}^{3,4}, & \tilde{\mathrm{M}}_{+}^{3,3}, & 0 \\ \tilde{\mathrm{M}}_{-}^{4,3}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{4,4} \end{bmatrix}, \tag{5.20a}$$

$$\mathbf{T}_{ab}(\Omega) \equiv \begin{bmatrix} \tilde{\mathrm{M}}_{-}^{3,1}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{3,2} \\ 0, & \tilde{\mathrm{M}}_{+}^{4,2}, & \tilde{\mathrm{M}}_{+}^{4,1}, & 0 \\ 0, & \tilde{\mathrm{M}}_{+}^{3,2}, & \tilde{\mathrm{M}}_{+}^{3,1}, & 0 \\ \tilde{\mathrm{M}}_{-}^{4,1}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{4,2} \end{bmatrix}, \tag{5.20b}$$

$$\mathbf{R}_{bb}(\Omega) \equiv \begin{bmatrix} \tilde{\mathrm{M}}_{-}^{1,1}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{1,2} \\ 0, & \tilde{\mathrm{M}}_{+}^{2,2}, & \tilde{\mathrm{M}}_{+}^{2,1}, & 0 \\ 0, & \tilde{\mathrm{M}}_{+}^{1,2}, & \tilde{\mathrm{M}}_{+}^{1,1}, & 0 \\ \tilde{\mathrm{M}}_{-}^{2,1}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{2,2} \end{bmatrix}, \tag{5.20c}$$

$$\mathbf{T}_{ba}(\Omega) \equiv \begin{bmatrix} \tilde{\mathrm{M}}_{-}^{1,3}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{1,4} \\ 0, & \tilde{\mathrm{M}}_{+}^{2,4}, & \tilde{\mathrm{M}}_{+}^{2,3}, & 0 \\ 0, & \tilde{\mathrm{M}}_{+}^{1,4}, & \tilde{\mathrm{M}}_{+}^{1,3}, & 0 \\ \tilde{\mathrm{M}}_{-}^{2,3}, & 0, & 0, & \tilde{\mathrm{M}}_{-}^{2,4} \end{bmatrix}, \tag{5.20d}$$

where $\mathbf{R}_{aa}(\Omega)$ represents the $\hat{a}_2 \rightarrow \hat{a}_1$ reflection, $\mathbf{T}_{ab}(\Omega)$ represents the $\hat{b}_{\mathrm{in}} \rightarrow \hat{a}_1$ transmission, $\mathbf{R}_{bb}(\Omega)$ represents the $\hat{b}_{\mathrm{in}} \rightarrow \hat{b}_{\mathrm{out}}$ reflection, $\mathbf{T}_{ba}(\Omega)$ represents the $\hat{a}_2 \rightarrow \hat{b}_{\mathrm{out}}$ transmission. Thus, the effect of filter cavity on the optical fields can be expressed as

follows:

$$\begin{bmatrix} \hat{a}_1(\Omega) \\ \hat{a}_1^\dagger(-\Omega) \\ \hat{a}_1(2\omega_m+\Omega) \\ \hat{a}_1^\dagger(2\omega_m-\Omega) \end{bmatrix} = \mathbf{R}_{aa}(\Omega) \begin{bmatrix} \hat{a}_2(\Omega) \\ \hat{a}_2^\dagger(-\Omega) \\ \hat{a}_2(2\omega_m+\Omega) \\ \hat{a}_2^\dagger(2\omega_m-\Omega) \end{bmatrix} + \mathbf{T}_{ab}(\Omega) \begin{bmatrix} \hat{b}_{\text{in}}(\Omega) \\ \hat{b}_{\text{in}}^\dagger(-\Omega) \\ \hat{b}_{\text{in}}(2\omega_m+\Omega) \\ \hat{b}_{\text{in}}^\dagger(2\omega_m-\Omega) \end{bmatrix}, \tag{5.21a}$$

$$\begin{bmatrix} \hat{b}_{\text{out}}(\Omega) \\ \hat{b}_{\text{out}}^\dagger(-\Omega) \\ \hat{b}_{\text{out}}(2\omega_m+\Omega) \\ \hat{b}_{\text{out}}^\dagger(2\omega_m-\Omega) \end{bmatrix} = \mathbf{R}_{bb}(\Omega) \begin{bmatrix} \hat{b}_{\text{in}}(\Omega) \\ \hat{b}_{\text{in}}^\dagger(-\Omega) \\ \hat{b}_{\text{in}}(2\omega_m+\Omega) \\ \hat{b}_{\text{in}}^\dagger(2\omega_m-\Omega) \end{bmatrix} + \mathbf{T}_{ba}(\Omega) \begin{bmatrix} \hat{a}_2(\Omega) \\ \hat{a}_2^\dagger(-\Omega) \\ \hat{a}_2(2\omega_m+\Omega) \\ \hat{a}_2^\dagger(2\omega_m-\Omega) \end{bmatrix}. \tag{5.21b}$$

In the idealised analysis in Section 5.2, the idler channel was ignored under resolved-sideband limit. The full analysis here will keep both signal and idler channels. Further linking the filter cavity to the main interferometer, propagation in the arm cavity will be considered for both channels to form a closed-loop transfer function. GW strain

signal will be added to the signal channel in the arm cavity and the final input-output relation takes the following form:

$$\begin{bmatrix} \hat{b}_{\rm out}(\Omega) \\ \hat{b}^{\dagger}_{\rm out}(2\omega_m - \Omega) \end{bmatrix} = \mathbf{M}(\Omega) \begin{bmatrix} \hat{b}_{\rm in}(\Omega) \\ \hat{b}^{\dagger}_{\rm in}(2\omega_m - \Omega) \end{bmatrix} + \mathbf{v}(\Omega)\, h(\Omega)\,, \tag{5.22}$$

where $\mathbf{M}(\Omega)$ is a 2×2 matrix describing the transformation of the ingoing field to the outgoing field, and $\mathbf{v}(\Omega)$ is a 2×1 column vector describing the response of the outgoing field to the GW signal, with the $\mathbf{v}(\Omega)^{2,1}$ component describing the signal leading into the idler channel $\hat{b}^{\dagger}_{\rm in}(2\omega_m - \Omega)$. The exact expressions for $\mathbf{M}(\Omega)$ and $\mathbf{v}(\Omega)$ are quite complicated, and in the subsection that follows, we will use the quantum noise spectral density to illustrate their frequency dependence.

5.3.2 Noise Spectral Density

The resulting noise spectral density from the full analysis is shown in Fig. 5.6 and Fig. 5.7. In Fig. 5.6, the optical spring effect is shown to play an important role in modifying the sensitivity. As analysed in Section 5.3.1, if left unattended, the optical spring will cause an additional resonance in the sensitivity via shifting the central frequency of signal sidebands in the filter cavity response. After compensating such a shift, the full analysis and the idealised Hamiltonian analysis produce a similar result for most of frequencies, given the same set of parameters. However, they start to deviate at low frequencies, where the full analysis shows a higher noise spectrum.

The deviation from idealised analysis is because of the coupling to the idler channel at low frequencies, as shown in Fig. 5.7, where a larger amount of signal will leak to the idler channel than what is output in the signal channel. Thus, the idler channel

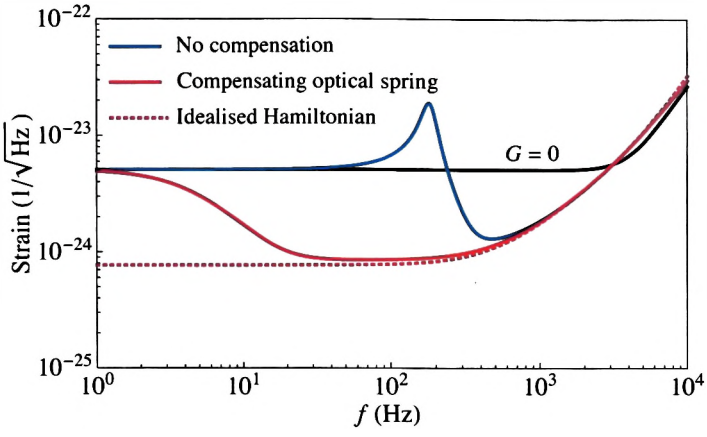


FIGURE 5.6: The signal-referred shot-noise spectral densities from the full analysis (red and blue) in comparison with the one obtained from idealised Hamiltonian (red dotted curve). Compensating the frequency shift of the mechanical oscillator due to the optical spring effect has a significant influence on the sensitivity.

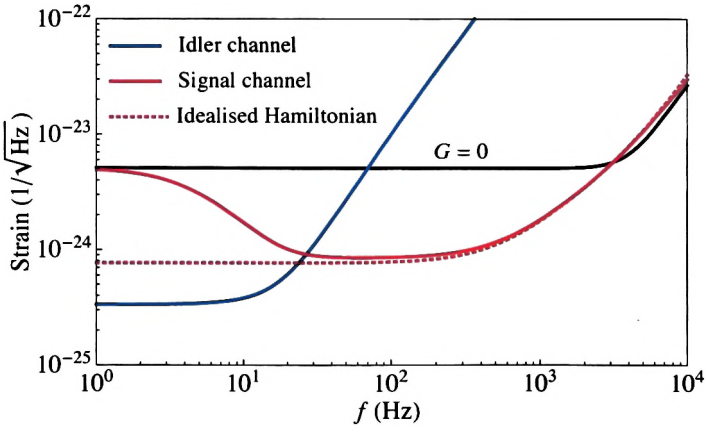


FIGURE 5.7: The noise spectral density for the idler channel, which can be read out by using a local oscillator at $\omega_0 + 2\omega_{m_0}$ (blue), in comparison with the signal channel (red). At low frequencies, the signal information contained in the idler channel is even more than that contained in the signal channel.

has a much better sensitivity at low frequencies. We can extract the signal information contained in the idler channel either by using the heterodyne readout with the pump field at $\omega_0 + \omega_m$ as the local oscillator or using an additional at $\omega_0 + 2\omega_m$ for the homodyne readout of the idler channel. We will discuss the details about the additional

homodyne readout scheme in Section 5.5, and show the final sensitivity by optimally combining the two readout schemes.

As we will show in Section 5.3.3, the parameters in Table 5.1 is within the stability regime. Therefore, we can indeed obtain sensitivity improvement without sacrificing the stability.

5.3.3 Stability Analysis

In this section, we use the Nyquist technique [168] to analyse the system stability. The criterion is a diagrammatic approach to test the stability of a system by only using the open-loop transfer function, even with delay.

In the idealised analysis in Section 5.2, $G \approx \omega_s$ defines the boundary for stability. Further increasing G through increasing the pumping power of the filter cavity will destabilise the system. In the full analysis when the idler channel is also included, we have a more complicated multi-input-multi-output (MIMO) system. In the sideband picture, given the 2-dimensional open-loop transfer matrix $\mathbf{M}_{\mathrm{OL}}(\Omega)$, the real and imaginary part of the determinant of $\mathbf{I} + \mathbf{M}_{\mathrm{OL}}(\Omega)$, where $\mathbf{I}$ is the 2-dimensional identity matrix, should not enclose the origin of the complex plain. This is limited to the case where $\mathbf{M}_{\mathrm{OL}}(\Omega)$ does not contain elements that have unstable poles, and we therefore also need to be careful about the point where to extract the open-loop transfer function, even though the stability of the final closed-loop transfer function is independent of such a choice. We choose the interface near ITM as highlighted by the red dashed line on the right panel of Fig. 5.4, where the top part is the arm cavity and the bottom is consisted of the mechanical oscillator and SRM.

The open-loop and the close-loop transfer matrices in our case are:

$$\mathbf{M}_{\mathrm{OL}}(\Omega) = \mathbf{M}_{\mathrm{opt}}(\Omega)\mathbf{M}_{\mathrm{cav}}(\Omega), \quad \mathbf{M}_{\mathrm{CL}}(\Omega) = (\mathbf{I} + \mathbf{M}_{\mathrm{OL}}(\Omega))^{-1}. \tag{5.23}$$

The optomechanical transfer matrix $\mathbf{M}_{\mathrm{opt}}(\Omega)$ is defined as:

$$\mathbf{M}_{\mathrm{opt}}(\Omega) = e^{i\Omega\tau_{\mathrm{SRC}}/2} \begin{bmatrix} \mathcal{T}_{\mathrm{opt}} & \mathcal{T}_{\mathrm{opt}} + \sqrt{R_{\mathrm{SRM}}} \\ -\sqrt{R_{\mathrm{SRM}}} - \mathcal{T}_{\mathrm{opt}} & -2\sqrt{R_{\mathrm{SRM}}} - \mathcal{T}_{\mathrm{opt}} \end{bmatrix}, \tag{5.24}$$

where

$$\mathcal{T}_{\mathrm{opt}} = -\sqrt{R_{\mathrm{SRM}}} + \frac{i(1+\sqrt{R_{\mathrm{SRM}}})^2 g^2 \tau_{\mathrm{SRC}}\, \omega_{m_0}}{2[\Omega(\Omega - 2\omega_m) + i\gamma_m(\Omega - \omega_m)]}. \tag{5.25}$$

Similarly, the passive cavity propagation matrix is:

$$\mathbf{M}_{\mathrm{cav}}(\Omega) = \begin{bmatrix} \mathcal{T}_{\mathrm{cav}}(\Omega) & 0 \\ 0 & \mathcal{T}^*_{\mathrm{cav}}(2\omega_m - \Omega) \end{bmatrix}, \tag{5.26}$$

where

$$\mathcal{T}_{\mathrm{cav}}(\Omega) = \frac{e^{i\Omega\tau_{\mathrm{arm}}} - \sqrt{R_{\mathrm{ITM}}}}{1 - e^{i\Omega\tau_{\mathrm{arm}}}\sqrt{R_{\mathrm{ITM}}}}. \tag{5.27}$$

In Fig. 5.8, we show the resulting Nyquist plot for the nominal parameters in Table 5.1. Because the mechanical frequency is not an integer number of the FSR of the arm cavity, the contour is quite complicated due to the phase factor $e^{2i\omega_{m_0}\tau_{\mathrm{arm}}}$ of the idler channel. Nevertheless, we only need to check whether the origin is enclosed or not to make a firm claim on the stability, while the complex feature of the contour does not matter. It turns out that, compensating the optical spring shift not only has a significant impact on the sensitivity, as shown in Fig. 5.6, but also on the stability. As shown in the plot on the right hand side of Fig. 5.8, With the nominal values for the

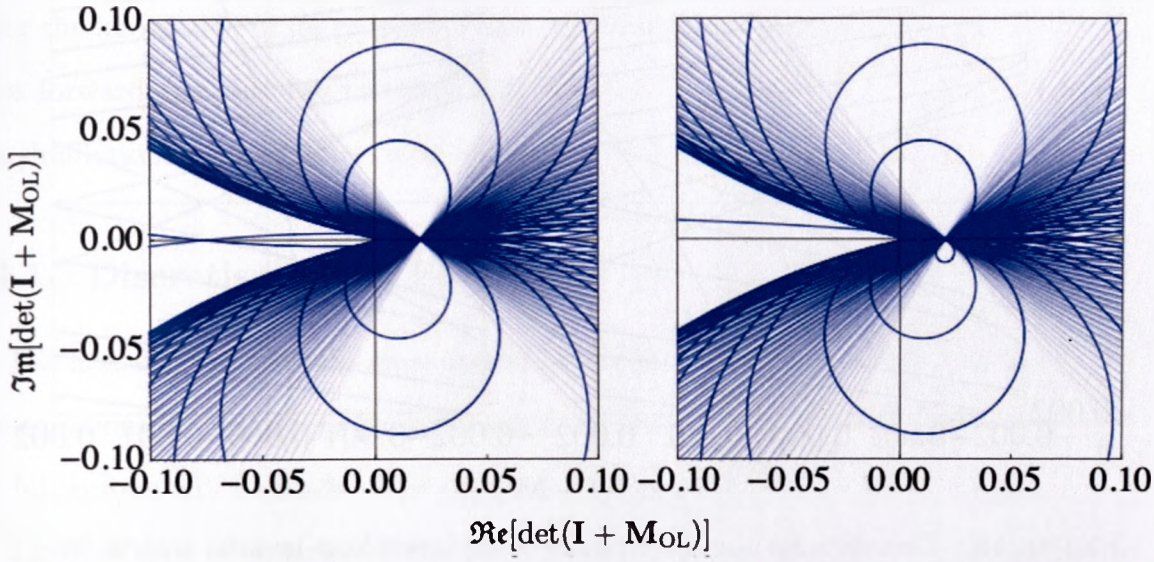


FIGURE 5.8: The plots for the determinant of the open-loop transfer matrix, given the nominal parameters. The left plot shows the case without compensating the optical spring, while the right one has the optical spring compensated and the contour does not enclose the origin. The opacity of the curve is intentionally made smaller when the magnitude of the frequency is large, which is to highlight the relevant low frequencies.

parameters, the system only becomes stable after the optical spring shift is accounted for. This feature can be understood as follows: to meet the PT-symmetry condition, the parametric interaction with the mechanical mode must be achieved by the correct optical sidebands in the filter cavity; thus, only after the optical spring is correctly compensated, it is possible and meaningful to discuss the stability issue.

In Fig. 5.9, we further show the effect of the optomechanical coupling rate G on the stability, after the necessary optical spring compensation. Indeed, the system is stable when $G \leq \omega_s$, which is consistent with the idealised Hamiltonian analysis [165].

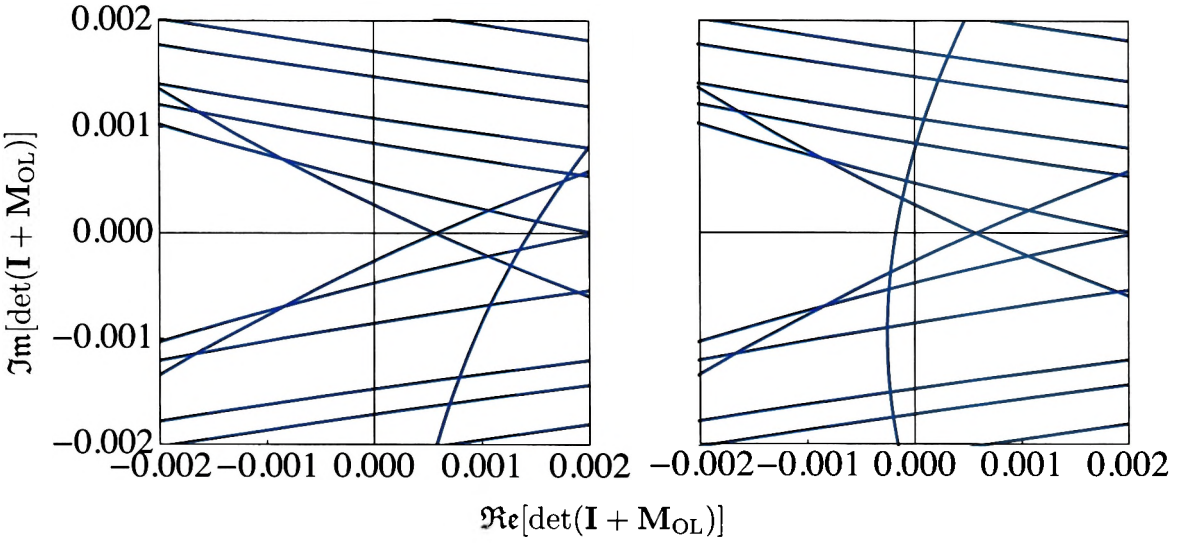


FIGURE 5.9: The plots for the determinant of the open-loop transfer matrix for $G = \omega_s$ (left) and $G = 1.01\,\omega_s$ (right), both with the optical spring shift compensated. For clarity, we only show a tiny regime around the origin, but we have checked that only the right plot has the origin enclosed when zooming out.

5.4 Time-domain analysis

In this section, we perform a numerical time-domain simulation to confirm the enhancement in sensitivity as derived above. The purpose of remaining in the time domain is to capture behaviour that may have been lost due to approximations used in the frequency-domain analysis. It is also much easier for us to capture non-linear behaviour in the time domain, as we have no need to perform Fourier transforms in this analysis. The primary limitation of this method is the relative lack of insight into the physics that we can obtain, as it is difficult to 'break open' the simulation and understand the behaviour of individual parts in isolation. As such, our time-domain approach is a powerful complementary tool to the analysis performed thus far and not a replacement.

We will begin by considering Eq. 5.11a–5.11g. The principle of this simulation is to

appropriately discretise the equations and enable evolve the system forward in time using the knowledge of its previous state. The mechanism through which the system steps forward and the way to recalculate relevant quantities are at the very heart of this simulation and, therefore, it is worthwhile to delve into some details.

5.4.1 Discretisation of Equations

The first step towards the time domain simulation is the discretisation of equations. The equations in our system can subdivided into three distinct types. The treatment of the full system will be made clear after the functionality of each type.is explored.

The first type of equation comprises all quantities that are related to others in the way of time-delay. For instance, the field quantities that propagate from one optical component to another. The general form of this type of equations reads:

$$a_2(t) = a_1(t - \tau), \tag{5.28}$$

where $\tau = L/c$ is the time delay caused by the propagation across the space of distance L, with c being the speed of light. This type of equations is discretised at time step n as

$$a_2[n] = a_1[n - n_d], \tag{5.29}$$

where n_d is time-step delay (steps in the past at which the field must be evaluated). It is given by $n_d = \tau/\Delta t$, where Δt is the simulation time increment. This type of equations encompasses the forward evolution of the field quantities, whereby a field at a new simulation step can be calculated only with reference to the previous value of another field, which will have been calculated at their step.

FIGURE 5.10: Field interactions at a mirror.

The second type of equation comprises interactions of fields at an optical component. This treatment follows a very standard approach, where the fields at a mirror (as defined in Fig. 5.10) are given by:

$$a_2(t) = t_m a_4(t) + r_m a_1(t), \tag{5.30a}$$

$$a_3(t) = t_m a_1(t) - r_m a_4(t), \tag{5.30b}$$

where r_m and t_m are the mirror amplitude reflectivity and transmissivity respectively. The discretisation of these fields is rather trivial:

$$a_2[n] = t_m a_4[n] + r_m a_1[n], \tag{5.31a}$$

$$a_3[n] = t_m a_1[n] - r_m a_4[n]. \tag{5.31b}$$

However, care must be taken in choosing the order to calculate these fields, as the field quantities are all initially unknown. The procedure for calculating these fields will be discussed later in due course.

The final consideration is the treatment of Eq. 5.12. It can be approached in many ways according to the choice of discretisation for time derivative. Based on its

well-studied nature [169], a "symmetric" approach is implemented:

$$\ddot{x}[n] + \gamma_m \dot{x}[n] + \omega_{m_0}^2 x[n] = \frac{F_{\text{rad}}[n]}{m}, \tag{5.32a}$$

$$\text{with } \dot{x}[n] = \frac{x[n+1] - x[n-1]}{2\Delta t}, \tag{5.32b}$$

$$\text{and } \ddot{x}[n] = \frac{x[n+1] - 2x[n] + x[n-1]}{\Delta t^2}. \tag{5.32c}$$

Note that substituting the latter two equations into Eq. 5.32a allows for the $x[n+1]$ term to be rearranged in terms of quantities at previous time steps and, therefore, it can be used to evolve the position of the mechanical oscillator without further issue.

5.4.2 Connecting the System

After categorizing the relative equations, the next to consider is connecting the discretized equations into a consistent loop. The most challenging feature of the system is the signal-recycling structure in the optical cavities, as the signal sent "downstream" from the simulation inputs returns and combines with the upstream signal in a feedback-like structure. For a system with many cavities, the cycling of signal could become very complicated to handle.

Considering that the signal must traverse the space between optical components, which induces a time delay, it allows us to break up the system into several independent "compartments". As the system is evolved forwards in time, when we consider fields at step n, we can rely on the values for all steps $< n$ that have already been calculated. This naturally means that all terms of the Eq. 5.28 type can be calculated immediately. To tackle the issue with Eq. 5.30, we can treat the right-hand side quantities $a_{1,4}(t)$ as the "inputs" to the mirror, which are already available, as the "outputs" or left-hand side of Eq. 5.28. This results in a logic flow, as shown in Fig. 5.11. At the start of a

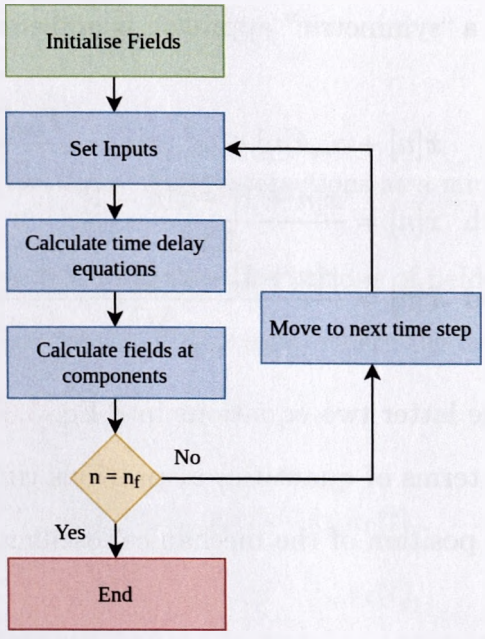


FIGURE 5.11: Logical flowchart of the time domain simulation.

new simulation step, we calculate all the "outputs" of the Eq. 5.28 type first, which gives all of the "inputs" quantities of the Eq. 5.30 types. We then calculate the "outputs" of all Eq. 5.30 types, which then affect the input fields in a neighbouring compartment that links the whole system together. Note that these links occur at a future time step and thus does not need to be calculated now, avoiding the trap of circular logic. This approach greatly simplifies the calculation performed in each simulation cycle. By compartmentalising the equations in this way, we can create several independent systems with very few equations each. It is thus very easy to add extra components to the system, without the need to rewrite any existing elements in the simulation. Furthermore, the system is well set up for optimisation techniques such as parallelisation.

The key issue in our approach is the large number of calculations to be performed, as well as the large number of quantities that need to be stored. To suitably capture the system behaviour, the time increment in simulation needs to be smaller than the shortest

time-scale of the system, which in our case is the short cavity traversal time around 100 ns. Although this MHz level of sampling is far above the frequency range of interest (1 Hz–10 kHz), it is necessary to capture the behaviour of components that affect the signal response within that frequency band. This will lead to a lot of waste as we are simulating about 1000 times more data than what we are interested in. Furthermore, we want to simulate for at least 1 second, which consists of 10 million samples, resulting in rather a lot of data for a computer to store.

A further limitation is the requirement for integer values of n_d in the discretisation of Eq. 5.29. For all cavities in the system, the traversal times must be some integer multiple of Δt. To enable the free choice of cavity lengths, particularly when there are many cavities in the system, Δt may have to be even smaller than the shortest cavity traversal time. This is not a big issue in our case as cavity number can be limited to be two. All the optical components are separated by a short delay time τ_s, except for the arm cavity with a long delay time τ_a. For convenience, we can choose τ_a to be an exact multiple of τ_s, without any loss of physical insight.

5.4.3 Step-Response Stability Analysis

The steady-state response is the most essentials in the sensitivity analysis of the system. As the simulation loop relies on past values of the system, initialisation is required at some finite time from which the simulation starts. To resolve this issue, all quantities in $t < 0$ are set to be 0, which corresponds to the off state. After the simulation begins, the building-up process of the steady-state leads to an initial transient effect, which might negatively affect the response estimate. Luckily, in our simulation, this initialisation takes about 0.03 s, which is sufficiently small compared to the total runtime and will not cause any problem.

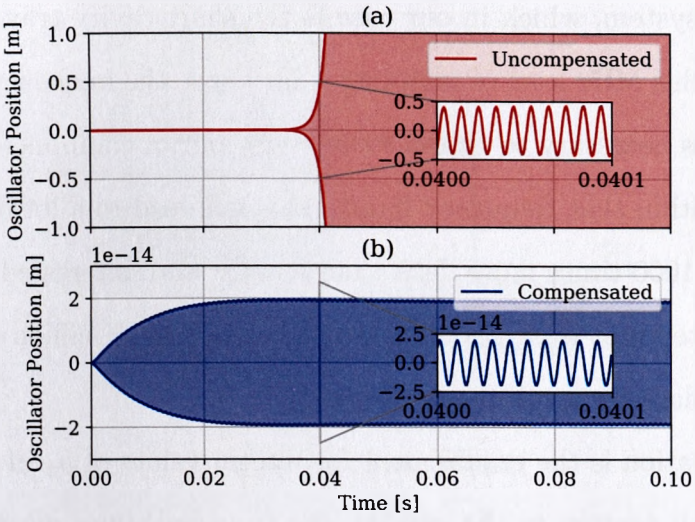


FIGURE 5.12: In-loop step-response of the position of the undamped mechanical oscillator. **(a)** When the pump field is offset by ω_{m_0} only. **(b)** When the pump field is further offset by compensating the optical spring. The lighter solid fill contains rapid oscillations with the dark bounding lines representing the amplitude envelope. The inserts of (a) and (b) show the zoomed-in views of two sections in the time-series, where the individual oscillations can be seen.

Interestingly, the transient behaviour is more of an opportunity than a problem. The time-domain response with custom inputs can be used to gain insight into some aspects of the system that can not be easily studied using the frequency-domain approach. One such test we can perform is to determine the stability of the system using its step response. Inserting a constant signal from $t = 0$ (step-function input) to excite the system dynamics, the resulting time-domain behaviour will clearly show whether the system is stable (settling at a constant value or a steady oscillation) or unstable (growing indefinitely towards an infinite amplitude). As shown in Fig. 5.12, the position response of the mechanical oscillator in the cases with or without compensating the optical spring clearly shows the stability or instability. It further confirms the result shown in Fig. 5.8 that the system can be stabilised by optical spring compensation.

5.4.4 Numerical Noise Spectral Density

To check the consistency with the frequency-domain approach in Section 5.3.2, the noise spectral density in this time-domain simulation will be determined in a two-step approach. The noise power spectral density in the strain signal measurement is given by:

$$S_{hh}^{n}(\Omega) = \frac{S_{bb}^{n}(\Omega)}{T(\Omega)T(\Omega)^{*}}, \tag{5.33}$$

where $S_{bb}^{n}(\Omega)$ is the noise power spectral density of the appropriate quadrature of $\hat{b}_{\text{out}}$ (see Fig. 5.4) and $T(\Omega)$ is the $h \to \hat{b}_{\text{out}}$ transfer function. The strain noise can thus be obtained through two separate simulations: one to obtain $T(\Omega)$ and another to obtain $S_{bb}^{n}(\Omega)$.

For a preliminary consistency check, we compare the optical transfer function $T(\Omega)$ by injecting a noiseless random signal with a constant spectral density into the channel of gravitational wave strain. By measuring the spectrum $S_{bh}(\Omega)$ of the output variable $\hat{b}_{\text{out}}$, the transfer function $T(\Omega)$ can be determined through the relation between cross- and power spectral density:

$$S_{bh}(\Omega) = T(\Omega)S_{hh}(\Omega). \tag{5.34}$$

Two simulations are implemented only with a difference in mechanical quality factor, as shown in Fig. 5.13, and a great agreement has been confirmed between the numerical and analytical approaches. In Fig. 5.14 we further show the optical transfer function of the mechanically undamped system for both the signal and idler channels, which also shows good agreements.

The second simulation is implemented by injecting all relevant noises with their own spectral densities in the absence of strain signal. This provides $S_{bb}^{n}(\Omega)$ and, together with

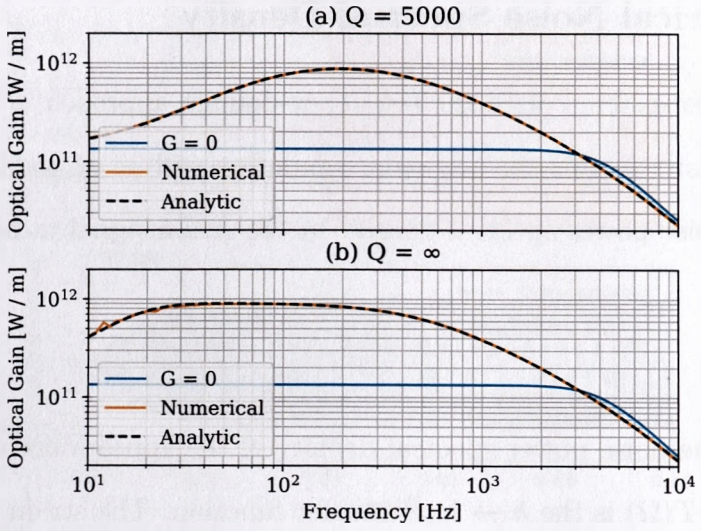


FIGURE 5.13: Transfer functions for optical gain with optical-spring compensated. The pump-off ($G = 0$) gain is shown in contrast to the pump-on gain for **(a)** a finite mechanical quality factor of 5000, and **(b)** an undamped oscillator. The numerical, time-domain approach shows very good agreement with the analytical, frequency-domain approach.

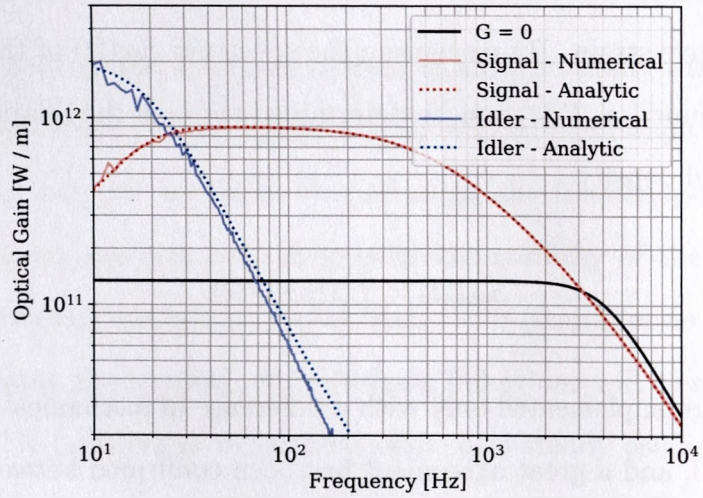


FIGURE 5.14: Optical gain transfer functions for the signal and idler channels in the undamped, optical-spring-compensated system.

the optical transfer function $T(\Omega)$ obtained in the first simulation, allows to complete the calculation of strain noise spectral density according to Eq. 5.33. The final results

of the time-domain simulation are shown in Fig. 5.15, which consistently corresponds to the frequency-domain analysis in Fig. 5.7.

5.5 Optimal readout scheme

In this section, we will investigate the readout scheme to optimally blend the outputs from the signal and idler channels for achieving the best signal sensitivity. Note that the approach we apply here can in general be used for any correlated signals.

Consider a generic case with two noisy signals which measure the same observable x, $\tilde{z_1}$ and $\tilde{z_2}$, such that $\tilde{z_i} = T_i x + \tilde{n_i}$, $i = 1, 2$, with T_i referring to the transfer function of x to the signal, and $\tilde{n_i}$ being the corresponding noise. These signals can be normalised such that $z_i = \tilde{z_i}/T_i = x + n_i$

In this case x is the gravitational wave signal, where in each measurement there is differing sensitivity in different frequency bands such as that shown in Fig. 5.14. In the frequency domain representation, the normalised signals z_i, can be blended together with filters frequency dependent filters, $\alpha(\Omega)$ and $1 - \alpha(\Omega)$, which sum to one over the entire frequency range:

$$Z(\Omega) = \alpha(\Omega)\, z_1(\Omega) + (1 - \alpha(\Omega))\, z_2(\Omega), \tag{5.35}$$

with total noise given by $n(\Omega) = \alpha(\Omega)\, n_1(\Omega) + (1 - \alpha(\Omega))\, n_2(\Omega)$. The power spectrum of the total normalised noise can be represented as:

$$\begin{aligned} S_{nn}(\Omega) =& |\alpha(\Omega)|^2 S_{n_1 n_1}(\Omega) + |1 - \alpha(\Omega)|^2 S_{n_2 n_2}(\Omega) \\ &+ \alpha(\Omega)(1 - \alpha^*(\Omega)) S_{n_1 n_2}(\Omega) \\ &+ (1 - \alpha(\Omega))\alpha^*(\Omega) S_{n_2 n_1}(\Omega), \end{aligned} \tag{5.36}$$

with $S_{n_i n_j}(\Omega)$ being the noise spectrum of $(i = j)$ or correlation between $(i \neq j)$ $n_1(\Omega)$ and $n_2(\Omega)$. Minimising Eq. 5.36 with respect to the conjugate of the filter $\alpha(\Omega)$, i.e. $\frac{\partial S_{nn}(\Omega)}{\partial \alpha^*(\Omega)} = 0$, the optimal filter can be constructed for each frequency Ω independently:

$$\alpha(\Omega) = \frac{S_{n_2 n_2}(\Omega) - S_{n_2 n_1}(\Omega)}{S_{n_1 n_1}(\Omega) + S_{n_2 n_2}(\Omega) - S_{n_1 n_2}(\Omega) - S_{n_2 n_1}(\Omega)}. \tag{5.37}$$

Applying the filters to Eq. 5.35, we can obtain the the power spectrum of the total signal $Z(\Omega)$:

$$\begin{aligned} S_{ZZ}(\Omega) =& |\alpha(\Omega)|^2 S_{z_1 z_1}(\Omega) + |1 - \alpha(\Omega)|^2 S_{z_2 z_2}(\Omega) \\ &+ \alpha(\Omega)(1 - \alpha^*(\Omega)) S_{z_1 z_2}(\Omega) \\ &+ (1 - \alpha(\Omega))\alpha^*(\Omega) S_{z_2 z_1}(\Omega), \end{aligned} \tag{5.38}$$

whose noise power spectrum is the same as $S_{nn}(\Omega)$ in Eq. 5.36. When there is correlation between n_1 and n_2, the noise across the applied high and low pass filters $\alpha(\Omega)$ and $1 - \alpha(\Omega)$ can be further improved. When the cross-correlation terms are zero, the final noise spectrum would be determined by the ability to create the optimal blends between two filters, such as high order roll off of the filters with phase margin at the blending frequency. The advantage of this signal post-processing scheme is to reduce the stability requirement if the system is operating in a closed feedback loop.

In the simulation discussed in Section 5.4, the signal can be extracted from both the signal channel around ω_0 and the idler channel near $\omega_0 + 2\omega_{m_0}$. Applying the optimal blending strategy of Eq. 5.37, the overall noise spectrum can take advantage of whichever channel with the better behaviour at one frequency. As shown in Fig. 5.15, the idler channel contributes more below 25 Hz, while the signal channel dominates at higher frequencies above 25 Hz, and the optimal blending curve follows the better one

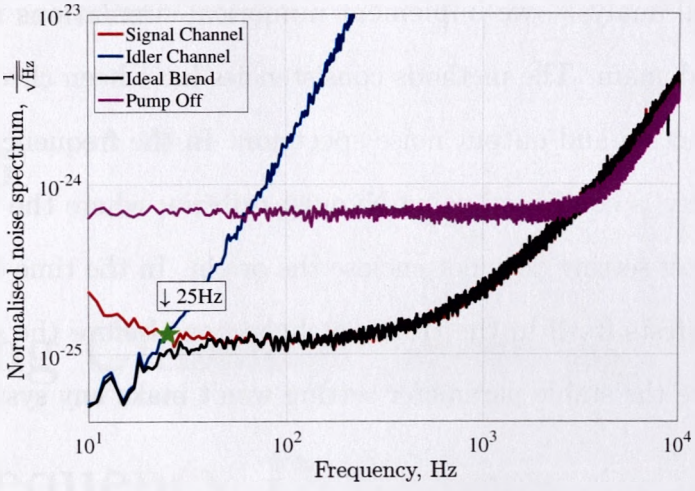


FIGURE 5.15: The normalized noise spectrum for the signal and idler channels. The signal channel (red) recovers the same response as the pump off case (magenta) which is limited by the bandwidth of the cavity at frequencies of a few kHz and above, which shows there is no degradation of the signal or the bandwidth of the cavity. The idler channel (blue) has greater sensitivity below 30Hz. The ideal blend (black) has greater sensitivity than both the signal and idler channels across the blend frequency range of approximately 25 Hz (green star)

at all frequencies.

5.6 Conclusions

In this paper, we analyse in detail the realization of a PT-symmetric interferometer [165] using an optomechanical filter cavity. We go beyond the idealised Hamiltonian analysis under single-mode and the resolved-sideband approximations, and consider the real physical parameters in the filter cavity setup. We prove that, after compensating the optical spring in the filter cavity, the stability and sensitivity improvement stated in the original proposal [165] remains valid in the realistic settings, even when a portion of the signal gets mixed into the idler channel.

To perform full analysis, we implement numerical simulations in both frequency-domain and time-domain. The methods consistencies have been checked with both the signal transfer function and output noise spectrum. In the frequency-domain analysis, the system stability is confirmed using Nyquist criteria, where the Nyquist plot with the stable parameter setting does not enclose the origin. In the time-domain simulation, the stability manifests itself in the transient behaviour before the system reaches its steady-state, where the stable parameter setting won't make any system quantity go to infinity.

Considering the unavoidable leakage of the signal into the idler channel in the actual setups, we further constructed the blending scheme for the output of the two channels. Applying the optimal output filter, one can take advantage of the channel with better behaviour at each frequency, to obtain optimal sensitivity at all frequencies. Another scheme with an auxiliary mechanical cooling beam is under investigation, where the information in the mechanical oscillator can be read out supplementally. Together with the idler channel, they can provide more possibilities for optimized blending readout.

Chapter 6

Modelling Classical Noise in High-Frequency Detectors

The thermal noise models employed in contemporary noise budgets for interferometric detectors are based on work the fundamental insights of [73], which lead to the approach discussed in some detail in Section 2.3. However, the low-frequency validity of the analytic equations that are derived from this approach breaks down for devices with a detection band that approaches close to the internal resonances of the test mass substrate. In this chapter we present work on alternate classical noise models (with a particular focus on substrate thermal noise) extended to frequency regions where the contemporary approach loses its validity. This work is presented in the context of a high-frequency detector design that seeks to detect the gravitational wave signals emitted by the collapse of a post-merger neutron star. This application provides a clear justification for expanding the detection band of GW detectors towards high-frequencies, and thus shows that updated classical noise models that are valid at higher frequencies will become increasingly necessary as the field of GW detection progresses.

The content of this chapter is based on a reformatted copy of the paper *T. Zhang, J. Smetana, Y. Chen et al.* "Toward observing neutron star collapse with gravitational

wave detectors" [170]. *T. Zhang* was the principal author and was responsible for writing the majority of the sections and editing the whole paper. I was responsible for writing Section 6.5, which presents my work on classical noise modelling. The paper provides the necessary context for the discussion of high-frequency classical noise models that is then expanded upon in the supplementary Section 6.6, which contains my original or unpublished work.

6.1 Introduction

On August 17, 2017, the *Advanced LIGO* and *Advanced Virgo* detectors observed a gravitational wave signal, GW170817, from a binary neutron star system [147]. Around 1.7 seconds after the coalescence, the Fermi Gamma-ray Burst Monitor observed a gamma-ray burst, GRB170817A [171], and initiated electromagnetic follow-up observations [172]. This unprecedented observational campaign motivated a series of studies, exploring the implications for astrophysics, dense nuclear matter, gravitation, and cosmology [147]. For example, GW170817 has been used as a 'standard siren' to measure the Hubble constant [173–178]. It is also useful for probing the star equation of state with tidal deformability and mode resonance measurements [179–182].

Despite these groundbreaking multi-messenger observations from GW170817, much is still unknown about what happened in the 1.7 seconds between the merger time when the gravitational wave frequency became too high to be detectable by *Advanced LIGO/Advanced Virgo*, and the occurrence time of the gamma-ray burst, followed by a whole set of electromagnetic transients from the X-ray to the radio band. Generally speaking, a post-merger remnant star may eventually collapse to a black hole and remains as a massive neutron star. Its fate and the timing of the possible collapse

depends on many factors, including the neutron star equation of state (EOS) which is largely unknown [183]; the dissipation and redistribution of angular momentum through gravitational radiation, magnetic field instabilities, and viscous effects; as well as how much thermal support is created in the violent merger, and how quickly the remnant cools due to neutrino emission, etc. [184]. Many of these factors depend on unknown microphysics, as well as complicated, nonlinear, and multiscale dynamics, which are rather preliminary in theoretical exploration. However, what happens in this stage is essential for generating the subsequent electromagnetic transients, so that its description is connected to the answers of many important open problems. For example, even though the multi-messenger observations of GW170817 have provided strong evidence that (at least some) short gamma-ray bursts come from binary neutron star mergers, the underlying burst engine and mechanism are still unknown [185]. Though a leading explanation is that an accretion powered jet is launched by a black hole+disk remnant, there are other proposals that the gamma-ray burst is launched by a magnetar [186–188]. Furthermore, if the post-merger remnant does collapse to a black hole, its lifetime also affects how massive the accretion disk is, how long it is irradiated by neutrino emission from the post-merger star, and the evolution of the magnetic fields. These factors all impact the subsequent electromagnetic signals, including the synthesis of r-process material and the associated kilonovae transient [189, 190]. On the other hand, assuming the collapse does happen, a measurement of its timing may be combined with the electromagnetic observations to constrain existing electromagnetic transient models, or to motivate new ones. The timing information can also be used to constrain the neutron star EOS [191, 192], especially in the regime of higher densities and temperatures than in the ones in the inspiral stage. Therefore, it is interesting to explore the design of a detector that is capable of identifying the prompt or delayed collapse of binary neutron

star merger remnant. As we shall see, this task is particularly challenging, as the frequencies of the collapse gravitational wave signals are generally higher than the main post-merger signals, and the amplitudes are weaker.

So far, searches for post-neutron star merger GW signals have used limited modelling of the physics, and yielded only upper bounds on potential signals. This is due both to the theoretical uncertainties in the post-merger waveform, and to the limited sensitivity of current observatories at these high frequencies. In particular, following GW170817, searches were performed on *Advanced LIGO/Advanced Virgo* data using either generic unmodeled algorithms, or by assuming a simple power-law spin-down waveform, but were at least a factor of 40 less sensitive than required for the theoretically expected signal [193].

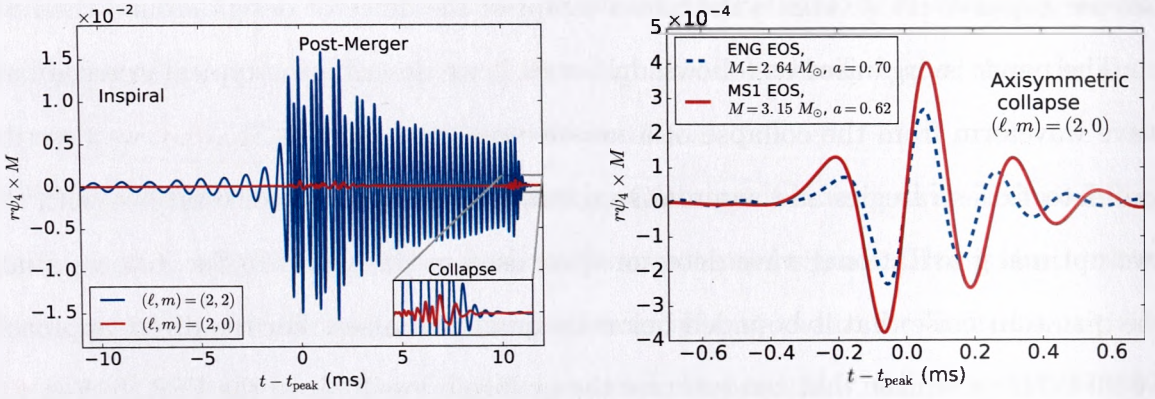
Here, we propose two different methods of identifying or constraining the timing of the collapse. In the indirect approach, we assume that we already have an accurate estimate for the post-merger waveform of hypermassive/supramassive neutron stars, so that we may infer the post-merger collapse from a non-detection of the post-merger signal with sufficient sensitivity. In the direct approach, we search for the collapse signal within the post-merger data using the collapse waveform, with a successful detection naturally giving rise to the timing information. We discuss the performance of both methods with the detector designs explored in this work. The detector design for indirect approach follows the conventional idea in Ref. [149, 194], which takes the coupled arm cavity and signal recycling cavity at 2 kHz–3 kHz. For the direct approach, the detector is a narrow band detector targeting on exploring the sensitivity around the frequency of arm cavity free spectral range (FSR). It is largely different from the current existing and future planned broadband gravitational wave detectors, *e.g.* *Advanced LIGO* [13], *Advanced Virgo* [14], *KAGRA* [15], *Einstein Telescope* [90, 195, 196] and

Cosmic Explorer [155]. This is the first attempt of the detector design around the FSR.

The paper is organised as follows: in Sec. 6.2, we describe the typical gravitational wave waveform from the collapse of a massive neutron star. In Sec. 6.3, we describe two detection strategies and summarise the resulting signal-to-noise rations with the two optimal gravitational wave detectors proposed in this work; In Sec. 6.4, we study the quantum noise that is bounded below by quantum losses. For the direct approach, we introduce a scheme that can saturate the quantum loss limit at the FSR. In Sec. 6.5, we discuss the classical noises in the two detectors.

6.2 Collapse Waveform

Depending on its total mass, a binary neutron star system may promptly collapse to a black hole after merger, or form a long-lived post-merger remnant that only collapses after redistributing/losing angular momentum to gravitational wave emission and magneto-vicious effects and/or radiative cooling [198]. A third possibility is that the neutron star EOS can support a stable remnant. At the onset of the collapse, the remnant star may still have an oscillating quadrupole moment due to its binary origin, so that the collapse process is not fully axisymmetric. As a result, we expect both the $\ell = 2,\ |m| = 2$ (or "22") and $\ell = 2,\ m = 0$ (or "20") spherical harmonics to contribute to the gravitational wave emission. The 22 part of the waveform mainly describes the continuous evolution and damping of the remnant quadrupole moment during the collapse process, and the 20 component comes from the collapsing, angular-averaged remnant star, which is non-spherical due to its spin. Right after the binary neutron star collision, the quadrupole moment emission of the remnant should play a significant role in the collapse waveform; at late times the remnant quadrupole moments damp



(a) (b)

FIGURE 6.1: Panel (a): An example binary neutron star merger waveform from Ref. [197] (in particular the Newman-Penrose scalar ψ_4 from the simulation with the HB EOS). In this case the post-merger remnant collapses to a black hole after $\sim 10\,\text{ms}$. There is an $\ell = 2$, $m = 0$ gravitational wave component associated with the collapse (inset), though at this point it is still very small compared to the $\ell = 2$, $m = 2$ component which dominates throughout. Panel (b): The $\ell = 2$, $m = 0$ component of the waveform from the collapse of unstable, uniformly rotating neutron stars. We show two examples with the "ENG" and "MS1" EOSs and different star masses. In the latter case, the $m = 0$ component is about two-thirds the amplitude found in the left panel. We also find that the pre- and post-peak portions of the collapse waveform contribute roughly equally to the total power, and have similar frequencies.

out, in which case the collapse waveform should be dominated by the 20 piece. This is illustrated in the left panel of Fig. 6.1, where we show these two contributions for an sample binary neutron star merger simulation from Ref. [197]. In this case, the collapse happens roughly 10 ms after merger, and the 22 mode contribution is still much larger than that of the 20. The former has been well studied in numerous cases using binary neutron star simulations [199, 200], while the latter has been less well-studied since it is typically small. In order to characterize the lower bound on the expected collapse gravitational wave signal—occurring when the remnant is able to completely reach rigid rotation and cool prior to forming a black hole—we construct several unstable equilibrium solutions describing uniformly rotating neutron stars, and simulate their

collapse. We show the resulting gravitational waveforms for two examples with masses and spins chosen to match that expected from a binary merger and different EOSs in the right panel of Fig. 6.1. These simulations are carried out with the code described in Ref. [201], using the same methods as in Ref. [202] (in particular, the equilibrium star solutions are constructed using the RNS code [203]). We find that the amplitude of the 20 mode is not significantly smaller (within a factor of 2) from the merger case, and is characterized by the black hole quasinormal mode ringing, with comparable contribution from the pre- and post-peak power. The gravitational waves are also similar (after normalizing by total mass) to those found in Refs. [202, 204], which considered the (in some cases induced) collapse of lower mass rotating neutron stars with various EOSs, indicating these results are relatively insensitive to the particular EOS, etc. We will use this to guide the waveform model described below.

We start by describing the post-merger waveform model mathematically. Despite the uncertainties in the neutron star's EOS and the complicated magneto-hydrodynamical processes governing the post-merger dynamics, it is generally accepted that the peak mode is one of the most robust features in the post-merger waveform [205–207]. This peak mode is associated with the emission from the rotating quadrupole moment of the remnant star, which usually dominates the waveform spectra after the first several milliseconds following the collision.

We phenomenologically model the peak mode signal as a decaying sinusoid for a source at distance d:

$$h = A\frac{50\mathrm{Mpc}}{d}\sin(2\pi f t)e^{-\pi f t/Q}e^{-\pi f \tau/Q}\,, \tag{6.1}$$

where the amplitude A, mode frequency f, and quality factor Q all depend on the star mass and EOS, and we take the convention that the waveform amplitude is specified

for an observer on axis. Here, the delay time τ is the time separation between the star merger and the collapse, i.e., the collapse time. This is the part of the post-merger waveform that we would have detected if there were no collapse happening. If we pick a typical EOS "TM1" and choose an equal-mass binary with total mass $M = 2.7M_\odot$, the parameters take the values

$$A \approx 2.5 \times 10^{-22}, \quad f \approx 2.8\,\mathrm{kHz}, \quad Q \approx 34\,. \tag{6.2}$$

We shall use this set of values as "canonical" values for later demonstration. Values for more EOSs can be found in Ref. [208]. For example, the amplitude changes by about a factor of two for the 5 EOSs considered in Ref. [208], ranging from the stiff side to the soft side of the EOS spectrum.

The collapse waveform contains two parts: the 22 mode and 20 mode. As the post-collapse waveform is very close to black hole ringdown, the 22 and 20 mode frequencies and damping rates are well approximated by black hole quasinormal mode frequencies. For a black hole with mass $\sim 2.7M_\odot$ and dimensionless spin $a \sim 0.7$ (typical for a post-collapse black hole), the corresponding mode frequencies and quality factors are

$$\begin{aligned} &f_{22} \sim 6.4\,\mathrm{kHz}, \quad f_{20} \sim 4.7\,\mathrm{kHz} \\ &Q_{22} \sim 3.2, \quad Q_{20} \sim 2.5\,. \end{aligned} \tag{6.3}$$

If we allow a range of spins 0.65–0.80 for prompt [209–212] and delayed collapses [213], and $\sim$ 20% variation between the mass thresholds of prompt and uniform-rotation collapses [214,215], the variation in the quasinormal frequency is approximately 30%, and the variation of quality factor is $\leq$ 20% [216].

Based on the discussion above, we shall take as a representative value for the

amplitude of the 20 mode $A_{20} \sim 0.8 \times 10^{-23}$, though it is reasonable to expect at least a factor of 2 variation depending on the specific EOS and star mass. On the other hand, if we make the observation that the amplitude of ψ_4 for the 22 mode is roughly continuous before and after the collapse [217], we can estimate A_{22} as $A_{22} \sim A(f/f_{22})^2$. Notice that this expression may underestimate the 22 mode amplitude if the collapse happens in the early post-merger phase ($\leq 20\,\mathrm{ms}$), where the peak mode only constitutes a sub-dominant fraction of the post-merger spectrum[1]. Other short-lived modes may also deform the quadrupole moments and affect the post-collapse 22 mode.

In summary, the post-collapse 22 waveform is

$$h_{22} = A_{22}\frac{50\mathrm{Mpc}}{d}\sin(2\pi f_{22}t)e^{-\pi f_{22}t/Q_{22}}e^{-\pi f\tau/Q}\,. \tag{6.4}$$

The collapse 20 waveform is

$$h_{20} = A_{20}\frac{50\mathrm{Mpc}}{d}\sin(2\pi f_{20}|t|)e^{-\pi f_{20}|t|/Q_{20}}\,. \tag{6.5}$$

We note that the axisymmetric part of the star is not sensitive to the collapse time, so that in this model, h_{20} is assumed to be independent of τ.

Finally, we mention that these simple types of waveforms are similar to those used currently in searches, e.g. of black hole ringdown or unmodelled bursts, and so could be easily utilized for that purpose. However, in this study will not attempt to quantify how different the physical waveform would be from the simple model described here, nor what the loss in sensitivity would be from modelling systematics. Instead, in the remainder we will take the above collapse waveforms as exact.

[1]The signal-to-noise ratio of the peak mode is roughly 1/2–1/3 of the signal-to-noise ratio of the whole post-merger waveform [149]

6.3 Detection strategies

In order to determine the collapse time, there are at least two detection strategies based on the measurement of gravitational waves.

The first method ("indirect" approach), assumes that we can accurately predict the post-merger waveform based on the inspiral measurement. Such an assumption relies on future advances in theoretical modelling of the post-merger processes and gravitational-wave measurements of the star EOS. If this assumption is satisfied, starting from any delay time τ, we can compare the inferred waveform signature with the actual waveform. If with enough confidence we no longer detect the inferred emission from the hypermassive neutron star starting from time τ, the remnant star must have collapsed.

The second method ("direct" approach) is to directly identify the signal of the collapse waveform in the entire post-merger process, either the 20 component or the 22 component, or both. The threshold signal-to-noise ratio (SNR) depends on the time width of the search window and the requirement on the false alarm rate. As the duration of the collapse waveform is on the order of milliseconds, successful detection with the direct method will be able to pin down the collapse time to $\mathcal{O}(\text{ms})$ accuracy.

For both strategies, we calculate the achievable SNR of a detector for the three types of waveform mentioned earlier: the post-merger peak mode, the collapse 20 mode and the collapse 22 mode. The detector strain sensitivity to a particular source in the sky relies on the antenna response, $F_{+/\times}$, which is a function of source location and gravitational wave signal frequency. A detailed calculation of the antenna response can be found in Ref. [218]. When the gravitational wave signal frequency is much smaller than the arm cavity FSR, the long wavelength approximation is applicable, under which the square root of the antenna power pattern, $\sqrt{|F_+|^2 + |F_\times|^2}$, is maximal and equals to 1 for gravitational waves at normal incidence, *i.e.* the polar angle $\theta = 0$.

Considering a long baseline facility, important features of the antenna response appear around the frequency of the arm cavity FSR. Fig. 6.2 shows the square root of the antenna power pattern, $\sqrt{|F_+|^2 + |F_\times|^2}$ at frequency of the arm cavity FSR. Here, $\sqrt{|F_+|^2 + |F_\times|^2}$ is 0 for $\theta = 0$, while the maximal response is around 0.25 for a signal source at $\theta = 45°, \phi = 45°$. The maximal SNR is characterised as

$$\mathrm{SNR} = 2\sqrt{\int \frac{|h(f)|^2 |(|F_+|^2 + |F_\times|^2)|_{\max}}{S_{hh}} df}\,, \tag{6.6}$$

where S_{hh} is the signal-referred power spectral density for the noise. In the following discussion, the SNR is calculated for the binary sources at 50 Mpc and with total mass $2.7 M_\odot$. The "canonical" parameters are listed in Eq. 6.2 and Eq. 6.3.

In Fig. 6.3, we give two detector designs, for the indirect and direct approaches respectively. For the former, the high-frequency (HF) detector design follows the one proposed in Ref. [149], which assumes the current existing technologies but with 4 MW arm power and 10 dB observed squeezing. As it turns out, the optimal length for measuring the peak mode of the post merger waveform is around 20 km, which gives SNR of 2.3 for the peak mode with $\tau = 10$ ms as shown in Fig. 6.4. This is consistent with the finding in Ref. [149], depending on the variation in the quasinormal frequency within 30% and variation in the quality factor within 20%. The resulting SNR can vary from 1.0 to 4.3. And allowing a factor of 2 variation of waveform amplitude, SNR can vary from 0.5 to 8.6. However, the SNR for 20 or 22 mode of collapse waveform from this design are far from allowing a direct detection of the collapse.

For the direct approach, an alternative design making use of the detector sensitivity at the first FSR turns out to be a better option. For such an FSR detector design, to achieve an SNR ~ 1, we have to assume 10 MW arm power and 10 dB observed squeezing.

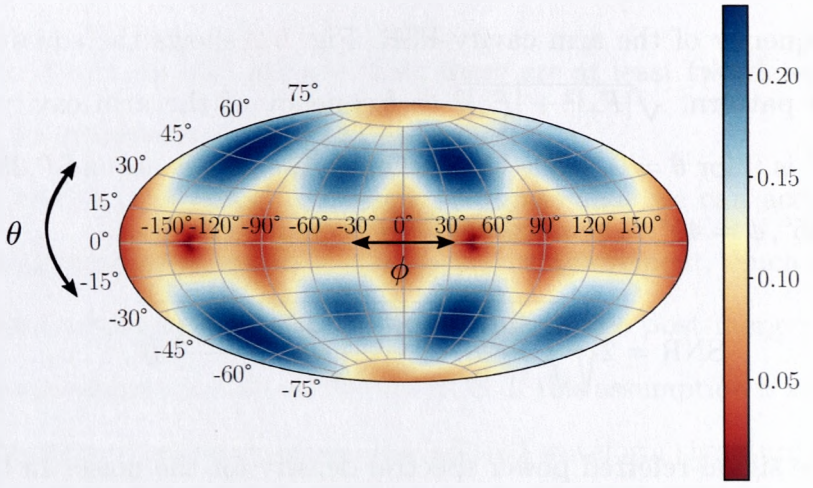


FIGURE 6.2: The square root of the antenna power pattern of a gravitational wave detector for signals from sky location described by θ and ϕ at frequency of the arm cavity FSR.

We also implement a new quantum technique based upon the active optomechanical filter, which allows us to achieve a loss-limited quantum sensitivity around the FSR. The classical noises are estimated based on proposed technology for future gravitational wave detectors, *e.g. Voyager* [92], *Einstein Telescope* [90,195,196] and *Cosmic Explorer* [155]. As shown in Fig. 6.4, the optimal length is 34 km, giving an SNR of 0.8 for the 20 mode of the canonical collapse waveform. We also show that by considering the variation in the quasinormal frequency and the quality factor, within 30% and 20% respectively, the SNR vary from around 0.6 to around 0.95. Taking into account a factor of 2 variation of the waveform amplitude, the FSR detector gives constant SNR for the 20 mode in the range of 0.3–1.9. The SNR for the 22 mode with $\tau = 10$ ms is small, unless a short collapse time is expected. In fact, the SNR of measuring the peak mode of the post-merge waveform with the indirect appraoch is expect to increase with a shorter collapse time. As shown in Fig. 6.5, in which the shaded areas indicate the boundary

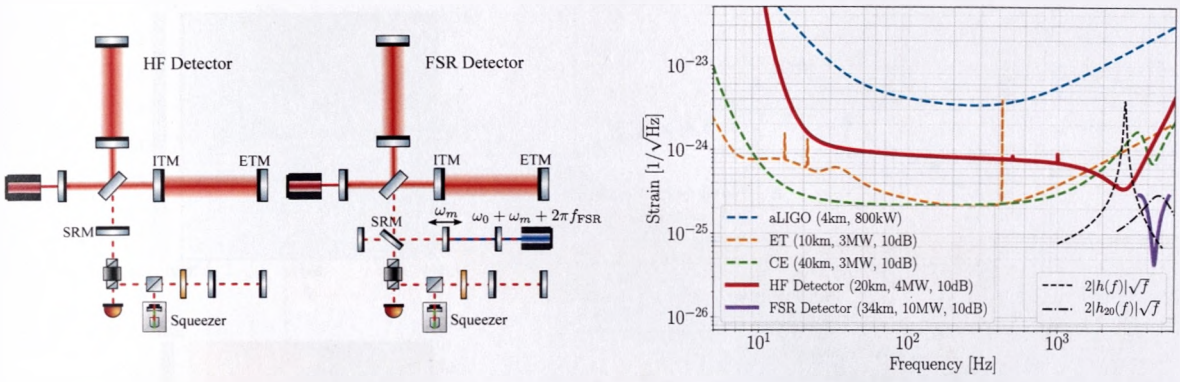


FIGURE 6.3: The left part of the figure shows two gravitational wave detector configurations, the 20 km HF detector and the 34 km FSR detector, which are optimal for detecting the peak mode of the post-merger waveform and the 20 mode of the collapse waveform. The right part of the figure shows the strain sensitivity of the two configurations in comparison with *Advanced LIGO*, Einstein Telescope and Cosmic Explorer. The strain sensitivity of the FSR detector is for sources at $\theta = 45°, \phi = 45°$, where the antenna response is at a maximum. The arm cavity powers of the HF and FSR detectors are 4 MW and 10 MW, respectively. The assumed observed squeezing levels are 10 dB. The other detector parameters are shown in Table 6.1. The amplitudes of the peak mode and 20 mode with parameters listed in Eq. 6.2 and 6.3 are shown as dashed and dot-dashed curves, respectively.

of the SNR taking into account variation in quasinormal frequency, quality factor and amplitude, the HF detector is preferred for the indirect approach when $\tau < 16$ ms.

In the following sections, we will discuss the quantum noise and classical noise of the two detector designs.

6.4 Quantum noise and quantum limit from optical losses at the FSR

At high frequencies, the gravitational wave detector sensitivity is critically limited by quantum noise. The quantum noise consists of quantum radiation pressure noise and quantum shot noise. The former, which results from the optomechanical interaction between the light and the mirror, dominates at low frequencies. Above the interferometer

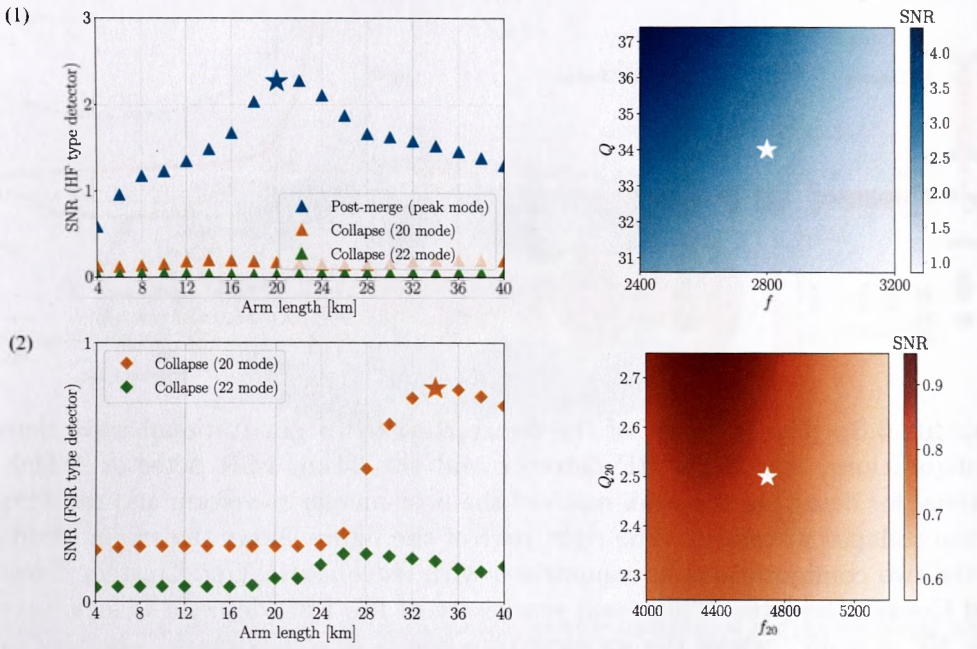


FIGURE 6.4: (1) Left: The SNR of the HF type detector for the peak mode of the post-merger waveform, and the 20 and 22 modes of the collapse waveform as a function of arm length. The waveform parameters are given as the "canonical" values in Eq. 6.2 and 6.3 and a delay time of 10 ms. The maximal SNR (2.3) for the peak mode is at around 20 km marked with the star, which is chosen for the arm length of the HF detector. The other detector parameters are listed in Table. 6.1; Right: With the HF detector, the SNR for the peak mode of the post-merge waveform which has a 30% variation in the qausinormal frequency f and a 20% variation in the quality factor. The white star corresponds to the star in the left part. (2) Left: The SNR of the FSR type detector for the 20 and 22 modes of the collapse waveform as a function of arm length. The maximal SNR (0.8) for the 20 mode is at around 34 km marked with the star, which is chosen for the arm length of the FSR detector. The waveform parameters are given as the "canonical" values in Eq. 6.3 and a delay time of 10 ms. The other detector parameters are listed in Table. 6.1; Right: With the FSR detector, the SNR of the 20 mode of the collapse waveform which has a 30% variation in the qausinormal frequency f_{20} and a 20% variation in the quality factor. The white star corresponds to the star in the left part.

bandwidth, the quantum shot noise dominates. Given the arm cavity power, the quantum shot noise limited sensitivity can also be further tailored by choosing parameters of the signal recycling cavity (SRC), *i.e.* the cavity length and finesse, which shapes the detector response to gravitational waves. By tuning the SRC parameters, we can

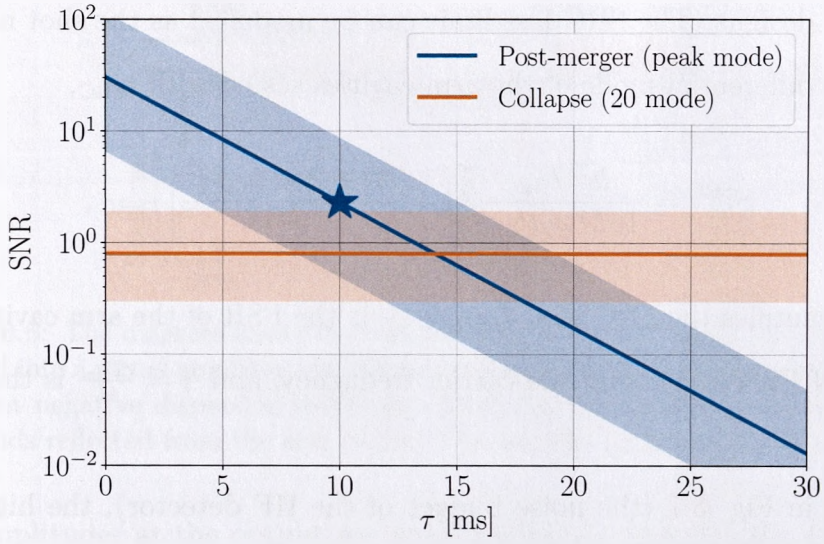


FIGURE 6.5: The SNR of the HF detector to peak mode of the post-merge waveform and FSR detector to the 20 mode of the collapse waveform as a function of delay time. The solid lines are for waveform parameters given in Eq. 6.2 and 6.3. The shaded areas give boundaries of minimal and maximal SNR taking into account the variation in the qausinormal frequency, the quality factor and a factor of 2 variation of the amplitudes of waveforms. At $\tau = 10\,\mathrm{ms}$, the SNR for peak mode is in the range of 0.5–8.6, for 20 mode is in the range of 0.3–1.9.

optimise the quantum sensitivity for the figure of merit at each arm length. Further reduction of the quantum noise can be realised by using squeezed light.

However, the quantum noise cannot be infinitely suppressed but is limited by a lower bound from the optical losses, especially including arm cavity loss and SRC loss [219]. Furthermore, the external output loss will limit the observed squeezing level. We assume 10 dB of observed squeezing, which implies a maximum of 10% output loss.

The arm cavity loss limit can be calculated as

$$S_{hh}^{\mathrm{arm}} = \frac{\hbar c^2}{4L^2\omega_0 P_{\mathrm{arm}}}\epsilon_{\mathrm{arm}}\,, \tag{6.7}$$

where c is the speed of light, $\hbar$ is the Plank constant, L is arm length, P_{arm} is the arm circulating power and ϵ_{arm} quantifies the arm cavity loss. $\epsilon_{\mathrm{arm}} = 100\,\mathrm{ppm}$ is assumed in

both detector designs. The SRC loss limit can be modelled as the shot noise spectral density of the differential mode of the arm cavities scaled with ϵ_{SRC},

$$S_{hh}^{\mathrm{SRC}} = \frac{\hbar c^2 T_{\mathrm{itm}}}{16 L^2 \omega_0 P_{\mathrm{arm}}} \left[1 + \frac{(\Omega - 2n\pi f_{\mathrm{FSR}})^2}{\gamma^2}\right] \epsilon_{\mathrm{SRC}}\,, \tag{6.8}$$

where ϵ_{SRC} quantifies the SRC loss, $f_{\mathrm{FSR}} = \frac{c}{2L}$ is the FSR of the arm cavity, n denotes the number of FSR away from the carrier frequency, and $\gamma = \frac{cT_{\mathrm{itm}}}{4L}$ is the arm cavity bandwidth.

As shown in Fig. 6.7 (the noise budget of the HF detector), the high frequency quantum noise dip is limited by the SRC loss, which is assumed to be 1000 ppm. The dip around 3 kHz is shaped through the coupled cavity resonance between the arm cavity and the SRC in the same way as in Ref. [149]. At the dip frequency, the sidebands resonate in the SRC, i.e. the accumulated phase reflecting from the arm cavity and traveling through a round trip in a long SRC equals to 2π. The resonance frequency is approximately, $\frac{c\sqrt{T_{\mathrm{itm}}}}{2\sqrt{L L_{\mathrm{SRC}}}}$, where L_{SRC} is the SRC length. The depth of the dip is determined by the SRC finesse. The choices of parameters for the HF detector is shown in Table. 6.1.

Towards higher frequencies, the SRC loss limit increases until Ω approaches the first FSR, where the noise returns to the level at zero frequency according to Eq. 6.8. To enable the best achievable sensitivity, an addition scheme which allows for saturating the loss limit is required. As shown in Ref. [220], a proper arrangement of an optomechanical filter cavity allows the quantum noise to saturate the loss limit by enhancing the signals inside the SRC, without sacrificing the detector bandwidth defined by the arm cavity. The mechanism functions by cancelling the phase gained by the signal sidebands reflecting from the arm cavity through a negative dispersion from the optomechanical filter cavity. Therefore, the signal sidebands will resonate in the SRC at all frequencies

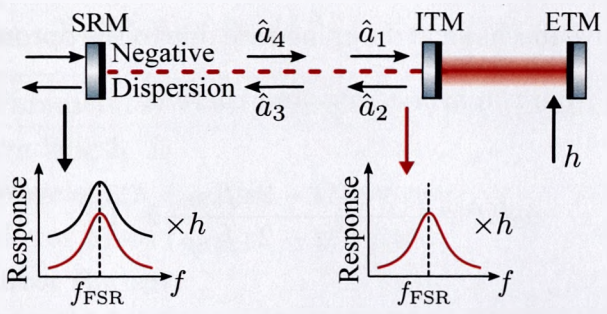


FIGURE 6.6: The diagram shows the coupled arm cavity and SRC. The SRC can be converted into a signal amplifier around the arm cavity FSR frequency with frequency dependent negative dispersion inside for cancelling the accumulated phase of the sidebands reflected from the arm cavity. The amplification factor is $1/\sqrt{2T_{\mathrm{SRM}}}$.

and their amplitudes at the output are amplified by a constant gain, $1/\sqrt{T_{\mathrm{SRM}}}$ with T_{SRM} being the power transmissivity of the signal recycling mirror.

As we will illustrate, a scheme, similar to the one considered in Ref. [220], can also work in our case for saturating the SRC loss limit around the FSR frequency. As shown in Fig. 6.3, an optomechanical filter cavity is placed inside the signal recycling cavity. This filter cavity is pumped with a laser at frequency $\omega_0 + \omega_m + 2\pi f_{\mathrm{FSR}}$, where ω_m is the mechanical resonant frequency of the mechanical oscillator. Compared to Ref. [220], the pump laser frequency is shifted by one FSR. When there is

$$\Omega - 2\pi f_{\mathrm{FSR}} \ll \gamma_f \ll \omega_m \,, \tag{6.9}$$

which is the so-called resolved-sideband regime, the input/output relation of the optomechanical filter for the sideband at Ω is

$$\hat{a}_4(\Omega) \approx -\frac{\gamma_{\mathrm{opt}} - i(\Omega - 2\pi f_{\mathrm{FSR}})}{\gamma_{\mathrm{opt}} + i(\Omega - 2\pi f_{\mathrm{FSR}})}\hat{a}_3(\Omega) \,, \tag{6.10}$$

where $\gamma_{\rm opt}$ is the negative mechanical damping rate due to the optomechanical interaction. The sidebands reflect from the arm cavity and there is

$$\hat{a}_2(\Omega) \approx \frac{\gamma + i(\Omega - 2\pi f_{\rm FSR})}{\gamma - i(\Omega - 2\pi f_{\rm FSR})}\hat{a}_1(\Omega)\,, \tag{6.11}$$

When $\gamma_{\rm opt} = \gamma$, the phases from reflecting off the arm cavity and the optomechanical filter cavity cancel out. Therefore, the signal sideband at Ω becomes resonant, and the amplitude is amplified by $1/\sqrt{T_{\rm SRM}}$. However, the sideband at $-\Omega$ does not acquire the negative dispersion according to Eq. 6.10. Therefore, when measuring the combination of sidebands around $\pm f_{\rm FSR}$ using homodyne detection, we will get a factor of $\sqrt{2}$ degradation of the signal amplitude, similar to the detuned SRC case. The resulting quantum noise spectral density around the frequency of arm cavity FSR can be calculated as

$$S_{hh}^{\rm FSR} = \frac{\hbar c^2 T_{\rm itm} T_{\rm SRM} e^{-2r}}{8L^2\omega_0 P_{\rm arm}}\left[1 + \frac{(\Omega - 2n\pi f_{\rm FSR})^2}{\gamma^2}\right] + S_{hh}^{\rm SRC} + S_{hh}^{\rm arm}\,, \tag{6.12}$$

where $e^{-2r} = 0.1$ represents 10 dB of observed squeezing. In the design of the FSR detector, we adopt $T_{\rm SRM} = 0.0015$, with 10 dB squeezing and the assumption of 1000 ppm SRC loss. The quantum noise amplitude result is a factor of $\sqrt{1.3}$ greater than the SRC loss limit. The total quantum noise is shown in Fig. 6.8. The other detector parameters are listed in Table. 6.1.

TABLE 6.1: Parameters of the detector

Parameters	"HF Detector"	"FSR Detector"
Arm length, L	20 km	34 km
Wavelength	1064 nm	1550 nm
Mirror Mass	100 kg	500 kg
Mirror Radius	0.24 m	0.34 m
Mirror thickness	0.26 m	0.6 m
ITM/ETM RoC	14 000 m	28 900 m
Arm circulating power, P_{arm}	4 MW	10 MW
ITM transmittivity	0.06	0.06
SRM transmittivity	0.1	0.0015
SRC length	100 m	50 m
SRC loss	1000 ppm	1000 ppm
Arm cavity loss	100 ppm	100 ppm
Squeezing level	10 dB	10 dB

6.5 Classical noises

In terms of fundamental classical noises in a gravitational wave detector, the main contributions at kHz frequencies will be from residual gas phase noise and mirror thermal noises.

6.5.1 Residual gas noise

Even though we operate in high vacuum, there will still be residual gas pressure. The statistical variations of multiple molecular species disturb the laser's phase when they pass across the beam [221]. To calculate this noise, we take the vacuum requirement designed for the Einstein Telescope and specify the residual gas contents, H_20, H_2, N_2, with partial pressures of 2.5×10^{-9}, 5×10^{-9}, 2.5×10^{-10} [90,195,196] respectively, for both the HF and FSR detectors. Our strategy resembles very closely the implementation in the noise modelling tool GWINC (Gravitational Wave Interferometer Noise Calculator) [222],

which uses the approach of Ref. [223]. We simply reinstated the numerical integration of Eq. 1 in Ref. [223] rather than maintaining the first order approximation, which strongly diverges above 1 kHz. Our treatment does pose a greater computational burden but is necessary to accurately model the residual gas noise.

6.5.2 Mirror thermal noise

The mirror thermal noise consists of coating and substrate brownian thermal noise [224], which, according to the fluctuation-dissipation theorem [61, 62, 225], come from the thermal fluctuations resulting from the mechanical dissipation in the substrate and the multiple layers of high and low index coatings [226, 227]. They are proportional to the temperature and inversely proportional to the beam size on the test masses.

In the HF detector, the sensitivity above 1 kHz is only limited by the quantum shot noise. This allows us to use the existing *Advanced LIGO* technology. Thus the HF detector is designed to be operated at room temperature, with a 1064 nm laser, fused silica substrates and Ta_2O_5/SiO_2 coatings. Since the thermal noise is not limiting, a relatively smaller beam size allows smaller and lighter test masses. In our case, we choose a 100 kg mirror with mirror radius of curve (RoC), radius and thickness as shown in Table. 6.1, which are determined by satisfying a beam clipping loss of $\sim$ 1 ppm (The ratio of mirror size and beam size is $\sim$2.63).

In the FSR detector, the thermal noise will limit the sensitivity. We design the FSR detector at a temperature of 120 K, with a 1550 nm laser, silicon substrates and amorphous silicon (a-Si) / silicon dioxide (SiO_2) optical coatings. It is worthwhile to discuss the calculation of substrate Brownian noise. Via the fluctuation-dissipation theorem, the substrate Brownian noise is related to the mechanical conductance of the substrate itself. As an extended body, the mirror's mechanical conductance is

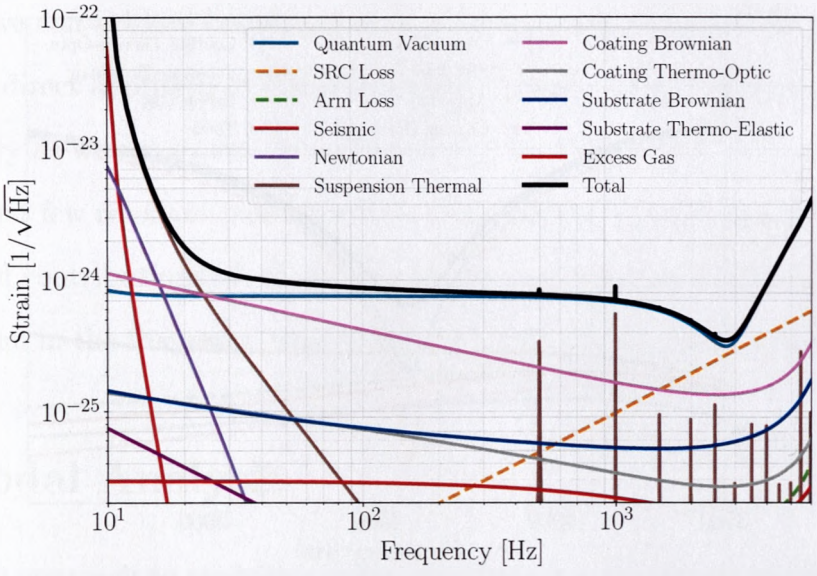


FIGURE 6.7: Noise Budget of the HF detector. The detector strain sensitivity is for optimal orientated sources at $\theta = 0°, \phi = 0°$. The detector parameters are listed in Table. 6.1.

affected by its vibrational eigenmodes, which affect the shape of the substrate surface that the laser impinges upon. The vibrational eigenfrequencies lie in the vicinity of a few kHz and the lowest frequency is inversely proportional to the mirror size. In the current operating and proposed detectors, the detector's sensitivity in the kHz range is limited by quantum shot noise. In such a case it is common to employ the approach of Ref. [68, 75, 228]. However, this is not the case for the FSR detector, since the quantum noise in the FSR detector is around or below $10^{-25}/\sqrt{\text{Hz}}$ around 4 kHz–5 kHz and the thermal noise around the substrate resonances is not negligible. It is thus imperative that we model the contribution of each vibrational eigenmode individually, as we are operating in a regime where the low frequency approximation breaks down. This is further exacerbated by us increasing the mirror size in the long baseline detector, which lowers the resonant frequencies yet further for a given fixed aspect ratio. Thus, to minimise the substrate thermal noise, we focus primarily on shifting the resonances

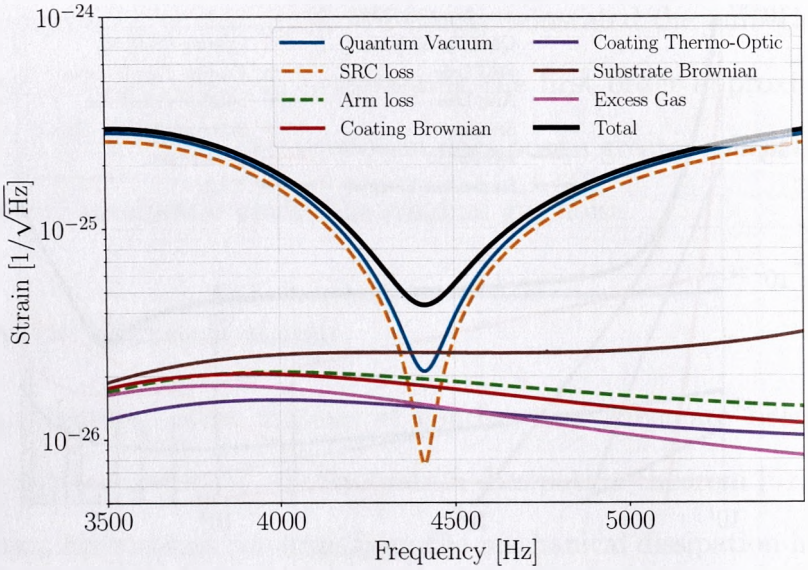


FIGURE 6.8: Noise Budget of the FSR detector. The detector strain sensitivity is for optimal orientated sources at $\theta = 45°, \phi = 45°$. The detector parameters are listed in Table. 6.1.

to higher frequencies. We achieve this in part by our choice of aspect ratio, which we empirically determine by a finite element method (FEM) to approximately maximise the frequency of the first resonant mode. Consequently, we adopt an aspect ratio of 1.75 in our work. Further optimisation can be achieved by minimising the mirror size (to raise the resonant frequencies) while guaranteeing beam clipping loss of no more than 1 ppm. The adopted mirror geometry parameters are listed in Table. 6.1. With our optimised mirror design for the FSR detector, we shift all frequencies to be above 6 kHz.

To accurately determine the substrate Brownian noise, we require a method that models the noise with reference to the individual contributions of each vibrational eigenmode. Therefore, we implement the approach of [72], utilising our FEM to determine the mode frequencies and the corresponding surface deformation. This technique is known to be more computationally intensive and, with the first 100 modes

considered, we can achieve convergence to approximately 80% of the noise amplitude given by the direct approach at low frequencies. However, a higher level of convergence is unnecessary as we are interested in the region around the first resonance where the tails of the first few resonant peaks begin to dominate the substrate thermal noise above the combined contributions of all subsequent modes. For a sufficient modelling of the noise spectrum in the frequency band of interest, only the first few modes are necessary.

6.6 Modal Analysis

The usual approach to modelling substrate thermal noise utilises the direct approach pioneered by *Yu. Levin* and follows from the analysis discussed in some detail in Section 2.3.2. However, the analytic result that is ultimately derived [73–75] is only valid for frequencies sufficiently below the first internal resonant mode of the test masses. Due to the combined effect of heavier mirrors and a shift towards higher frequency detection, the internal resonances of the substrate begin to encroach on the detection band. Consequently, this low-frequency approximation ceases to be valid. In particular, the FSR detector would operate in a narrow band very close to the forest of internal resonances and thus care must be taken in tuning the mirror parameters to avoid coinciding with a resonant peak. To this end, we employ a finite element model (FEM) to determine the parameters of the mirror resonances, such as their mode shape and corresponding elastic potential energy. The modal analysis is then combined with the approach of [72] to obtain the thermal noise contribution.

Each internal mode can be modelled as a point-mass harmonic oscillator, which follows the standard and well-understood dynamics explored in Section 2.2. Via the balanced, oscillating exchange between elastic potential energy and kinetic energy,

we can equate the two energy terms at their corresponding maximum values, $U_{\max} = \frac{1}{2} m_n {\omega_n}^2 {x_{\max}}^2$, where the RHS represents the maximum kinetic energy. This equation can be adapted to model the point-mass approximation to the extended body, yielding [72]

$$U_{n,\max} = \frac{1}{2} \alpha_n m {\omega_n}^2 {\Delta l_n}^2, \tag{6.13}$$

where m is now the mass of the real substrate, which is constant for each mode, a_n is the effective mass coefficient and Δl_n is the apparent displacement. The terms a_n and Δl_n are convenient for writing the equation in a simple and familiar form. However, they contain the complexity which allows us to reduce the three-dimensional phenomenon into a one-dimensional model. Thus the most challenging part of the analysis is determining these parameters accurately.

The apparent displacement, Δl_n, contains within its definition the interaction between the probe beam and the mode shape of the particular substrate resonance. The displacement of the test mass is inferred from a power fluctuation of the laser beam caused by a phase shift applied by the moving substrate. The beam has an extended spatial intensity distribution, which means that the power fluctuation is naturally lesser—and thus the inferred (apparent) displacement should be lesser—if the phase shift to the field is applied in a region where the beam is less intense. This results in the definition [72]

$$\Delta l_n = \int_S f(\rho, \theta) u_k(\rho, \theta) \mathrm{d}\sigma, \tag{6.14}$$

where f is the intensity distribution of the beam profile normalised over the substrate surface ($f \propto \exp(-\rho^2/w^2)$ for a gaussian profile with a beam radius w at the substrate surface) and u_k is the local displacement of the substrate surface in the direction of the local wavevector, $\boldsymbol{k}$. This can be intuitively understood as stating that the complex,

distributed displacement of the vibrating surface of the extended substrate is identical, from the beam's perspective, to a uniform translation by Δl_n of a rigid body.

To fully connect the behaviour of an extended body to the point-mass equivalent, we must introduce the effective mass coefficient, α_n. A particular mode shape may only produce a slight apparent displacement or even no displacement at all, as is the case for all modes with a non-zero angular degree. As a result, a slight apparent displacement may be the product of a highly energetic oscillation. To make this small displacement consistent with the large amount of stored energy, we may increase the *effective* mass of the mirror in this particular vibrational mode, which is related to the true mass by $m_{\text{eff}} = \alpha_n m$.

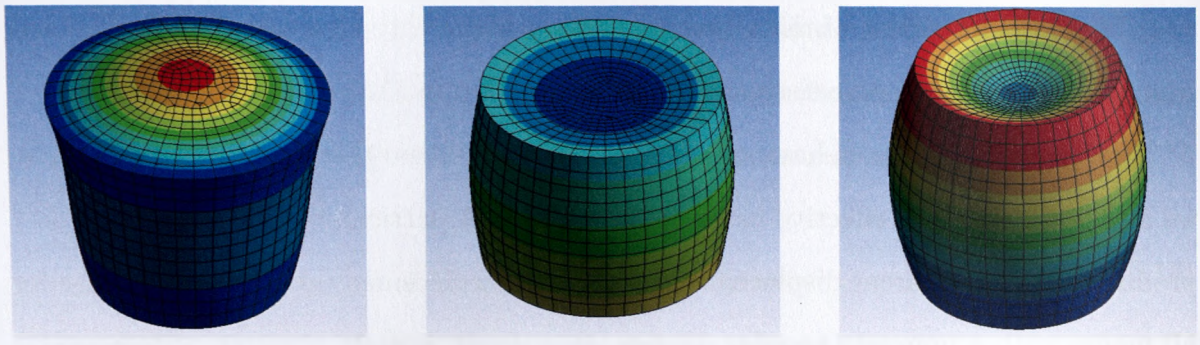
The total thermal noise PSD in displacement can then be obtained as the sum of the thermal noise contribution from each mode independently, given by

$$S_{x,n}(\Omega) = \frac{4k_B T}{\alpha_n m \Omega} \left(\frac{{\omega_n}^2 \phi_n(\Omega)}{(\Omega^2 - {\omega_n}^2)^2 + {\omega_n}^4 {\phi_n}^2(\Omega)} \right), \tag{6.15}$$

which can be derived with reference to the fluctuation-dissipation theorem (Section 2.3). Due to the simplicity of the point-mass oscillator treatment, the equation governing the thermal noise can be written very succinctly. The challenge is in determining the key parameters, ω_n, α_n and Δl_n, particular to each mode.

Although analytical solutions to some aspects of the system exist, for example, the frequencies of the cylinder's internal resonances are known [229], we employed a finite element analysis technique to obtain all three parameters. We used the FEA software Ansys to perform the modal analysis. Ultimately, the analysis was not just important for calculating the thermal noise but also to somewhat optimise the parameters of the detector, such as the diameter and thickness of the substrate, and the beam radius.

The first aspect of the optimisation we performed was on the optimal geometry of



(a) 15.3 kHz (b) 15.5 kHz (c) 17.4 kHz

FIGURE 6.9: The shape of the three lowest-frequency modes with non-zero beam overlap for an optimised aspect ratio of 1.75. Modes (a) and (c) are approximately analogous to the fundamental mode of a vibrating, circular membrane, with the oscillations of the central anti-node being in-phase and in opposition for the two modes respectively. Mode (b) corresponds to a radially symmetric expansion and contraction along the cylinder's axis of symmetry. The resonant frequencies specified under each mode are for a test mass radius of 10 cm.

the test mass. As can be seen from the behaviour of Eq. 6.15, there is a very rapid rise in the thermal noise around the resonances at ω_n. Due to this effect, it is generally much more important to carefully tune the position of these resonances to be far out of the detection band even if this comes at a cost of higher losses elsewhere. The resonances of the substrate are altered through changes in both the thickness and radius of the test mass but they do not shift uniformly with a change in either quantity in isolation as one resonance may reduce whilst another increases. Instead, we may parametrise the geometry in terms of the test mass radius, governing the overall length-scale of the test mass, and the unit-less aspect ratio, given by the ratio of the thickness to the radius. We know that upon scaling all mirror dimensions by a common scale factor, a, the resonances all shift uniformly by the factor $1/a$. Therefore, we can simplify the optimisation as we have determined that all resonances shift inversely with the radius for a constant aspect ratio.

Consequently, the optimisation only needs to be performed on the aspect ratio.

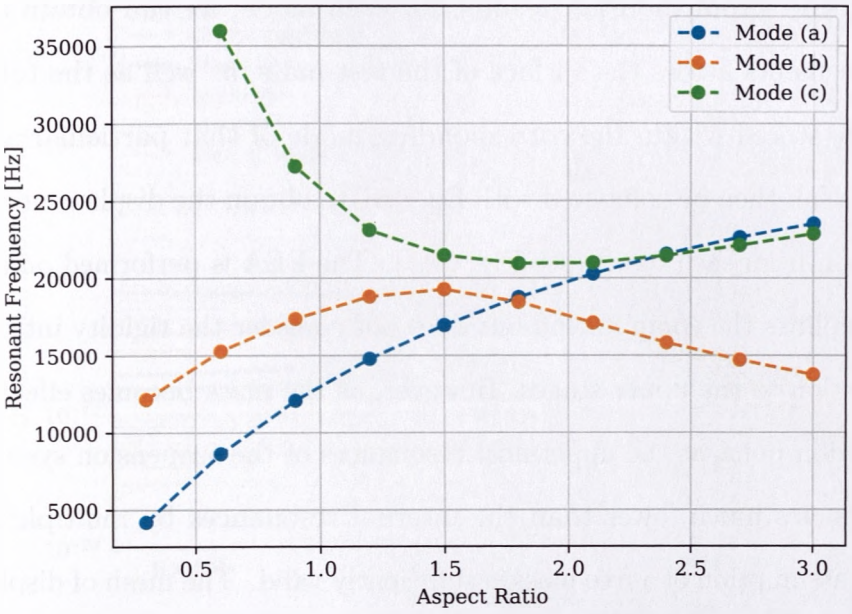


FIGURE 6.10: The resonant frequencies of the three modes shown in Fig. 6.9 over a range of aspect ratios. Each data point was obtained through the FEA modal analysis performed individually for each new substrate geometry defined by the aspect ratio. Due to the time-consuming nature of this process, the parameter space is rather sparsely sampled. Despite this limitation, we can determine the key result, which is the approximate aspect ratio of 1.75 at which the two lowest resonant frequencies coincide.

Within the reasonable range of aspect ratios, three modes compete for the lowest frequency, and these are shown in Fig. 6.9. We varied the aspect ratio to determine how these resonances change, as shown in Fig. 6.10. Of particular interest is that Mode (a) is the lowest frequency mode for small aspect ratios and decreases in frequency with decreasing aspect ratios, whilst Mode (b) is the lowest frequency mode for large aspect ratios and whose frequency decreases with increasing aspect ratio. This results in an optimal aspect ratio where these two resonances match, thus maximising the lowest resonance. This happens at approximately an aspect ratio of 1.75, which matches the value we proposed in Section 6.5.

Once the aspect ratio has been optimised we can turn to the calculation of the thermal noise itself. From the FEA results, for each mode, we can obtain a mesh of sampled displacements across the surface of the test mass, as well as the total elastic potential energy stored within the corresponding mode of that particular amplitude. These quantities can then be combined with Eq. 6.15 to obtain the displacement thermal noise contribution from each mode (see Fig. 6.11). The FEA is performed on a free test mass, which simplifies the computation but does not consider the rigidity introduced by the fibre connection to the upper stages. However, as the mass becomes effectively free beyond a transition point at the uppermost resonances of the suspension system, which occur at frequencies much lower than the internal resonances by multiple orders of magnitude, the assumption of a free mass is sufficiently valid. The mesh of displacements provides us with a numerical sampling of the function u_k in Eq. 6.14. By assuming a standard gaussian beam profile, we can then numerically determine the value of Δl_n using Eq.6.14. The elastic potential energy obtained from the FEA gives us the value of $U_{n,\mathrm{max}}$ in Eq. 6.13, which we can rearrange to give us α_n. Therefore, we have obtained the two key unknowns in Eq. 6.15 and thus we can numerically compute the thermal noise curve for different parameter choices, and produce the substrate thermal noise model presented in Fig. 6.8. This is a much more computationally intensive process than the direct approach, as this process involves performing the integral in Eq. 6.14 for every change in the substrate or beam parameters. In the worst case, if the geometry of the substrate is fundamentally changed, the FEA needs to be performed again, making the process substantially more time-consuming.

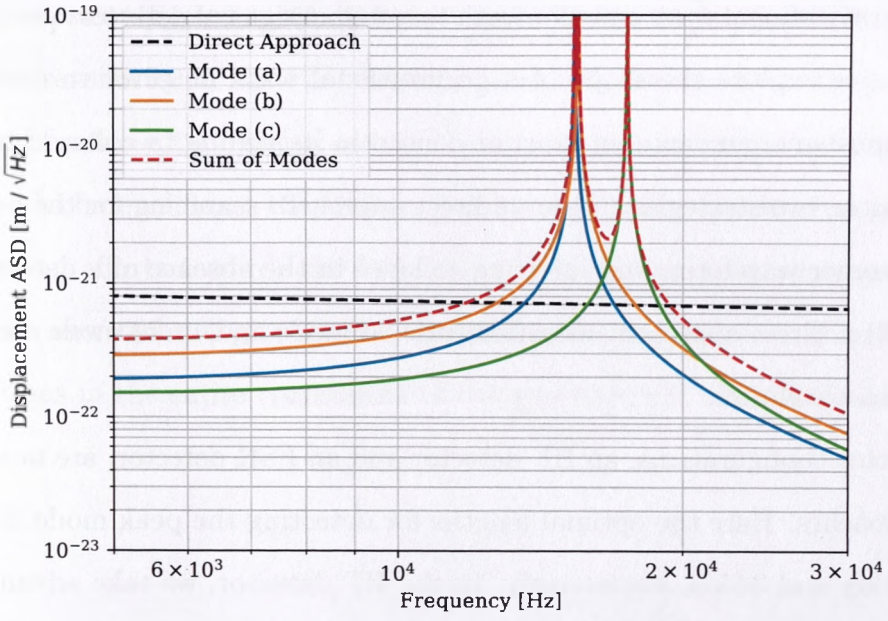


FIGURE 6.11: Thermal noise in displacement contributed by the three highlighted modes in Fig. 6.9. The red dashed line shows the total thermal noise from the three modes considered and the black dashed line shows the result of Eq. 2.45 applied to this system. The total thermal noise from just the first three modes is already approaching the analytic model in black at low frequencies. The convergence between the two thermal noise approaches is likely much closer than the figure suggests due to the infinite half-plane model used in the creation of the black trace, which has the effect of overestimating the true thermal noise. Around and after the substrate resonances, the modal analysis captures the effect of the resonant peaks on the thermal noise which the direct approach misses completely.

6.7 Summary

In the era of multi-messenger gravitational-wave astronomy, we now have unprecedented opportunities to better understand the strong field gravitational dynamics and complicated astrophysical processes of a binary neutron star merger, with gravitational wave observations and a slew of electromagnetic transients that subsequently arise from the post-merger remnant. The nature of the post-merger remnant and the lifetime it survives are the key factors for answering the existing open questions, such as the engine

of short gamma-ray bursts. Previous works have discussed possible ways to probe the post-merger gravitational wave emission with targeted designs of detectors [148,149,194]. In this work, we explore the scope of a gravitational wave detector to measure the lifetime of a post-merger remnant star, or constrain its timing to collapse to a black hole. We propose two strategies: (1) an indirect approach: searching for the peak mode of the post merger waveform, and inferring collapse in the absence of a detection after some time; (2) a direct approach: measuring the 20 mode and/or 22 mode components of the collapse waveform.

Two detector configurations, an HF detector and an FSR detector, are modelled for the two approaches. Here the optimal lengths for detecting the peak mode and the 20 mode are 20 km and 34 km, respectively. In the HF detector, we take advantage of a coupled arm cavity and SRC resonance to boost the sensitivity at 2 kHz–3 kHz, which is now limited by the SRC loss. In the FSR detector, our strategy is to make use of the detector sensitivity at the first FSR, which leads to a long baseline design. To achieve a loss-limited quantum sensitivity at the FSR frequency, we use an optomechanical detector with negative dispersion to cancel the phase of the sideband gained upon reflection from the arm cavity. This converts the SRC into a signal amplifier. The FSR detector is eventually limited by both quantum and classical noise, in particular, residual gas noise and mirror thermal noise. The frequency dependent antenna response is properly calculated, without using the long wavelength approximation, which is not applicable for our long baseline detector.

While the NS collapse would be essentially undetectable with current gravitational wave detectors, with these optimized detector designs, we realise a SNR of 0.5–8.6 in the indirect approach and a SNR of 0.3–1.9 for directly detecting the 20 mode of the collapse waveform. In particular, it seems that a SNR of 1.9 in the direct approach

is probably not significant enough to claim detection, as compared to the threshold SNR $\sim$ 5 commonly used in black hole ringdown studies [230]. In the current noise budget the sensitivity of the FSR detector is already limited by both quantum losses and several fundamental classical noises, *i.e.* the mirror thermal noises and the residual gas noise which depends on the facility. Further improving the detector sensitivity not only relies on efforts towards the laser power enhancement for shot-noise reduction, but also on new technologies for the classical noises suppression. These are ongoing investigations in the entire community of the gravitational-wave detection.

Chapter 7

Compact Interferometry for Precision Displacement Measurement

The suppression of seismic noise in contemporary GW detectors is achieved through a multi-stage process including both active control and passive isolation. The active feedback relies on a network of low-noise inertial sensors, the quality of whose readout is essential for the precise control of the active platforms that support the test masses. Interferometric displacement sensors have the ability to improve the readout of these inertial sensors and, generally any position readout in the broader seismic control scheme. In this chapter we present an optically and mechanically simple and compact interferometric displacement sensor with multi-wavelength sensing range.

The content of this chapter is a reformatted copy of the paper *J. Smetana, R. Walters, S. Bauchinger et al.* "Compact Michelson interferometers with subpicometer sensitivity" [231] of which I was the principal author. I wrote the initial draft and edited the paper with editorial contribution from *D. Martynov.* The paper presents the experimental results of laboratory work conducted collaboratively by *R. Walters, S. Bauchinger, D. Martynov* and myself.

7.1 Introduction

Interferometric displacement readout has played a pivotal role in expanding the boundaries of precision measurement. In the field of GW astronomy, interferometric readout had shown its utility in detecting GW signals [16, 140, 141]. The cornerstone of its success, in part, lies in two areas: (a) the exceptional sensitivity achievable over distances significantly smaller than the wavelength of light used and (b) versatility in design across vastly different length scales. To the first point, contemporary GW detectors such as *Advanced LIGO* [13] and *Advanced Virgo* [14] can measure some of the weakest signals in the universe in the form of GW strain. In particular, *Advanced LIGO* is capable of measuring displacements as fine as $2 \times 10^{-20}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ in the most sensitive frequency region around 100 Hz [17]. To the latter point, interferometric measurement sees widespread use in investigations of quantum phenomena but also spans km-scales in terrestrial GW detection and million-km-scales in *LISA*—the future space-based GW observatory [33].

We present the investigation of a compact, interferometric sensor for the precise measurement of picometer-scale displacements of macroscopic objects. Displacement sensing of gram-scale optics through to kilogram-scale test masses is widespread in the GW community due to the strict requirements on suppression of terrestrial displacement noise (such as seismic noise) in ground-based GW detectors. Displacement sensing and control has been achieved through numerous means, such as capacitive position sensors [13], coil-readout geophones such as the Sercel L-4C and optical shadow sensors such as the BOSEM [50]. The displacement sensitivity of these devices is limited across a broad band by the readout noise of these sensors. When used in feedback control, the sensitivity of these instruments imposes a fundamental floor on the level of suppression that is achievable. For example, *Advanced LIGO* is limited in its sensitivity below

30 Hz by controls noise [17], which can be improved with more sensitive devices. The astrophysical case for expanding sensitivity towards lower frequency in contemporary GW detectors is broad, ranging from improved early-warning systems [232] for multi-messenger detections [233] to the routine observation of heavier astronomical objects such as intermediate-mass black holes [234].

The future role of interferometric displacement sensors in achieving the sensitivity improvements necessary for observing these astrophysical phenomena has already been explored in detail in [88]. The 5 Hz–30 Hz frequency band in *Advanced LIGO* is dominated by a range of technical noises, such as angular controls noise caused by low frequency (<1 Hz) ground vibrations (see [17,88] and references therein). Inertial sensors with an interferometric readout, such as 6D seismometers [56] and next generation beam rotation sensors [235,236], have the potential to significantly reduce the test mass motion, simplify the lock acquisition process [237], and reduce the required bandwidth of angular control loops. Furthermore, interferometric sensors can efficiently damp the suspensions without limiting the design sensitivity of GW detectors. The sensors achieve the required improvement of sensing noise by a factor of 100 compared to the currently utilised shadow sensors [50,88].

Outside the field of GW detection, compact interferometry has seen use across a wide range of applications. It has been applied in the readout of torsion balance experiments [236], which can be used in investigations into areas of fundamental physics, such as short-range force measurements or bosonic dark matter searches [238]. Compact interferometers have also been used as reference sensors in the development of other displacement measuring devices [239]. The potential for compact interferometers to serve as absolute distance measurement devices [240] has led to their use in alignment sensing at the *Large Hadron Collider* [241].

The history of compact interferometer development reaches back several decades and includes a broad range of designs, topologies and varying levels of sophistication based on the technology available at the time. A comprehensive review of compact interferometry can be found in [242] and expounds all of the different strategies undertaken in this sphere. Compact interferometry is not a recent concept as exemplified by the device by de la Rue et al. [243] dating from 1972. However, the design of a high performance compact interferometer is non-trivial and often falls into a particular performance niche. The design can be tuned for high linearity, in one case 5 pm, which was achieved by Weichert et al. [244]. Other designs are tuned for high sensitivity within the desired frequency range, such as the *LISA Pathfinder* interferometer, achieving a sensitivity of $1\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at 10 mHz [245].

The term 'compact' in reference to these devices can, however, be rather misleading depending on the application. Many of the existing compact interferometers still occupy an area of hundreds of square-centimeters, which is unacceptably large for applications where space is at a premium, such as the sensing of the *Advanced LIGO* quadruple suspension [42]. Some devices conform to the label of 'compact' better than others, for example the design by Royer et al. [246] is confined to an area of only 8 cm × 5 cm. Further development has led to interferometers designed with the direct application to seismic displacement sensing in mind. An example of this, the HoQI sensor [57], has shown substantial improvements in the sensitivity of commercial geophones by replacing the electronic coil readout and thus eliminating the dominant source of noise [247].

The high sensitivity is in part due to the very sharp response of the interferometer, being able to cover the full signal range in a mere quarter-wavelength of displacement. A typical value of 266 nm for a simple Michelson interferometer, is significantly smaller than that achievable by alternate technology, such as the 0.7 mm linear range of the

BOSEM [50]. The cost of this sharp response is a significant reduction in range, limiting the use of many interferometric devices to specialised uses over sub-wavelength range. However, we can implement multi-quadrature readout techniques [248], coupled with a fringe-counting algorithm, to expand the range over multiple interference fringes (read [242] for an overview of the many multi-fringe readout techniques available). These techniques maintain the precision of the instrument, whilst overcoming its most significant limitation.

In this paper, we build a compact displacement sensor out of commercially available products, such SmarAct C01 sensing heads, NKT Adjustik laser, and Thorlabs fibre beam splitters. We implement a multi-quadrature readout scheme with fringe counting in the digital domain, which can be readily compared to the analogue-domain equivalent demonstrated by numerous other multi-fringe interferometers. The scheme makes use of a purely digital implementation of the laser frequency modulation, subsequent signal demodulation, filtering and arm phase extraction. The key strengths of this design are that the physical system is highly compact, mechanically simple, robust and easy to deploy where the necessary computational infrastructure is already present. We use only commercial components, with the primary component, the sensing head, occupying an area of merely $1.3\,\mathrm{cm} \times 0.4\,\mathrm{cm}$ (excluding any necessary mounting brackets). We achieve a high level of sensitivity due to an appropriate choice of laser and through additional frequency stabilisation. In Section 7.2 we briefly outline the principles of deep frequency modulation and present the experimental layout. Section 7.3 is devoted to the experimental results of our fixed-mirror measurement with the accompanying noise budget and the subsequent improvements we achieved using a frequency-stabilising feedback system.

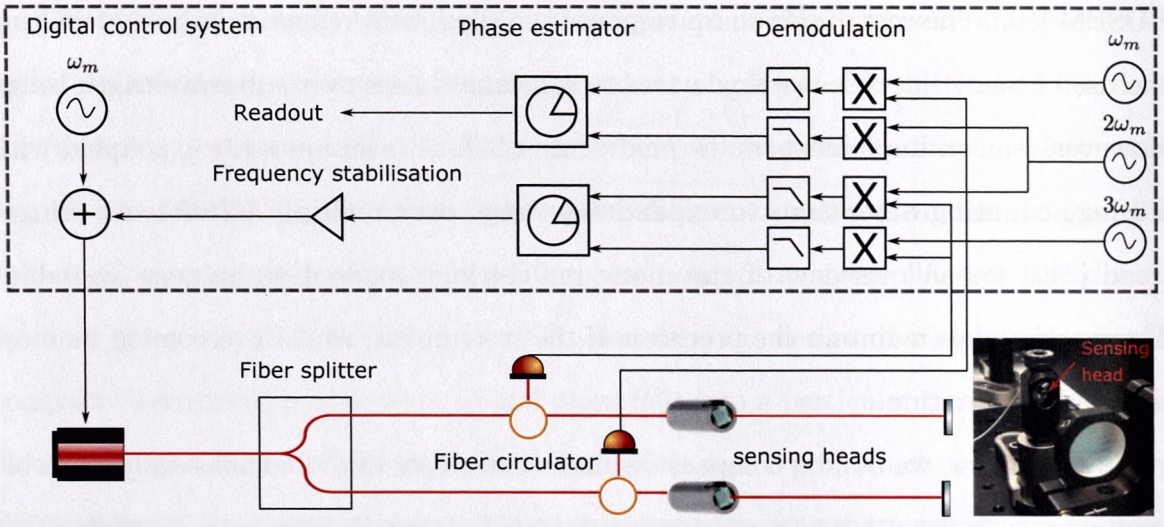


FIGURE 7.1: Layout of the experiment and of the digital signal processing. A digitally sourced modulation is applied to the laser, which is then connected to both sensing heads. The raw DC diode signal is digitised and demodulated in the digital domain at multiple harmonics. An algorithm is applied to extract the displacement from the resulting signals.

7.2 Optical layout

The scheme for achieving our measurement is shown in Fig. 7.1. A single laser is connected to two separate sensing heads mounted on separate base plates. Each sensing head comprises a simple Michelson interferometer with an internal, fixed-length reference arm and an open port, which forms the external measurement arm. Each measurement beam is then reflected from a fixed mirror mounted on the respective base plate and aligned to maximise the fringe contrast. The interference pattern is detected by a photodiode with a responsivity of $0.9\,\mathrm{A\,W^{-1}}$. It is mounted on the output of a fibre circulator which connects together the output of the laser, the sensing head and the photodiode. The signal is passed through a $68\,\mathrm{k\Omega}$ transimpedance amplifier followed by a 7 kHz low-pass anti-aliasing filter. It is then digitised at a sampling rate of 64 kHz by a 20 V peak-to-peak analogue-to-digital converter (ADC). The signal is

down-sampled to 32 kHz once digitised for processing.

We modulate the laser frequency at $\omega_m = 2\pi \times 700\,\text{Hz}$ and amplitude of $A_m = 0.8\,\text{GHz}$ with an internal PZT actuator. The power observed at the photodetector is given by the equation [249]

$$P(t) = A[1 + C\cos(\phi(t) + m\cos(\omega_m t))], \tag{7.1}$$

where A is the power amplitude, C is the fringe contrast, $m = 4\pi A_m \Delta L/c$ is the unit-less modulation index, ΔL is the longitudinal imbalance of the Michelson arms and c is the speed of light. The variable arm phase, ϕ, is related to the test mass x displacement and laser wavelength, λ, via $\phi = 4\pi x/\lambda$. The signal in Eq. 7.1 can be expressed in terms of its Fourier series expansion [249]:

$$P(t) = P_0 + \sum_{n=1}^{\infty} 2CAJ_n(m)\cos(\phi + n\pi/2)\cos(n\omega_m t), \tag{7.2}$$

where J_n is the n^{th} order Bessel function of the first kind and $P_0 = A(1+CJ_0(m)\cos(\phi))$. The sum in Eq. 7.2 shows that the input modulation at ω_m manifests in the output spectrum at ω_m and all higher order harmonics. Furthermore, displacements of the test mass become up-converted to sidebands around each harmonic of ω_m.

The measured signal passes through a transimpedance amplifier and the raw signal is then digitised. The procedure gives us access to the raw power signal (up to the input anti-aliasing filter's cutoff frequency at 7 kHz) as given in Eq. 7.1. The extraction of the signal from the modulation harmonics relies on accurate demodulation, which in turn relies on a local oscillator that corresponds to the modulation source with high fidelity. The key advantage of the digital scheme is that the local oscillator is simultaneously used for the signal demodulation and laser modulation, which leads to a form of digital

homodyne readout. This compensates for timing jitter that could otherwise introduce a discrepancy between the modulation source and the local oscillator.

This signal is then demodulated at multiple harmonics of the input modulation frequency by multiplying the raw input signal by $\sin(k\omega_m t + \psi_k)$ and $\cos(k\omega_m t + \psi_k)$ for $k \in \{1, 2, 3\}$. The demodulation phase, ψ_k, represents the phase difference between the modulation signal at its generation and at the point of digitisation together with any time-varying fluctuations such as jitter noise. We tune ψ_k to maximise the signal in the cosine quadrature and reduce the timing fluctuations to a second order effect. The output is given by the equation

$$S_k = CAJ_k(m)\cos(\phi + k\pi/2) \tag{7.3}$$

We take the ratio of two harmonics to construct our elliptical Lissajous figure. We derive the arm phase for the first sensing head from the ratio of the first two signals, which is given by

$$\frac{S_1}{S_2} = \frac{J_1(m)}{J_2(m)}\tan(\phi). \tag{7.4}$$

To simplify the process further, we tune the modulation depth, A_m, to achieve the ratio of the Bessel functions in Eq. 7.4 equal to unity (circular Lissajous figure). This leads to improvements in linearity of the measurement as well as preventing the introduction of non-stationary noise. Since we can tune A_m only for one sensing head, we tuned the length between the second sensing head and the measured mirror to achieve the ratio of the Bessel functions for the second and third harmonics of the sensing head close to unity. The arm phase is extracted in real time through Eq. 7.4. Through digital demodulation, we obtain signals S_1 and S_2 in real time. We then apply the four-quadrant arctangent function to this pair of values per sample on the assumption

that the Lissajous figure is circular. Fluctuations and drifts in the parameters of the Lissajous figure introduce non-linear effects in the real-time phase readout. We analysed these effects by applying an ellipse fitting algorithm in post-processing to compare to the performance of the real-time phase extraction scheme. This comparison is addressed in Section 7.3.4. A more complex real-time phase extraction scheme that does not rely on fixed values of the ellipse parameters could later be adopted if more sensing heads need to be probed with various longitudinal imbalances or drifts become an issue [249].

7.3 Experimental results

The sensitivity analysis of our interferometric sensing heads was performed on a fixed test mass in order to effectively suppress the presence of ground motion and 'zero out' the displacement signal. The observed sensitivity to displacement is shown in Fig. 7.2b. The achieved sensitivity is $3 \times 10^{-13}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 1 Hz and is a factor of 300 better compared to the BOSEM, one of the currently used sensors in *Advanced LIGO* [50]. In this section, we discuss the noise analysis of the measured sensitivity.

7.3.1 Environmental disturbances

We mitigated the perturbation of the laser beams via air currents, which otherwise form the dominant noise source at frequencies below 1 Hz. We suppressed these effects by placing the components into a sealed, foam-padded box. The box has reduced the air currents by three orders of magnitude around 1 Hz. However, the noise is still seen below 100 mHz in Fig. 7.2b. At these frequencies, our insulation becomes less efficient due to the air pressure fluctuations in the laboratory. We expect that in-vacuum operation of the sensing heads will further suppress the noise below 100 mHz.

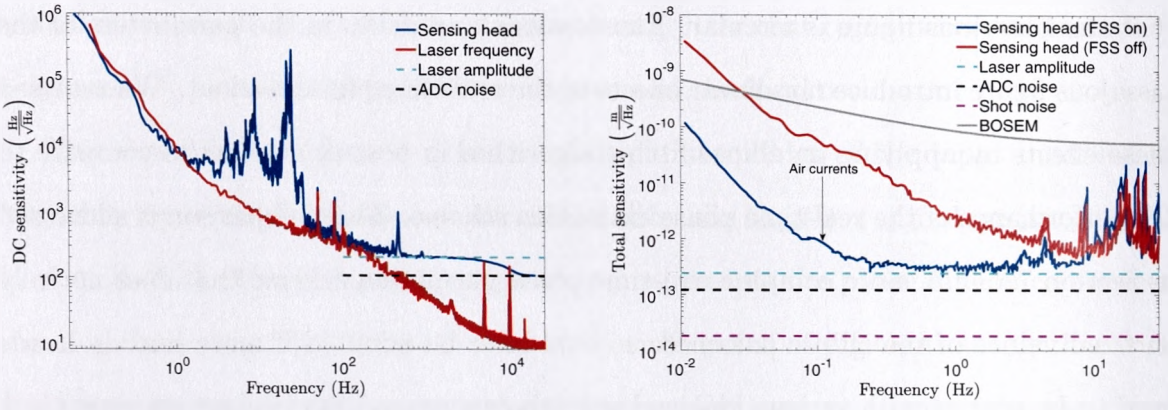


FIGURE 7.2: Panel (a) shows the DC sensitivity of the sensing heads without the modulation-demodulation process (calibration factor is $2.85 \times 10^{15}\,\text{Hz}\,\text{m}^{-1}$ in the setup). The DC sensitivity is compared with the laser frequency noise that was independently measured by beating two identical lasers. The two curves are consistent at low frequencies (below 1 Hz). The dashed lines represent the inferred noise levels from the two largest noise sources near the modulation frequency: laser amplitude noise and ADC noise. The former noise is suppressed above 7 kHz by the anti-aliasing filter. The laser frequency noise is below the level of $100\,\text{Hz}/\sqrt{\text{Hz}}$ and does not limit the spectrum around the modulation frequency. Panel (b) shows the total sensitivity of the sensing head with the heterodyne readout and frequency noise stabilisation (FSS). The FSS-on curve represents the extend of the sensitivity increase we achieve with the frequency stabilisation. The curve is limited by the residual air currents below 100 mHz. The dashed lines represent the noise limits of several prominent, incoherent noise sources and show that our new noise floor is consistent with the laser amplitude noise limit at 100 mHz–10 Hz. The BOSEM sensitivity curve [50] allows for a comparison between the sensitivity of the sensing head and contemporary sensing devices.

Seismic noise dominates the measured spectrum at 10 Hz–100 Hz because the common-mode-rejection of our setup vanishes above 10 Hz. The sensing heads do not share the same base plate and are physically separated by 20 cm. Above 100 Hz, our passive isolation system suppresses the coupling of the ground vibrations. We achieve the isolation by positioning the sensing heads on metal plates supported by rubber pads. The sealed box is also supported by another set of rubber pads.

7.3.2 Laser frequency noise

Our displacement sensitivity is strongly limited by the laser frequency noise below 10 Hz as shown in Fig. 7.2a. The figure also shows the readout signal calibrated to the units of $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ when no modulation is provided to the laser and the Michelson interferometer is in the gray state: the readout power is at half maximum. In this state, the detector is sensitive to the displacement noise of the test mass, laser frequency and intensity noise as well as the readout electronics.

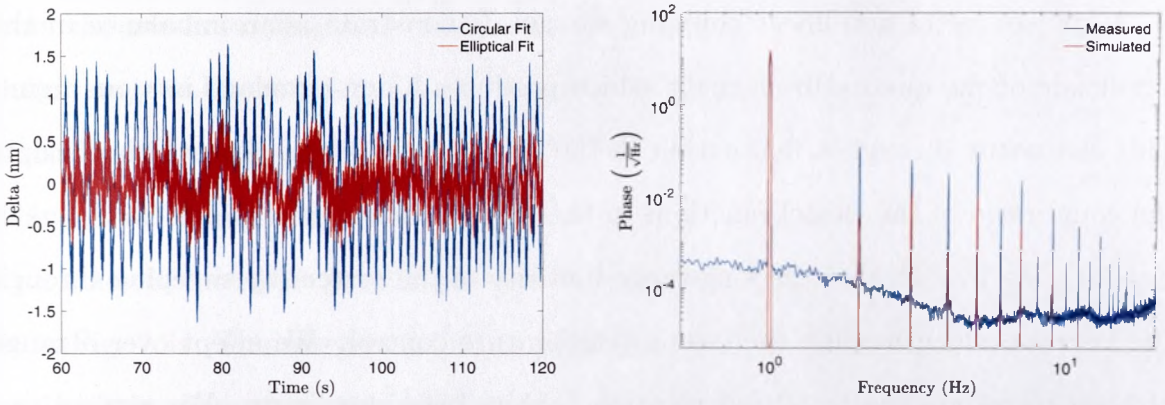
The frequency noise rises sharply towards lower frequencies and causes significant drifts in the displacement measurement over long periods of time. It is the most significant noise source in our region of interest and poses a significant limit to improvements beyond existing displacement sensor performance. We thus devoted our attention to suppressing this noise to improve the performance of the sensor in this key region.

Typically, the laser sensing noise in Michelson interferometers is suppressed by reducing the longitudinal imbalance between the two arms ($\Delta L = 0\,\mathrm{m}$). However, our scheme intentionally uses the imbalance to read out the signal at multiple harmonics of the modulation frequency, ω_m. Therefore, instead of reducing the imbalance, we follow the approach from [250, 251] and feed the reference head readout channel back to the laser modulation servo as shown in Fig. 7.1. This has the effect of rejecting signal, common to both sensing heads, from the second sensing head (the 'measurement' head). The displacement spectrum obtained in closed-loop operation is shown in Fig. 7.2b. Any true motion of the test mass that is detected by the measurement head will be independent of the reference head readout and thus will not be suppressed by the frequency stabilisation servo. Therefore, the noise suppression achieved in Fig. 7.2b represents a true improvement in the signal-to-noise ratio.

7.3.3 Down-converted noises

Aside from noises that naturally sit at the frequencies measured by the sensing head, there are noise sources which are shifted in frequency according to the choice of modulation frequency. The former noise types, such as laser frequency noise, manifest in the same way as the displacement signal, ϕ, in Eq. 7.1. Therefore, they undergo the same shift from sidebands around DC up to sidebands around ω_m and then shift back to DC upon demodulation. The latter types of noise, such as analogue-to-digital converter noise, laser intensity noise, transimpedance amplifier noise, and shot noise, couple directly into the series terms in Eq. 7.2. These noises, therefore, only undergo the down-conversion from the demodulation, which leads to displacement noise coupling of $n_x(\Omega) \propto n_i(\omega_m \pm \Omega)$ for a source i of this type.

This important distinction is a key practical consideration in our choice of modulation frequency. The region around the modulation frequency and its harmonics must be intrinsically lower in noise compared to the signal in the region of interest. Otherwise, the down-converted noise will dominate over the signal in the readout channel. For this very reason, we avoided placing the modulation frequency in the intermediate frequencies dominated by the forest of mechanical resonances. We choose the modulation frequency of 700 Hz because the laser frequency noise and seismic disturbances are below the broadband laser amplitude noise at this frequency as shown in Fig. 7.2a. In the unconstrained mirror application, such as in the readout of a multiple-degree-of-freedom seismometer [56], this factor will be of particular importance as the seismic measurement will become dominant over other noise sources and will itself be down-converted and potentially obscure the measurement at other frequencies. We are interested in enhancing the precision towards the low frequency limit and, fortunately, the ground motion inherently decreases towards the high frequencies. Thus, if we place



(a) (b)

FIGURE 7.3: Panel (a) shows the residual displacement from non-linear coupling observed when performing a sweep over the laser wavelength using temperature tuning. The circular fit represents a phase extraction algorithm that assumes balanced harmonics. The elliptical fit shows the improved linearity that we obtained through a more advanced algorithm that applies ellipse fitting to the Lissajous figure. Panel (b) shows the phase spectrum of a 1 Hz excitation of the laser wavelength using the PZT. The measured data represents the upper limit for the distribution of signal to higher harmonics due to non-linear coupling. The simulated curve represents our estimate of the spectrum that would correspond to the elliptical fit in panel (a).

the modulation at an appropriately high frequency, the down-converted ground motion should not pose an issue.

7.3.4 Linearity

Interferometric devices have an exceedingly narrow span of linearity in their raw readout (less than a quarter-wavelength). The long-range linearity is recovered via the construction of a Lissajous figure using the aforementioned dual-quadrature readout scheme. This technique relies on the correct determination of the centre position and eccentricity of the Lissajous figure in the phase extraction using the four-quadrant arctangent. If these ellipse parameters are determined poorly, the phase signal couples non-linearly into the readout channel, making non-linearities a particular area of concern in the literature [242].

A key source of non-linear coupling we can demonstrate is an imbalance in the amplitude of the quadrature signals, which produces a non-circular Lissajous figure. This can occur through a fluctuation in the modulation depth, A_m, which disrupts the equal ratio of the Bessel functions in the first two harmonics ($J_1(m) \neq J_2(m)$ in Eq. 7.4). We investigated the long-range linearity of the system by sweeping through the laser wavelength using the laser's temperature control. We swept over a range of 1 nm corresponding to 62 full rotations of the Lissajous figure. We applied two phase extraction algorithms in post-processing, shown in Fig. 7.3a, using a circular and elliptical fit. The former algorithm mimics our real-time phase extraction technique that assumes balanced harmonics, giving a residual RMS displacement of 0.68 nm over an effective displacement of 48 µm. The latter algorithm applies an elliptical fit to the many-period Lissajous figure constructed from the full sweep in order to correct for non-zero eccentricity and offsets in the centre of the Lissajous figure. Algorithms developed in [252,253] can be readily adopted to perform the ellipse fitting. The residual RMS in this case was reduced to 0.25 nm.

We also approached the investigation of non-linear coupling through another method, consisting of a controlled 1 Hz excitation of the laser frequency using the laser PZT. The amplitude of the oscillation was set at approximately half the Lissajous period in order to maximise the non-linear effect (we were limited in our choice of amplitude by the range of the PZT). The amplitude spectrum of the phase derived from the real-time readout, shown in Fig. 7.3b, contains a forest of peaks at multiples of 1 Hz, which represent the non-linear coupling of the pure 1 Hz input signal. The measured curve should be treated as an upper limit to the non-linearity of the system, as the non-linear response of the PZT is dominant over the non-linearity of the sensing head readout. We define a measure of the single-period non-linearity as the ratio of the RMS contribution

from the non-linear frequency components to the total RMS of the signal. Under this definition, the measured spectrum shows a non-linear measure of 1.2%. The value is similar to the non-linearity observed in [254] for TLB 6700 Velocity laser from Newport. The simulated curve in Fig. 7.3b contains the estimated spectrum for an excitation with a non-linear RMS contribution that has been matched to the RMS of the elliptical fit in Fig. 7.3a. This model exhibits a non-linear measure of 0.12%.

In applications that experience larger displacements, such as the position sensing of a suspension, the range of the signals will span multiple interference fringes, making the coupling of non-linearities more prominent than in the short displacement range of our fixed-mirror measurement. Naturally, this makes the topic of linearity particularly relevant.

7.4 Conclusion

Compact, interferometric position sensors are well positioned as candidates for the next generation of displacement readout in precision applications. Two such examples, the 6D seismometer [56] and the *Advanced LIGO* quadruple suspensions [42], can benefit from compact interferometers.

We presented a commercially available, interferometric displacement sensor which makes use of the sharp sub-fringe response of a Michelson interferometer. Using deep frequency modulation techniques, we were able to linearise the response over the range of a full fringe and a simple fringe-counting algorithm was implemented to extend the range over multiple fringes. Our fixed mirror measurement showed an improved level of sensitivity over existing shadow sensors. This, together with the sensor's compact design, means it is a suitable candidate for future upgrades and can be readily implemented

with little impact on existing infrastructure. The complexity has been effectively shifted to the digital domain. This does make the sensors more difficult to implement in situations where existing computational systems are lacking, but in the case of current GW detectors, this will not be an issue.

We were able to overcome the most significant initial sensitivity limitation, the coupling of laser frequency noise, through the use of one sensing head as a frequency reference. Particularly in multiple degree-of-freedom readout situations, where all sensing heads can receive input from the same laser source, this is a very low-impact solution to improving upon current displacement sensing methods. In a future upgrade, we will read out the signal from the anti-symmetric port of the Michelson interferometer. The approach will help suppress extra noise from the fibre between the sensing head and the photodetector [255].

Finally, we have discussed the occurrence of non-linearities in the presence of phase extraction algorithms that apply circular and fitted, static ellipse parameters. The application of our displacement sensor in, for example, suspended test masses may require the routine correction of the parameters of the Lissajous fit. This may arise due to long-term effects such as drifts in the modulation depth and timing jitter, but also due to a transformation of the Lissajous figure over high-RMS displacements, which will vary significantly depending on the specifics of the application.

A more detailed study on the impact of drifts in the parameters of the elliptical fit will be left for future work. Our simplified real-time phase extraction technique may not be suitable for high-RMS applications such as the displacement sensing of a suspended mass with no seismic pre-isolation. In this case, however, it is possible to employ a more advanced, dynamic algorithm that can apply parameter corrections in real-time and maintain linearity more robustly [249]. We are in the process of testing

the effectiveness of such an algorithm applied to the sensing heads presented.

The impact of non-linearities in the phase readout will be suppressed in applications such as the *Advanced LIGO* quadruple suspensions. This is due to their significantly lower RMS displacement, which is achieved by the previous stages of seismic isolation [52]. In this case, the RMS displacement should be lower than that simulated by our frequency sweep in Section 7.3.4, which makes our set-up suitable for such an application already, without significant expansion of the phase extraction scheme.

Chapter 8

Conclusion

Precision interferometry is a versatile tool used in many areas of scientific discovery. The ability to probe systems with light across vast scales, from trapped single atoms to kilogram-scale mirrors over kilometre ranges, couples with a state-of-the-art level of displacement sensitivity. Of particular note, interferometry is central in the development of gravitational wave detection—a field that has opened a new window upon the universe only in the last decade.

Despite the success of interferometry, there are many areas for further development, ranging from improvements in sensitivity, novel suggestions in detector design, and novel applications of interferometry in practical areas, as well as in settling questions of fundamental physics importance. The chapters presented in this work explore a variety of approaches, both theoretical and experimental, across all these developmental areas.

The experiment presented in Chapter 3 and 4 concerns a novel attempt at reaching the so-called standard quantum limit in interferometric displacement sensing. The contemporary configuration of gravitational wave detectors comprises large, suspended optics resonating a high-power electric field in the optical cavity that they form. In such a case, the relative displacement readout is susceptible to the detrimental effects of quantum noise. However, the presence of these qantum fluctuations also provides us

with the opportunity to investigate aspects of macroscopic quantum mechanics, such as test-mass entanglement via the radiation pressure effect, and test models of quantum gravity.

Chapter 3 is concerned with the motivation of the experiment, whose ultimate goal is to reach a quantum-noise-dominated level of sensitivity over a band surrounding the standard quantum limit at 100 Hz. To achieve this, it is essential to mitigate two prominent noise sources, these being seismic noise and thermal noise. The rest of Chapter 3 is concerned with the mitigation strategy of the former. The design of the suspension is performed using the linear Lagrangian formalism and we arrive at a four-stage passive suspension design. However, it is non-trivial to analyse a many-degree-of-freedom system such as this, mainly due to the range of cross-coupling and nonlinear effects that occur in the real experiment due to imperfections in design and assembly. This generally results in a reduction in the suspension's isolating performance, which puts a strain on our ambitious target of 8 orders of magnitude of seismic suppression. To plan ahead for a possible need to expand the isolation performance, we outline the design of an active stabilisation scheme of the cryostat frame itself, which will form a significant part of the ongoing seismic isolation development strategy. Overall, it is important to understand the behaviour of the suspension and much of future testing will focus on understanding the nuances of the suspension performance, which will undergo continuous, iterative improvement.

The discussion of the experiment continues in Chapter 4 with thermal considerations, particularly with the process for cooling the suspension to the 20 K operating point and the material choices to achieve minimal thermal noise. Our ability to reach the standard quantum limit will depend on the development status of low-noise coatings that have not yet seen widespread adoption, such as GaAs/AlGaAs or a-Si/SiO_2. Thus,

we also consider existing, well-tested coatings (Ta_2O_5/SiO_2) as a risk-mitigation strategy. However, it is unlikely that we will be able to reach the required level of performance with these coatings. Hence the success of the experiment is, in part, dictated by the near-future readiness of the experimental coatings. A key novelty of the experiment is the cryogenic operation of our suspended cavity using silicon substrates. The lower temperature will reduce thermal noise, which is the significant limiting factor that prevents existing experiment from reaching the quantum-level sensitivity currently.

The chapter concludes with a presentation of the current status of the experiment and discusses in some detail the specifics of the cavity alignment and control scheme. The final result shows the successful locking of a suspended, high-finesse ($>10^5$) cavity and its initial displacement spectrum. In the immediate future, a dual-cavity scheme will be implemented to produce a low-noise cavity spectrum in an attempt to measure the quantum back-action, or radiation pressure noise. Over longer timescales, the experiment will progress along the milestones laid out clearly in Chapter 3.

Beyond gaining an understanding of quantum noise and utilising it for an enhanced understanding of physics, there is broad scope for innovative techniques for suppressing the quantum noise below currently attainable levels. One such approach is to enhance the optical gain of an interferometer using an optomechanical filter. Chapter 5 proposes such an implementation by replacing one of the optics in the *Advanced LIGO* signal recycling cavity with a mechanical oscillator. The overall mechanism is to enhance the amplitude of the signal field by converting photons from a static pump, carefully offset in frequency from the main laser, into photons at the signal frequency. This is achieved via a feedback mechanism wherein the mechanical oscillator's position is perturbed by the radiation pressure interaction from the beat between the signal and pump fields, and the resulting motion modulates the pump field to coherently amplify the

amplitude of the signal field. We develop a dual frequency- and time-domain technique to independently verify this effect in a theoretical framework. This configuration is a novel example of a technique that suppressed the quantum noise below the nominal quantum limit imposed by the Heisenberg uncertainty principle in the conventional interferometer configuration. It is an example of how the manipulation of the signal field through non-linear processes can break out of the conventional quantum noise constraints, similar to the generation of squeezed states of light. Interestingly, however, the filter operates as a signal amplifier that does not affect the noise in the same way. As a result, the amplification boosts the signal-to-noise ratio, which is a novel aspect of this quantum amplification scheme.

The search for novel gravitational wave detector designs continues in Chapter 6 with a proposal for a high-frequency gravitational wave detector. Such a configuration is justified through the potential to detect novel high-frequency gravitational wave signals, produced during the collapse of a post-merger neutron star remnant. To characterise the noise performance of such high-frequency detectors, we must relax certain assumptions that are made in the derivation of contemporary classical noise models. Of particular interest in this chapter is the correct estimation of the substrate thermal noise in a frequency band that approaches the forest of substrate resonant modes. In such a case, the thermal noise sharply rises in spikes well above the broadband thermal noise baseline. The substrate dimensions must thus be carefully chosen to tune the locations of the resonant peaks away from the detection band. Even in the case of successful tuning, the tails of the thermal noise peaks are sufficiently broad that the contemporary models significantly underestimate the level of noise. Although the great breakthrough in thermal noise estimation was the move away from numerical analyses of the substrate modes into a low-frequency approximation that yielded an analytic equation, the future

calculations will have to return to the modal analysis. There are no simple analytic solutions for deriving the thermal noise around the resonant peaks of the substrate modes. As next generation detectors continue to expand the detection band to higher frequencies, and adopt heavier mirrors whose modes reach to lower frequencies, the forest of resonances will encroach into the detection band of the interferometers. Therefore, the analysis shown in this chapter provides a useful example of the future trajectory of thermal noise analysis.

The final experiment presented in Chapter 7 concerns the sensitivity measurement of a low-noise interferometric position sensor. We investigate a device by SmarAct, which we integrate into a digital, phase-readout scheme that uses deep frequency modulation techniques. This device imprints the displacement of an external mirror onto the sidebands of a carefully controlled frequency modulation. By demodulating the resulting signal as several harmonics of the base frequency, we gain access to multiple sinusoidal signals, which can be used to produce a Lissajous figure. From this, and with the addition of a fringe-counting algorithm, we can track the displacement of a mirror across multiple wavelengths of range. We measure the sensitivity of such a device at $0.3\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$, which is a factor of 300 better than the current displacement sensors used on *Advanced LIGO*. The device shows promise as a future upgrade to such devices and more broadly as a highly sensitive readout alternative in inertial sensors. The pinnacle of displacement sensitivity is achieved by interferometric sensors, most prominently by the gravitational-wave detectors mentioned repeatedly throughout this work. These detectors are anything but compact. However, they prove the efficacy of the interferometric principle in metrology. Therefore, sensors such as the one we present in this chapter represent the future of high-precision displacement devices and we show that they are ready to satisfy existing requirements for new sensors in highly

sensitive facilities such as the *Advanced LIGO* detectors and the *Large Hadron Collider*. Future developments will see devices like this grow in both performance and variety of applications as research and industry will continue to demand increasingly more precise diplacement sensors.

Appendix A

Relations in Interferometry

A.1 Michelson Interferometer—Output Power

The output power at the photodetector of a Michelson interferometer as shown in Fig. 1.1b can be derived by considering the standard field interaction at the beamsplitter and propagation through space [5]. For fields defined according to the schematic in Fig. A.1, the following set of equations is obtained

$$A_1 = r_{\mathrm{BS}} A_{\mathrm{in}}, \tag{A.1}$$

$$A_2 = A_1 \exp(-2ikx_1), \tag{A.2}$$

$$A_3 = t_{\mathrm{BS}} A_{\mathrm{in}}, \tag{A.3}$$

$$A_4 = A_3 \exp(-2ikx_2), \tag{A.4}$$

$$A_{\mathrm{out}} = t_{\mathrm{BS}} A_2 - r_{\mathrm{BS}} A_4, \tag{A.5}$$

where $r_{\mathrm{BS}} = t_{\mathrm{BS}} = 1/\sqrt{2}$ is the amplitude reflectivity and transmissivity respectively, $k = 2\pi/\lambda$ is the wavenumber, λ is the wavelength, and x_1 and x_2 are the distances from the beamsplitter to the north and east arm mirror respectively. The convention is used where a positive and negative sign in reflection off the beamsplitter is applied according

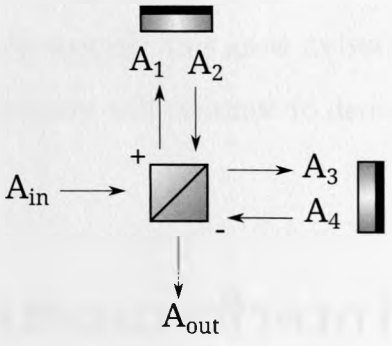


FIGURE A.1: Definitions of fields in a Michelson interferometer used in the output power derivation.

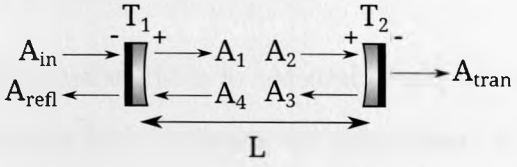


FIGURE A.2: Definitions of fields in a Fabry-Pérot interferometer used in the derivation of reflected and transmitted power.

to Fig. A.1. Solving these simultaneous equations yields the output field in terms of the input,

$$A_{\text{out}} = \frac{A_{\text{in}}}{2}[\exp(-2ikx_1) - \exp(-2ikx_2)]. \tag{A.6}$$

The output power thus becomes

$$P_{\text{out}} = |A_{\text{out}}|^2 = \frac{P_{\text{in}}}{2}[1 - \cos(2k(x_2 - x_1))]. \tag{A.7}$$

A.2 Fabry-Pérot Interferometer—Output Power

The power at the photodetectors in reflection and transmission of a Fabry-Pérot interferometer can be derived from the following set of equations

$$A_1 = t_1 A_{\text{in}} + r_1 A_4, \tag{A.8}$$

$$A_2 = A_1 \exp(-ikL), \tag{A.9}$$

$$A_3 = r_2 A_2, \tag{A.10}$$

$$A_4 = A_3 \exp(-ikL), \tag{A.11}$$

$$A_{\text{refl}} = t_1 A_4 - r_1 A_{\text{in}}, \tag{A.12}$$

$$A_{\text{tran}} = t_2 A_2, \tag{A.13}$$

which applies to fields as defined in Fig. A.2. Solving for the output fields, A_{refl} and A_{tran}, in terms of the input field gives

$$A_{\text{refl}} = \frac{r_2 \exp(-2ikL) - r_1}{1 - r_1 r_2 \exp(-2ikL)}, \tag{A.14}$$

$$A_{\text{tran}} = \frac{t_1 t_2 \exp(-ikL)}{1 - r_1 r_2 \exp(-2ikL)}. \tag{A.15}$$

The corresponding power measured by each photodetector becomes

$$P_{\text{refl}} = |A_{\text{refl}}|^2 = \frac{R_1 + R_2 - 2r_1 r_2 \cos(2kL)}{1 + R_1 R_2 - 2r_1 r_2 \cos(2kL)} P_{\text{in}}, \tag{A.16}$$

$$P_{\text{tran}} = |A_{\text{tran}}|^2 = \frac{T_1 T_2}{1 + R_1 R_2 - 2r_1 r_2 \cos(2kL)} P_{\text{in}}. \tag{A.17}$$

Appendix B

Relations in Interferometer Sensitivity

B.1 Offsets in Linear Systems

Consider the spring system in Fig. 2.3a. If a vertical offset is applied to the mass suspended by the spring, for example due to a constant gravitational force acting on the mass, this will lead to a static elongation of the spring and thus an increase in the stored potential energy. We can thus rewrite the potential energy term in the Lagrangian (Eq. 2.12) as

$$V = \frac{1}{2}k(z - z_g + z_{\mathrm{eq}})^2, \tag{B.1}$$

where z_{eq} is the constant equilibrium distance from the mass to the zero-point 'ground' point of the spring. The resulting contribution of this term in the EL equation is

$$kz(t) - kz_g(t) + kz_{\mathrm{eq}}. \tag{B.2}$$

Whilst the first two terms are functions of time, the final term evidently is not. Upon applying the Laplace transform to Eq. B.2, we notice that the third term (which becomes a zero-centred delta function) vanishes for all non-zero frequencies. Therefore, the dynamics of the system are not affected by an offset either due to a static force or a

translation of the coordinate system.

B.2 Direct Approach to Thermal Noise—Point-Mass Suspension

We can start with the general equation for the thermal noise in displacement of a mechanical system using the direct approach (Eq. 2.44,

$$S_x(\Omega) = \frac{8k_BT}{\Omega^2}\frac{W_{\text{diss}}}{{F_0}^2}. \tag{B.3}$$

For a single-stage isolation system modelled as a spring-suspended point-mass, we can obtain the following set of equations,

$$W_{\text{diss}} = \Omega U_{\text{max}}\phi, \tag{B.4}$$

$$U_{\text{max}} = \frac{1}{2}k_0{x_0}^2, \tag{B.5}$$

$$F_0 = m\Omega^2 x_0. \tag{B.6}$$

The homogeneous loss of the point-mass system is given by Eq. B.4 (Eq. 4.8 in [73]), where U_{max} is the elastic potential energy stored at peak deformation of the mechanical system. For a spring, this quantity is given by Eq. B.5, where x_0 is the peak extension/compression of the spring. If we confine the analysis to frequencies well above the resonance, then the system responds to an applied force approximately as a free mass, leading to Eq. B.6. Substituting Eq. B.5 and B.6 into Eq. B.4 yields

$$W_{\text{diss}} = \frac{k_0\phi{F_0}^2}{2m^2\Omega^3}. \tag{B.7}$$

Substituting this relation into Eq. B.3 results in

$$S_x(\Omega) = \frac{4k_B T k_0 \phi}{m^2 \Omega^5}. \tag{B.8}$$

This result is consistent with the result derived directly using the *fluctuation-dissipation theorem*, obtaining the PSD of the thermal force in Eq. 2.40 and using the free-mass approximation to transform into a displacement PSD.

Appendix C

Cryogenic Experiment

C.1 Optimised Blade Design

We wish to optimise the design of the blade springs to minimise their resonant frequencies. Starting with Eq. 3.24, we note which parameters we want to minimise and which maximise. However, we are constrained in the extent of to which we can optimise each parameter by (a) the space limitations within the cryostat, and (b) the yield limit of the material. Using Eq. 3.25, we note that, for example, increasing the length, l, can be compensated by an increase in the base width, a, to maintain constant stress.

We can consider the parameters in pairs, scale each quantity in a way that maintains constant stress, and see whether the resonant frequency benefits from the scaling. By considering each pair of parameters, we can determine which parameters should be optimised preferentially to guide our blade design. The results of this pair-wise consideration is shown in Table C.1. The first quantity is scaled to nominally decrease the resonant frequency. The second quantity is then scaled to compensate for the stress change. In the final column, the resulting scaling of the resonant frequency is shown. The scaling to a negative power means that it is preferable to optimise the former quantity over the latter.

TABLE C.1: Optimisation of Parameters

Parameter	Scaled By	Compensated With	Scaled By	Frequency Dependence
l	A	h	$A^{-\frac{1}{2}}$	$A^{-\frac{3}{4}}$
l	A	a	A	A^{-1}
l	A	m	A^{-1}	A^{-1}
h	A^{-1}	a	A^2	$A^{-\frac{1}{2}}$
h	A^{-1}	m	A^{-2}	$A^{-\frac{1}{2}}$
a	A^{-1}	m	A^{-1}	1

The result of the analysis provides some clear guidelines for the design. Firstly, it is optimal to maximise the length of the blade. When compensated with every other parameter, a longer blade always results in a lower resonant frequency. Once the length has been maximised, minimising the blade thickness is the next step in the optimisation. As counter-intuitive as it may seem, it is optimal to maximise the base width and minimise the attached mass, as this allows for an optimisation of the parameters that have a disproportionately more significant effect on the resonant frequency.

Our blade design is thus driven by these derivations. We begin the design process by maximising the blade length within the available space. For the top-most (CRYO) stage blades, we are limited by a central pillar containing the cryopump components. We thus arrange the blades in an equilateral triangle with the largest side length possible. On the subsequent stages, we are not obstructed in the middle and thus we design the blades to span the whole diameter of the available chamber cross-section in a cross shape. In this case we increase the blade number to four per stage. We then maximise the base width of the blades in order to allow for the thinnest blades possible. The result of this optimisation is the blade profile shown in Fig. 3.7.

C.2 Vertical Wire Stiffness

For a suspension wire of constant cross-sectional area, A_w, the stress and extension is uniform along the length of the wire. The Young's modulus, Y, is related to the stress, σ, and strain, δ, experienced by the wire, which is given by

$$Y = \frac{\sigma}{\delta} = \frac{F_v l}{A_w z}, \tag{C.1}$$

where F_v is the force applied to the wire, l is the wire length, and z is the vertical displacement of the suspended mass.

We can rearranging Eq. C.1 to yield the force and, noting Hooke's law ($F_v = k_{w,z} z$), we can determine the vertical spring constant as

$$k_{w,z} = \frac{Y A_w}{l}. \tag{C.2}$$

Appendix D

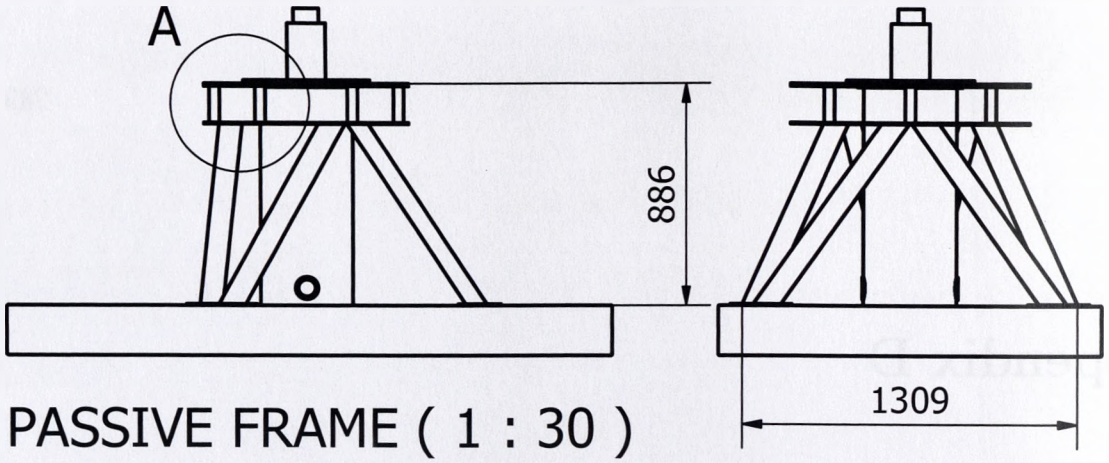
Technical Drawings

This appendix contains a collection of technical drawings relevant to the experiment explored in Chapters 3 and 4. These drawings are inserted here for completeness and as a point of reference for the key physical components of the experiment.

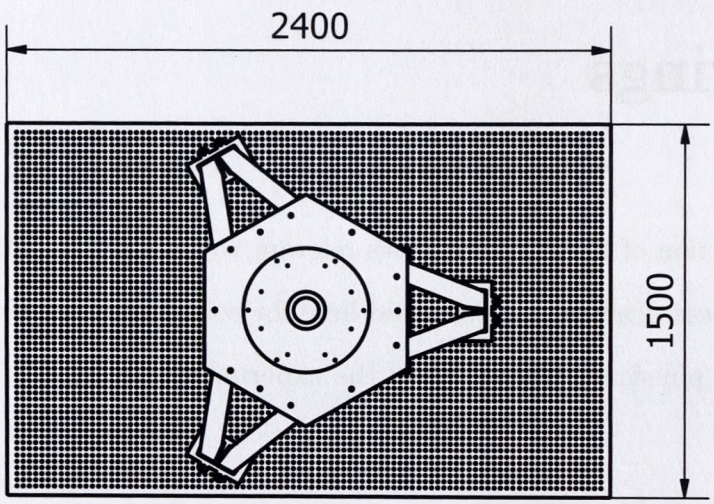
D.1 Cryostat Support Frame

The following drawings show a more detailed schematic of the two variants of the cryostat frame design. Both drawings show the same basic frame shape (highlighted in red) with the same overall dimensions seated on identical optical tables.

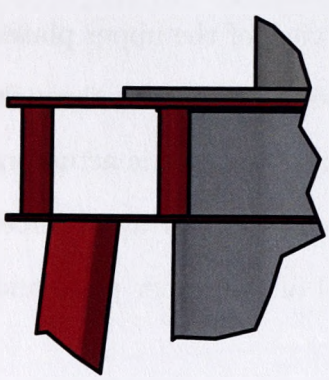
The details (A and B) on the drawings show an enlarged view of the upper plates attachment mechanism. In the passive frame design, the frame is attached through vertical pillars. The active frame design shows a series of angled piezoelectric actuators in a hexapod configuration. The top plate is further populated with an array of Sercel L-4C geophones showing a possible scheme for active control of the upper plate and cryostat.

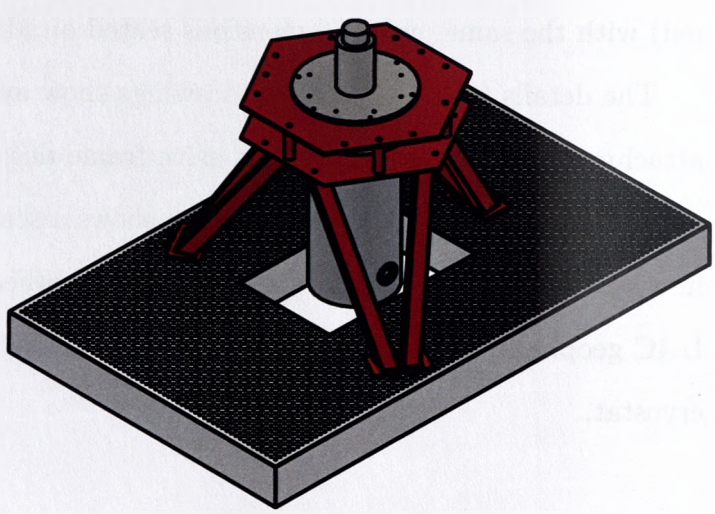


PASSIVE FRAME (1 : 30)

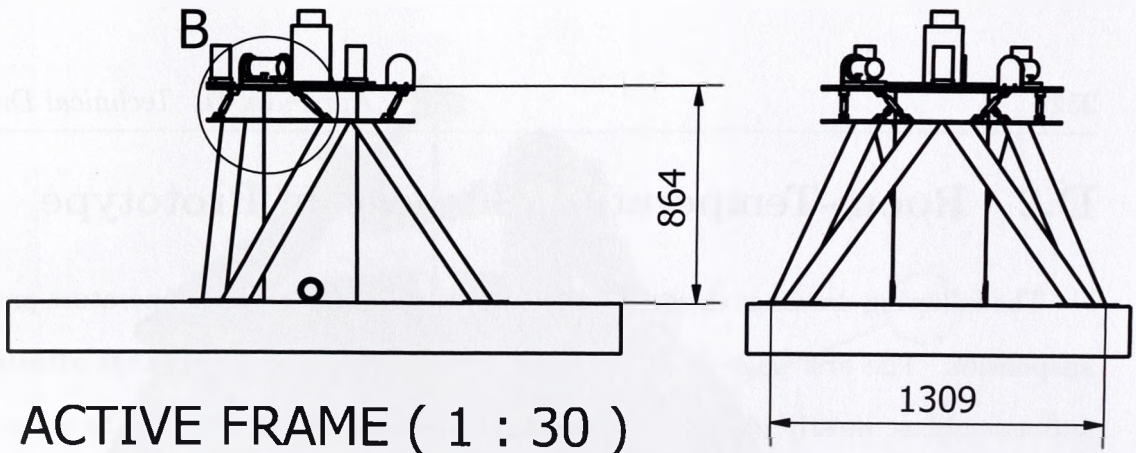


A (1 : 10)

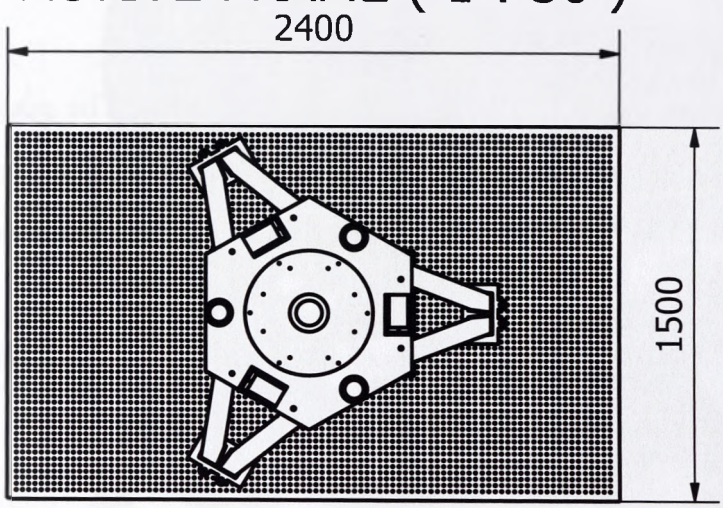




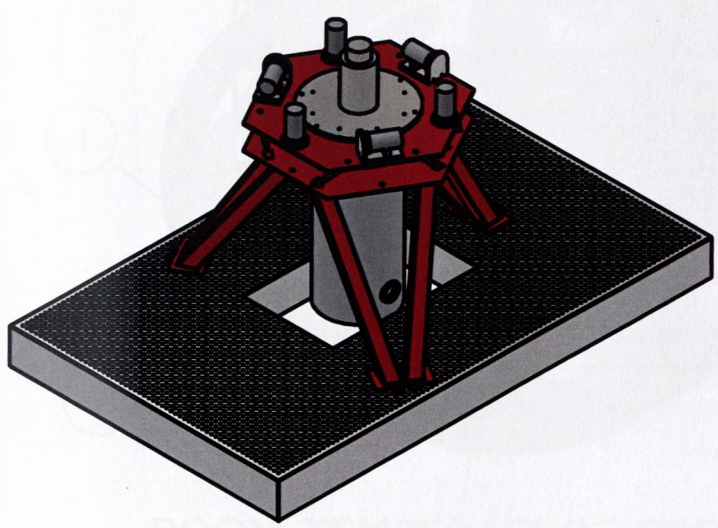
Designed by	Date
Jiri Smetana	13/09/2021
CRYOSTAT FRAME	1 / 2

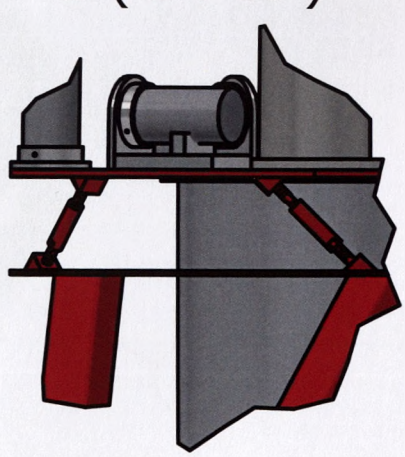


ACTIVE FRAME (1 : 30)



B (1 : 10)





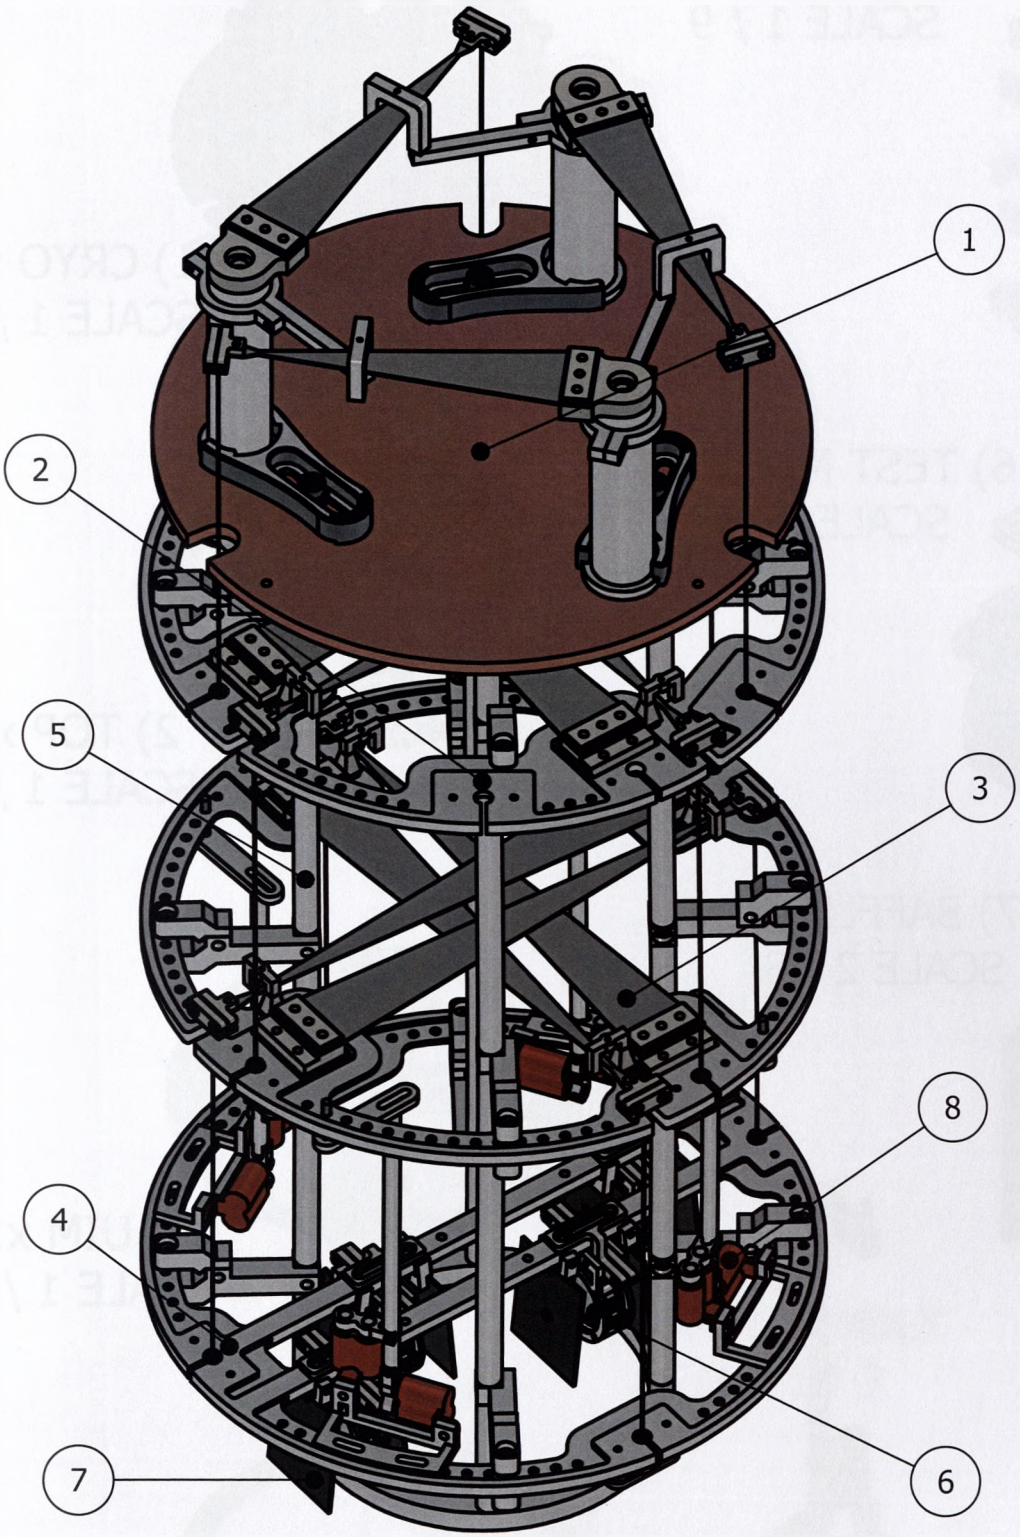
Designed by	Date
Jiri Smetana	13/09/2021
CRYOSTAT FRAME	2 / 2

D.2 Room-Temperature Suspension Prototype

The following drawings show a detailed schematic of the room-temperature prototype suspension. The first page shows an enlarged view of the assembly that identifies the sub-assemblies, mostly identified as the individual suspension stages. The second page shows the individual sub-assemblies separated from the complete design.

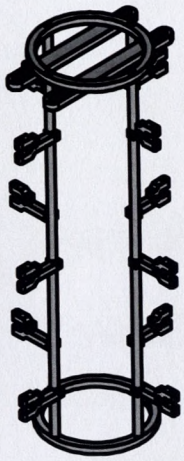
The CRYO stage hosts the upper blades on a copper cold plate. The plate itself is not a removable component but is rather permanently attached to the cryostat' central pillar. The pillar has been omitted from the drawing for clarity. Otherwise, the remainder of the components comprise original designs.

The subsequent pages show the individual parts that are to be joined to form each sub-assembly and thus can be used as a guide during construction.

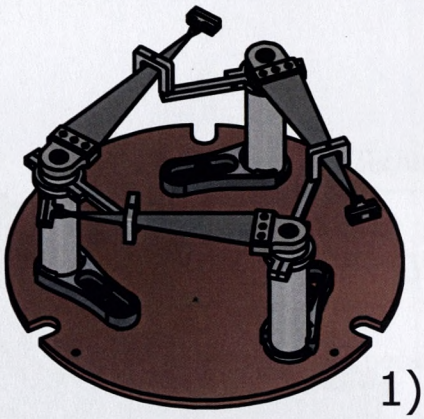


ROOM-TEMPERATURE PROTOTYPE
SCALE 1 / 3

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	1 / 7

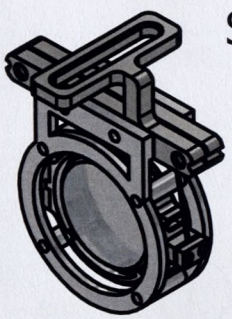


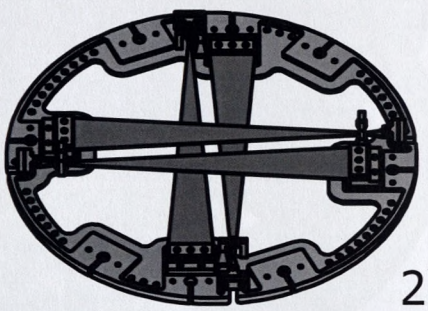
5) SAFETY x1
SCALE 1 / 9



1) CRYO x1
SCALE 1 / 5

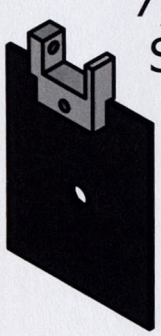
6) TEST MASS x 4
SCALE 2 : 3

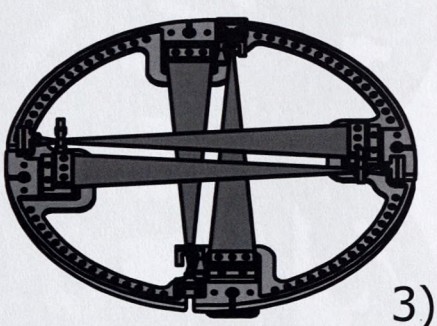




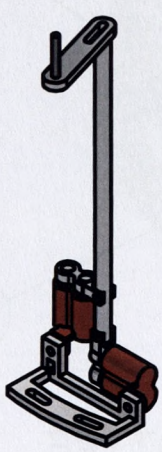
2) TOP x1
SCALE 1 / 5

7) BAFFLE x8
SCALE 2 : 3

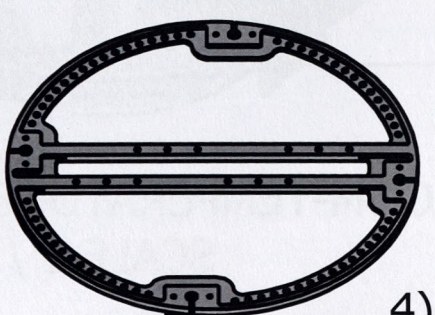




3) UIM x1
SCALE 1 / 5

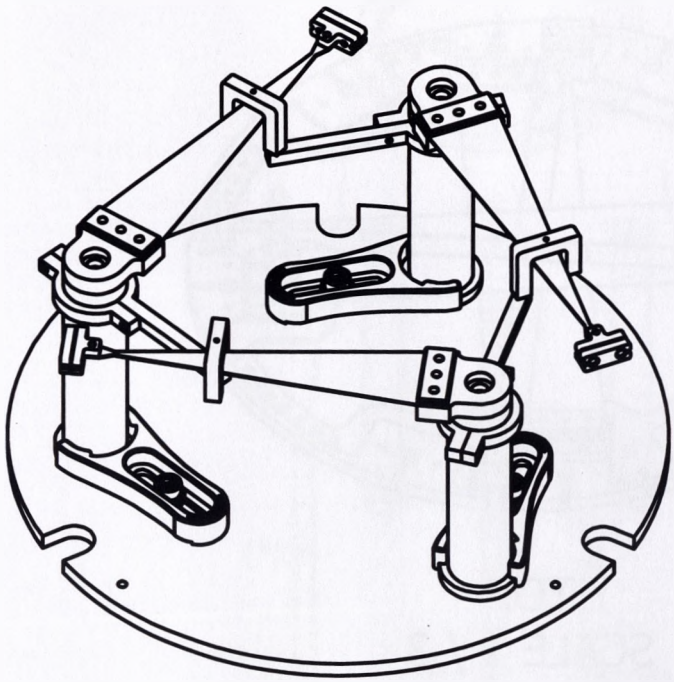


8) DAMPING x4
SCALE 1 / 3

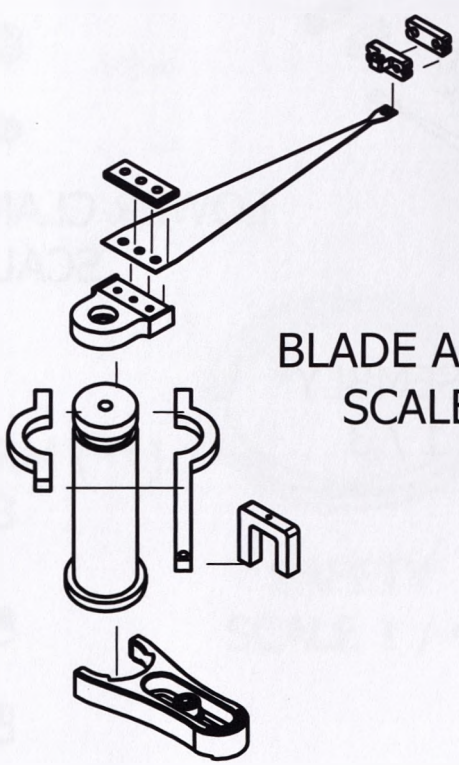


4) PUM x1
SCALE 1 / 5

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	2 / 7

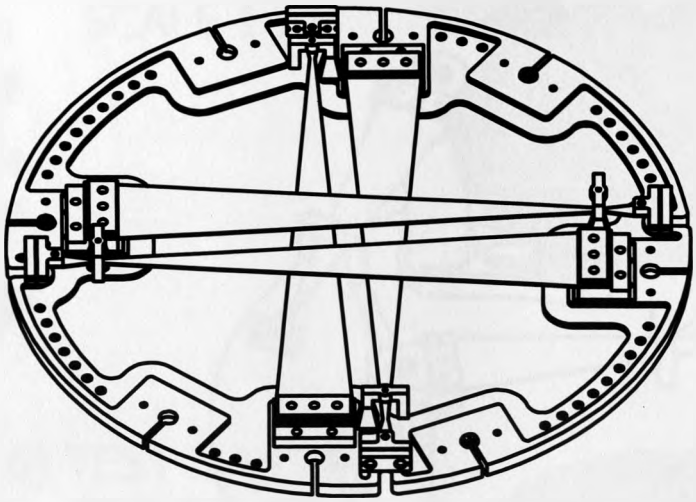


CRYO
SCALE 1 / 3

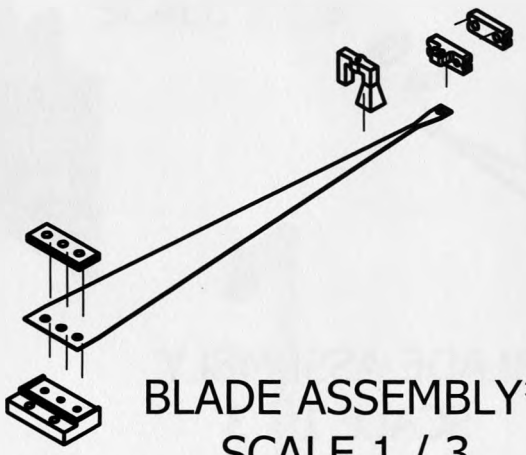


BLADE ASSEMBLY
SCALE 1 / 3

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	3 / 7



TOP
SCALE 1 / 3



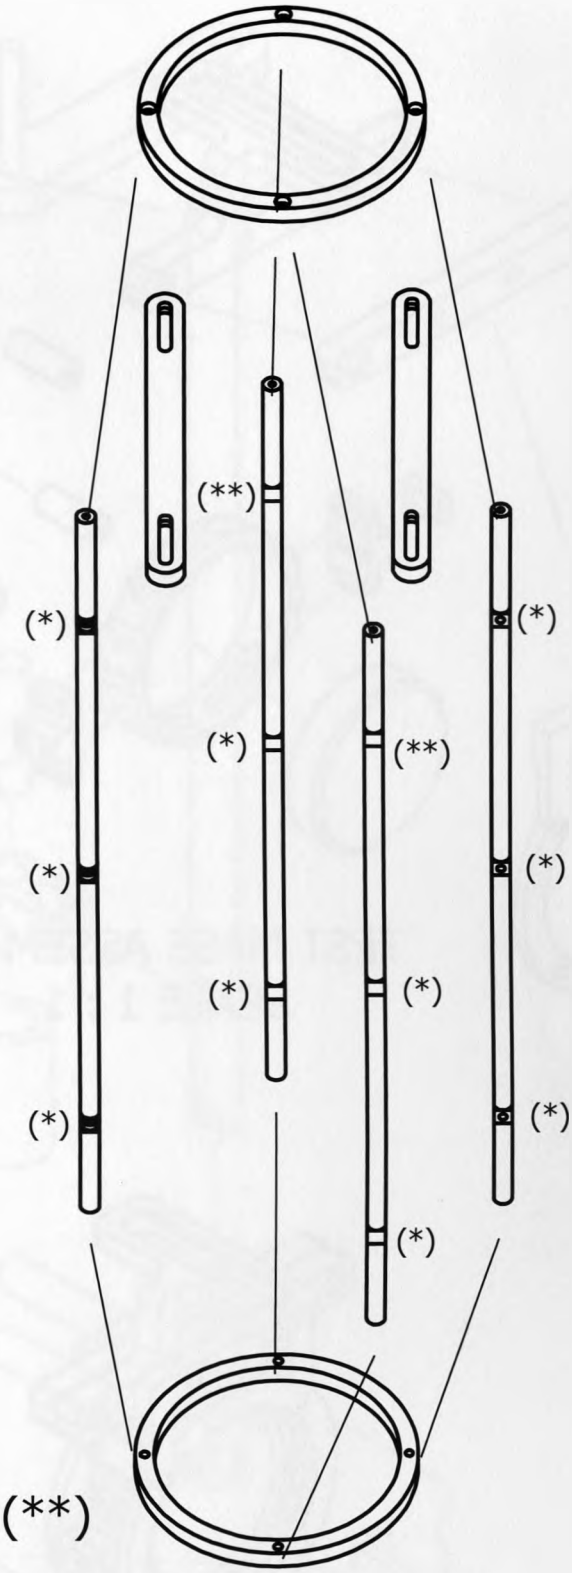
BLADE ASSEMBLY*
SCALE 1 / 3



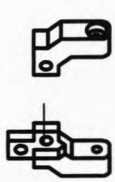
LOWER CLAMP ASSEMBLY*
SCALE 1 / 3

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	4 / 7

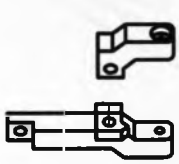
*(IDENTICAL ASSEMBLY FOR UIM & PUM)



SAFETY
SCALE 1 / 4

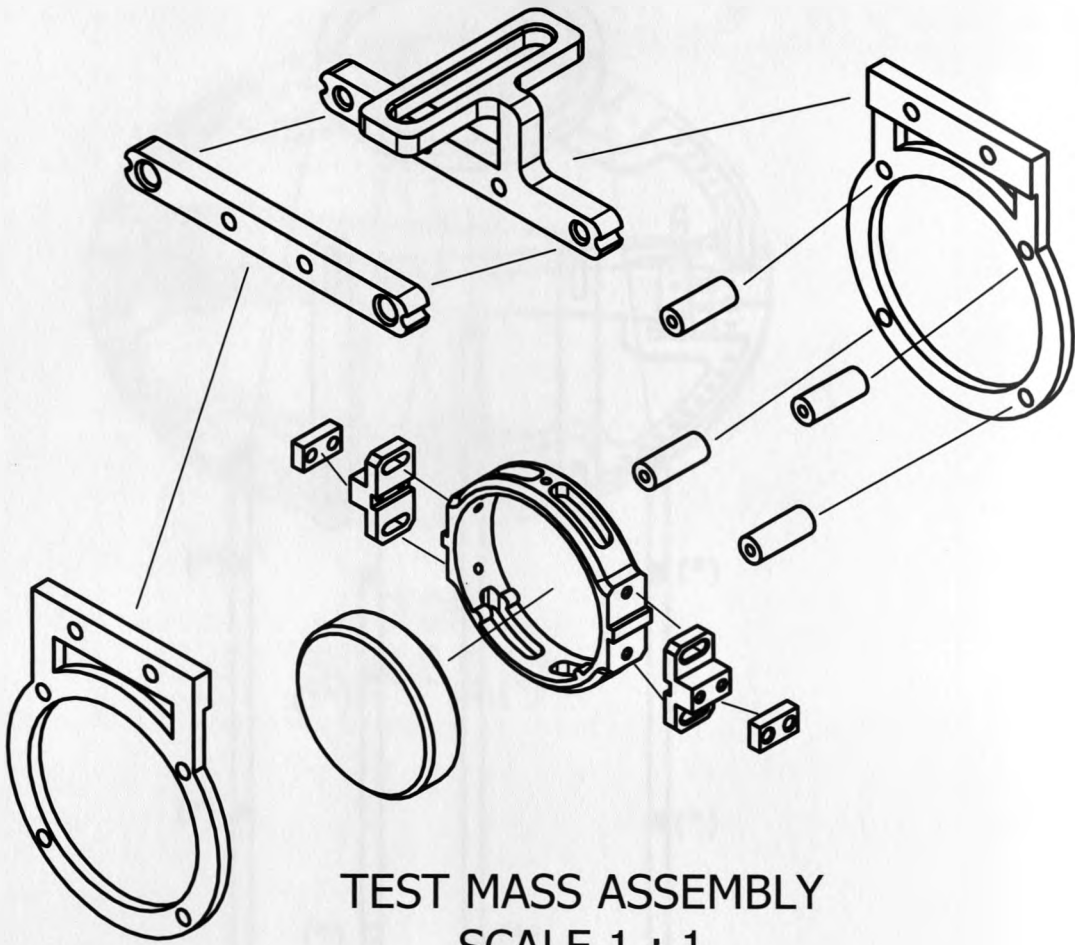


SHORT SAFETY ARM (**)
SCALE 1 / 3

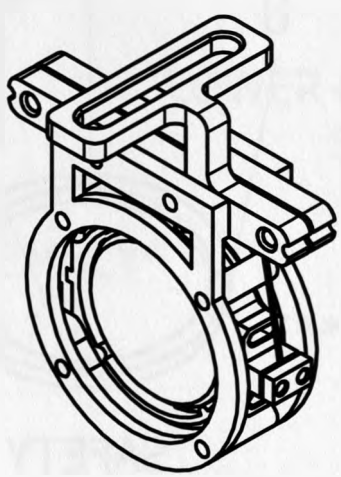


LONG SAFETY ARM (*)
SCALE 1 / 3

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	5 / 7

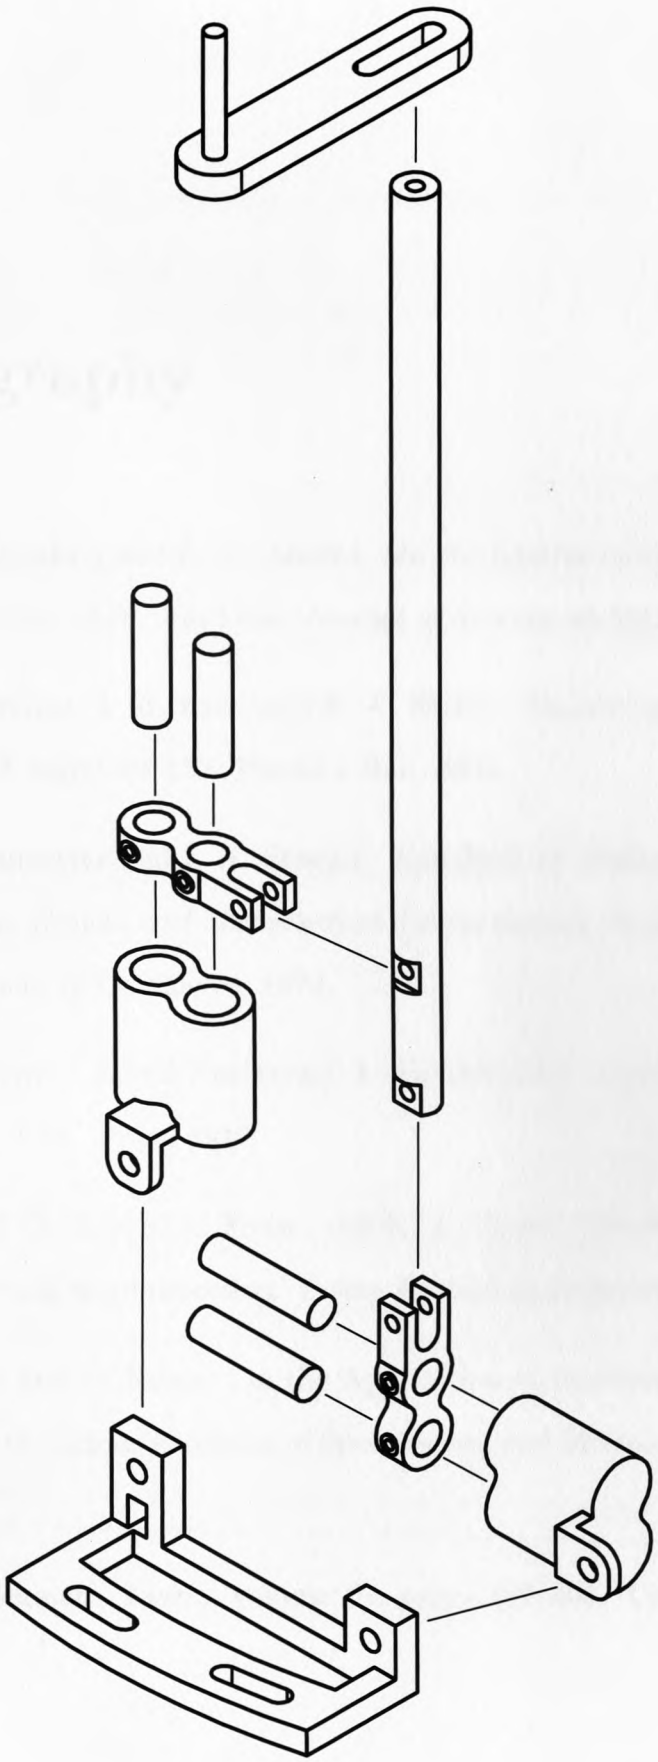


TEST MASS ASSEMBLY
SCALE 1 : 1



TEST MASS

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	6 / 7



DAMPING
SCALE 1 : 1

Designed by	Date
Jiri Smetana	24/09/2020
SUSPENSION	7 / 7

Bibliography

[1] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, s3-34(203):333–345, 1887.

[2] C. L. Phillips, J. M. Parr, and E. A. Riskin. *Signals, systems, and transforms*, chapter 3, pages 89–152. Prentice Hall, 2003.

[3] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, chapter 29, page 1020. United States Department of Commerce, 1972.

[4] C. L. Byrne. *Signal Processing A Mathematical Approach*, chapter 12, pages 169–191. CRC Press, 1947.

[5] C. Bond, D. Brown, A. Freise, and K. A. Strain. Interferometer techniques for gravitational-wave detection. *Living Reviews in Relativity*, 19(1):3, feb 2017.

[6] A. Pérot and C. Fabry. On the Application of Interference Phenomena to the Solution of Various Problems of Spectroscopy and Metrology. ApJ, 9:87, February 1899.

[7] A. E. Siegman. *Lasers*, volume 37, pages 462–466. University Science Books, 1986.

[8] J. Bechhoefer. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.*, 77:783–836, Aug 2005.

[9] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(2):97–105, June 1983.

[10] Daryl W. Preston. Doppler-free saturated absorption: Laser spectroscopy. *American Journal of Physics*, 64(11):1432–1436, 1996.

[11] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation.* W. H. Freeman, 1973.

[12] P. R. Saulson. If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves? *American Journal of Physics*, 65(6):501–505, 1997.

[13] J. Aasi, B. P. Abbott, R. Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, mar 2015.

[14] F. Acernese, M. Agathos, K. Agatsuma, et al. Advanced virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, dec 2015.

[15] T. Akutsu, M. Ando, K. Arai, et al. Kagra: 2.5 generation interferometric gravitational wave detector. *Nature Astronomy*, 3(1):35–40, 09 2019.

[16] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.

[17] A. Buikema, C. Cahillane, G. L. Mansell, et al. Sensitivity and performance of the advanced ligo detectors in the third observing run. *Phys. Rev. D*, 102:062003, Sep 2020.

[18] H. Miao, R. X. Adhikari, Y. Ma, B. Pang, and Y. Chen. Towards the fundamental quantum limit of linear measurements of classical signals. *Phys. Rev. Lett.*, 119:050801, Aug 2017.

[19] V. B. Braginsky. Classical and quantum restrictions on the detection of weak disturbances of a macroscopic oscillator. *Journal of Experimental and Theoretical Physics*, 26(4), April 1968.

[20] K. S. Thorne. Gravitational radiation. In S. W. Hawking and W. Israel, editors, *Three hundred years of gravitation*, chapter 9, pages 330–458. Cambridge University Press, Cambridge, 1 edition, 1987.

[21] H. Miao and Y. Chen. Quantum theory of laser interferometer gravitational wave detectors. In D. G. Blair, L. Ju, C. Zhao, and E. J. Howell, editors, *Advanced Gravitational Wave Detectors*, pages 277–297. Cambridge University Press, Cambridge, 2012.

[22] W. Schottky. On spontaneous current fluctuations in various electrical conductors. *Journal of Micro/Nanolithography, MEMS, and MOEMS*, 17(4):1 – 11, 2018.

[23] C. M. Caves. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Phys. Rev. Lett.*, 45:75–79, Jul 1980.

[24] V. B. Braginsky, F. Ya. Khalili, and K. S. Thorne. *Quantum Measurement.* Cambridge University Press, 1992.

[25] C. M. Caves, K. S. Thorne, R. W. P. Drever, V. D. Sandberg, and M. Zimmermann. On the measurement of a weak classical force coupled to a quantum-mechanical oscillator. i. issues of principle. *Rev. Mod. Phys.*, 52:341–392, Apr 1980.

[26] V. B. Braginsky, M. L. Gorodetsky, F. Ya. Khalili, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. Noise in gravitational-wave detectors and other classical-force measurements is not influenced by test-mass quantization. *Phys. Rev. D*, 67:082001, Apr 2003.

[27] C. M. Caves and B. L. Schumaker. New formalism for two-photon quantum optics. i. quadrature phases and squeezed states. *Phys. Rev. A*, 31:3068–3092, May 1985.

[28] K. J. Blow, R. Loudon, S. J. D. Phoenix, and T. J. Shepherd. Continuum fields in quantum optics. *Phys. Rev. A*, 42:4102–4114, Oct 1990.

[29] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum non-demolition interferometers by modifying their input and/or output optics. *Physical Review D*, 65(2):022002, dec 2001.

[30] H. Miao. *Exploring macroscopic quantum mechanics in optomechanical devices.* PhD thesis, University of Western Australia, 2010.

[31] J. Mizuno. *Comparison of optical configurations for laser-interferometric gravitational-wave detectors.* PhD thesis, AEI-Hannover, MPI for Gravitational Physics, Max Planck Society, 1995.

[32] Y. Enomoto, K. Nagano, and S. Kawamura. Standard quantum limit of angular motion of a suspended mirror and homodyne detection of a ponderomotively squeezed vacuum field. *Phys. Rev. A*, 94:012115, Jul 2016.

[33] P. Amaro-Seoane, H. Audley, S. Babak, et al. Laser interferometer space antenna. *ArXiv*, 2017.

[34] M. V. Plissi, K. A. Strain, C. I. Torrie, N. A. Robertson, S. Killbourn, S. Rowan, S. M. Twyford, H. Ward, K. D. Skeldon, and J. Hough. Aspects of the suspension system for geo 600. *Review of Scientific Instruments*, 69(8):3055–3061, 1998.

[35] E. Schwartz, A. Pele, J. Warner, et al. Improving the robustness of the advanced ligo detectors to earthquakes. *Classical and Quantum Gravity*, 37(23):235007, November 2020.

[36] C. C. Speake. Anelasticity in flexure strips revisited. *Metrologia*, 55(1):114–119, feb 2018.

[37] G. Cagnoli, J. Hough, D. De Bra, M. M. Fejer, E. Gustafson, S. Rowan, and V. Mitrofanov. Damping dilution factor for a pendulum in an interferometric gravitational waves detector. *Physics Letters A*, 272(1):39–45, 2000.

[38] T. Shimoda, N. Aritomi, A. Shoda, Y. Michimura, and M. Ando. Seismic cross-coupling noise in torsion pendulums. *Phys. Rev. D*, 97:104003, May 2018.

[39] H. Lück, P. Aufmuth, O. S. Brozek, K. Danzmann, A. Freise, S. Goßler, A. Grado, H. Grote, K. Mossavi, V. Quetschke, B. Willke, K. Kawabe, A. Rüdiger, R. Schilling, W. Winkler, Ch. Zhao, K. A. Strain, G. Cagnoli, M. Casey, J. Hough, M. Husman, P. McNamara, G. P. Newton, M. V. Plissi, N. A. Robertson, S. Rowan, D. I. Robertson, K. D. Skeldon, C. I. Torrie, H. Ward, B. F. Schutz, I. Taylor, and B. S. Sathyaprakash. The status of geo600. *AIP Conference Proceedings*, 523(1):119–127, 2000.

[40] H. Lück, C. Affeldt, J. Degallaix, A. Freise, H. Grote, M. Hewitson, S. Hild, J. Leong, M. Prijatelj, K. A. Strain, B. Willke, H. Wittel, and K. Danzmann. The upgrade of GEO 600. *Journal of Physics: Conference Series*, 228:012012, may 2010.

[41] M. V. Plissi, C. I. Torrie, M. E. Husman, N. A. Robertson, K. A. Strain, H. Ward, H. Lück, and J. Hough. Geo 600 triple pendulum suspension system: Seismic isolation and control. *Review of Scientific Instruments*, 71(6):2539–2545, 2000.

[42] N. A. Robertson, G. Cagnoli, D. R. M. Crooks, E. Elliffe, J. E. Faller, P. Fritschel, S. Go$szlig$ler, A. Grant, A. Heptonstall, J. Hough, H. L$uuml$ck, R. Mittleman, M. Perreur-Lloyd, M. V. Plissi, S. Rowan, D. H. Shoemaker, P. H. Sneddon, K. A. Strain, C. I. Torrie, H. Ward, and P. Willems. Quadruple suspension design for Advanced LIGO. *Classical and Quantum Gravity*, 19(15):311, aug 2002.

[43] S. Braccini, L. Barsotti, C. Bradaschia, et al. Measurement of the seismic attenuation performance of the VIRGO Superattenuator. *Astroparticle Physics*, 23(6):557–565, jul 2005.

[44] M. E. Husman, C. I. Torrie, M. V. Plissi, N. A. Robertson, K. A. Strain, and J. Hough. Modeling of multistage pendulums: Triple pendulum suspension for geo 600. *Review of Scientific Instruments*, 71(6):2546–2551, 2000.

[45] C. F. Da Silva Costa, C. Billman, A. Effler, S. Klimenko, and H.-P. Cheng. Regression of non-linear coupling of noise in LIGO detectors. *Classical and Quantum Gravity*, 35(5):055008, feb 2018.

[46] K. D. Hahn, E. M. Johnson, A. Brokken, and S. Baldwin. Eddy current damping of a magnet moving through a pipe. *American Journal of Physics*, 66(12):1066–1076, 1998.

[47] M. V. Plissi, C. I. Torrie, M. Barton, N. A. Robertson, A. Grant, C. A. Cantley, K. A. Strain, P. A. Willems, J. H. Romie, K. D. Skeldon, M. M. Perreur-Lloyd, R. A. Jones, and J. Hough. An investigation of eddy-current damping of multi-stage pendulum suspensions for use in interferometric gravitational wave detectors. *Review of Scientific Instruments*, 75(11):4516–4522, nov 2004.

[48] S. Rowan, S. M. Twyford, J. Hough, D.-H. Gwo, and R. Route. Mechanical losses associated with the technique of hydroxide-catalysis bonding of fused silica. *Physics Letters A*, 246(6):471–478, 1998.

[49] G. Cagnoli, L. Gammaitoni, J. Hough, J. Kovalik, S. McIntosh, M. Punturo, and S. Rowan. Very high Q measurements on a fused silica monolithic pendulum for use in enhanced gravity wave detectors. *Phys. Rev. Lett.*, 85:2442–2445, Sep 2000.

[50] L. Carbone, S. M. Aston, R. M. Cutler, A. Freise, J. Greenhalgh, J. Heefner, D. Hoyland, N. A. Lockerbie, D. Lodhia, N. A. Robertson, C. C. Speake, K. A. Strain, and A. Vecchio. Sensors and actuators for the advanced LIGO mirror suspensions. *Classical and Quantum Gravity*, 29(11):115005, may 2012.

[51] P. G. Nelson. An active vibration isolation system for inertial reference and precision measurement. *Review of Scientific Instruments*, 62(9):2069–2075, sep 1991.

[52] F. Matichard, B. Lantz, R. Mittleman, et al. Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, sep 2015.

[53] Oaul Lorrain. An optical lever for measuring very small rotational oscillations. *Optics and Lasers in Engineering*, 15(3):197–201, 1991.

[54] M. P. Ross, K. Venkateswara, C. A. Hagedorn, J. H. Gundlach, J. S. Kissel, J. Warner, H. Radkins, T. J. Shaffer, M. W. Coughlin, and P. Bodin. Low-Frequency Tilt Seismology with a Precision Ground-Rotation Sensor. *Seismological Research Letters*, 89(1):67–76, 11 2017.

[55] C. M. Mow-Lowry and D. Martynov. A 6d interferometric inertial isolation system. *Classical and Quantum Gravity*, 36(24):245006, nov 2019.

[56] A. S. Ubhi, J. Smetana, T. Zhang, S. Cooper, L. Prokhorov, J. Bryant, D. Hoyland, H. Miao, and D. Martynov. A six degree-of-freedom fused silica seismometer: design and tests of a metal prototype. *Classical and Quantum Gravity*, 39(1):015006, dec 2021.

[57] S. J. Cooper, C. J. Collins, A. C. Green, D. Hoyland, C. C. Speake, A. Freise, and C. M. Mow-Lowry. A compact, large-range interferometer for precision measurement and inertial sensing. *Classical and Quantum Gravity*, 35(9):095007, mar 2018.

[58] H. Tariq, A. Takamori, F. Vetrano, C. Wang, A. Bertolini, G. Calamai, R. DeSalvo, A. Gennai, L. Holloway, G. Losurdo, S. Márka, M. Mazzoni, F. Paoletti, D. Passuello, V. Sannibale, and R. Stanga. The linear variable differential transformer (lvdt) position sensor for gravitational wave interferometer low-frequency controls.

Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 489(1):570–576, 2002.

[59] G. Losurdo, G. Calamai, E. Cuoco, et al. Inertial control of the mirror suspensions of the VIRGO interferometer for gravitational wave detection. *Review of Scientific Instruments*, 72(9):3653–3661, sep 2001.

[60] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32:110–113, Jul 1928.

[61] H. B. Callen and T. A. Welton. Irreversibility and Generalized Noise. *Physical Review*, 83(1):34–40, jul 1951.

[62] H. B. Callen and R. F. Greene. On a Theorem of Irreversible Thermodynamics. *Physical Review*, 86(5):702–710, jun 1952.

[63] P. R. Saulson. Thermal noise in mechanical experiments. *Physical Review D*, 42(8):2437–2445, oct 1990.

[64] I. W. Martin and S. Reid. Coating thermal noise. In Gregory Harry, Timothy P. Bodiya, and Riccardo DeSalvo, editors, *Optical coatings and thermal noise in precision measurement*, chapter 4, pages 31–54. Cambridge University Press, Cambridge, 2012.

[65] A. V. Cumming, A. S. Bell, L. Barsotti, M. A. Barton, G. Cagnoli, D. Cook, L. Cunningham, M. Evans, G. D. Hammond, G. M. Harry, A. Heptonstall, J. Hough, R. Jones, R. Kumar, R. Mittleman, N. A. Robertson, S. Rowan, B. Shapiro, K. A. Strain, K. Tokmakov, C. Torrie, and A. A. van Veggel. Design and development of the advanced LIGO monolithic fused silica suspension. *Classical and Quantum Gravity*, 29(3):035003, feb 2012.

[66] G. González. Suspensions thermal noise in the LIGO gravitational wave detector. *Classical and Quantum Gravity*, 17(21):4409–4435, nov 2000.

[67] J.E. Logan, J. Hough, and N.A. Robertson. Aspects of the thermal motion of a mass suspended as a pendulum by wires. *Physics Letters A*, 183(2):145–152, 1993.

[68] G. I. González and P. R. Saulson. Brownian motion of a mass suspended by an anelastic wire. *The Journal of the Acoustical Society of America*, 96(1):207–212, 1994.

[69] V. B. Braginsky, V. P. Mitrofanov, and K. V. Tokmakov. On the thermal noise from the violin modes of the test mass suspension in gravitational wave antennae. *Physics Letters A*, 186(1):18–20, 1994.

[70] A. Gillespie and F. Raab. Suspension losses in the pendula of laser interferometer gravitational-wave detectors. *Physics Letters A*, 190(3):213–220, 1994.

[71] S. Rowan, R. Hutchins, A. McLaren, N. A. Robertson, S. M. Twyford, and J. Hough. The quality factor of natural fused quartz ribbons over a frequency range from 6 to 160 hz. *Physics Letters A*, 227(3):153–158, 1997.

[72] A. Gillespie and F. Raab. Thermally excited vibrations of the mirrors of laser interferometer gravitational-wave detectors. *Physical Review D*, 52(2):577–585, jul 1995.

[73] Yu. Levin. Internal thermal noise in the LIGO test masses: A direct approach. *Physical Review D*, 57(2):659–663, jan 1998.

[74] F. Bondu, P. Hello, and J.-Y. Vinet. Thermal noise in mirrors of interferometric gravitational wave antennas. *Physics Letters A*, 246(3-4):227–236, sep 1998.

[75] Y. T. Liu and K. S. Thorne. Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses. *Physical Review D*, 62(12):122002, nov 2000.

[76] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Physical Review D*, 78(10):102003, nov 2008.

[77] G. M. Harry, A. M. Gretarsson, P. R. Saulson, S. E. Kittelberger, S. D. Penn, W. J. Startin, S. Rowan, M. M. Fejer, D. R. M. Crooks, G. Cagnoli, J. Hough, and N. Nakagawa. Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings. *Classical and Quantum Gravity*, 19(5):897–917, mar 2002.

[78] I. W. Martin, R. Bassiri, R. Nawrodt, M. M. Fejer, A. Gretarsson, E. Gustafson, G. Harry, J. Hough, I. MacLaren, S. Penn, S. Reid, R. Route, S. Rowan, C. Schwarz, P. Seidel, J. Scott, and A. L. Woodcraft. Effect of heat treatment on mechanical dissipation in ta_2o_5 coatings. *Classical and Quantum Gravity*, 27(22):225020, oct 2010.

[79] J. Steinlechner, I. W. Martin, A. S. Bell, J. Hough, M. Fletcher, P. G. Murray, R. Robie, S. Rowan, and R. Schnabel. Silicon-based optical mirror coatings for ultrahigh precision metrology and sensing. *Phys. Rev. Lett.*, 120:263602, Jun 2018.

[80] Tara Fortier and Esther Baumann. 20 years of developments in optical frequency comb technology and applications. *Communications Physics*, 2(1):153, December 2019.

[81] J. A. Sidles and D. Sigg. Optical torques in suspended fabry–perot interferometers. *Physics Letters A*, 354(3):167–172, 2006.

[82] Huan Yang, Haixing Miao, Da-Shin Lee, Bassam Helou, and Yanbei Chen. Macroscopic quantum mechanics in a classical spacetime. *Phys. Rev. Lett.*, 110:170401, Apr 2013.

[83] Haixing Miao, Denis Martynov, Huan Yang, and Animesh Datta. Quantum correlations of light mediated by gravity. *Phys. Rev. A*, 101:063804, Jun 2020.

[84] J. Smetana and D. Martynov. Towards the standard quantum limit in a table-top interferometer. In *Proceedings of the GRavitational-waves Science&technology Symposium (GRASS) 2019.* Zenodo, April 2020.

[85] S. Slussarenko, M. M. Weston, H. M. Chrzanowski, L. K. Shalm, V. B. Verma, S. W. Nam, and G. J. Pryde. Unconditional violation of the shot-noise limit in photonic quantum metrology. *Nature Photonics*, 11(11):700–703, November 2017.

[86] P. Yang, B. Xie, and S. Feng. Subhertz interferometry at the quantum noise limit. *Opt. Lett.*, 44(9):2366–2369, May 2019.

[87] S. Davuluri and Y. Li. Shot-noise-limited interferometry for measuring a classical force. *Phys. Rev. A*, 98:043809, Oct 2018.

[88] H. Yu, D. Martynov, S. Vitale, M. Evans, D. Shoemaker, B. Barr, G. Hammond, S. Hild, J. Hough, S. Huttner, S. Rowan, B. Sorazu, L. Carbone, A. Freise, C. Mow-Lowry, K. L. Dooley, P. Fulda, H. Grote, and D. Sigg. Prospects for detecting gravitational waves at 5 hz with ground-based detectors. *Phys. Rev. Lett.*, 120:141102, Apr 2018.

[89] J. Cripe, N. Aggarwal, R. Lanza, A. Libson, R. Singh, P. Heu, D. Follman, G. D. Cole, N. Mavalvala, and T. Corbitt. Measurement of quantum back action in the audio band at room temperature. *Nature*, 568(7752):364–367, April 2019.

[90] S. Hild, M. Abernathy, F. Acernese, et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28(9):094013, apr 2011.

[91] Y. Michimura, K. Komori, Y. Enomoto, K. Nagano, A. Nishizawa, E. Hirose, M. Leonardi, E. Capocasa, N. Aritomi, Y. Zhao, R. Flaminio, T. Ushiba, T. Yamada, L.-W. Wei, H. Takeda, S. Tanioka, M. Ando, K. Yamamoto, K. Hayama, S. Haino, and K. Somiya. Prospects for improving the sensitivity of the cryogenic gravitational wave detector kagra. *Phys. Rev. D*, 102:022008, Jul 2020.

[92] R. X. Adhikari, K. Arai, A. F. Brooks, et al. A cryogenic silicon interferometer for gravitational-wave detection. *Classical and Quantum Gravity*, 37(16):165003, jul 2020.

[93] T. Westphal, G. Bergmann, A. Bertolini, M. Born, Y. Chen, A. V. Cumming, L. Cunningham, K. Dahl, C. Gräf, G. Hammond, G. Heinzel, S. Hild, S. Huttner, R. Jones, F. Kawazoe, S. Köhlenbeck, G. Kühn, H. Lück, K. Mossavi, J. H. Pöld, K. Somiya, A. M. van Veggel, A. Wanner, B. Willke, K. A. Strain, S. Goßler, and K. Danzmann. Design of the 10 m AEI prototype facility for interferometry studies. *Applied Physics B*, 106(3):551–557, mar 2012.

[94] H. Müller-Ebhardt, H. Rehbein, R. Schnabel, K. Danzmann, and Y. Chen. Entanglement of macroscopic test masses and the standard quantum limit in laser interferometry. *Phys. Rev. Lett.*, 100:013601, Jan 2008.

[95] R. Penrose. *The Emperor's New Mind*, chapter 8, page 475. Oxford University Press, 1989.

[96] V. B. Braginskii and Iu. I. Vorontsov. Quantum-mechanical limitations in macroscopic experiments and modern experimental technique. *Uspekhi Fizicheskikh Nauk*, 114:41–53, September 1974.

[97] C. M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, 23(8):1693–1708, apr 1981.

[98] R. Schnabel. Squeezed states of light and their applications in laser interferometers. *Physics Reports*, 684:1–51, 2017. Squeezed states of light and their applications in laser interferometers.

[99] J. Südbeck, S. Steinlechner, M. Korobko, and R. Schnabel. Demonstration of interferometer enhancement through einstein-podolsky-rosen entanglement. *Nature Photonics*, 14(4):240–244, April 2020.

[100] J. Aasi, J. Abadie, B. P. Abbott, et al. Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light. *Nature Photonics*, 7(8):613–619, August 2013.

[101] F. Acernese, M. Agathos, L. Aiello, et al. Increasing the astrophysical reach of the advanced virgo detector via the application of squeezed vacuum states of light. *Phys. Rev. Lett.*, 123:231108, Dec 2019.

[102] M. A. Arain and G. Mueller. Design of the advanced ligo recycling cavities. *Opt. Express*, 16(14):10018–10032, Jul 2008.

[103] A. Buonanno and Y. Chen. Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors. *Physical Review D*, 64(4):042006, jul 2001.

[104] A. Buonanno and Y. Chen. Scaling law in signal recycled laser-interferometer gravitational-wave detectors. *Phys. Rev. D*, 67:062002, Mar 2003.

[105] F. Ya. Khalili and Yu. Levin. Speed meter as a quantum nondemolition measuring device for force. *Phys. Rev. D*, 54:4735–4737, Oct 1996.

[106] V. B. Braginsky, M. L. Gorodetsky, F. Ya. Khalili, and K. S. Thorne. Dual-resonator speed meter for a free test mass. *Phys. Rev. D*, 61:044002, Jan 2000.

[107] P. Purdue. Analysis of a quantum nondemolition speed-meter interferometer. *Phys. Rev. D*, 66:022001, Jun 2002.

[108] P. Purdue and Y. Chen. Practical speed meter designs for quantum nondemolition gravitational-wave interferometers. *Phys. Rev. D*, 66:122004, Dec 2002.

[109] Y. Chen. Sagnac interferometer as a speed-meter-type, quantum-nondemolition gravitational-wave detector. *Phys. Rev. D*, 67:122004, Jun 2003.

[110] C. Gräf, B. W. Barr, A. S. Bell, F. Campbell, A. V. Cumming, S. L. Danilishin, N. A. Gordon, G. D. Hammond, J. Hennig., E. A. Houston, S. H. Huttner, R. A. Jones, S. S. Leavey, H. Lück, J. Macarthur, M. Marwick, S. Rigby, R. Schilling, B. Sorazu, A. Spencer, S. Steinlechner, K. A. Strain, and S. Hild. Design of a speed meter interferometer proof-of-principle experiment. *Classical and Quantum Gravity*, 31(21):215009, nov 2014.

[111] B. Dasgupta and T. S. Mruthyunjaya. The stewart platform manipulator: a review. *Mechanism and Machine Theory*, 35(1):15–40, 2000.

[112] T. Shimoda, S. Takano, C. P. Ooi, N. Aritomi, A. Shoda, Y. Michimura, and M. Ando. Torsion-bar antenna: a ground-based mid-frequency and low-frequency gravitational wave detector, 2018.

[113] A. S. Ubhi, L. Prokhorov, S. Cooper, C. Di Fronzo, J. Bryant, D. Hoyland, A. Mitchell, J. van Dongen, C. Mow-Lowry, A. Cumming, G. Hammond, and D. Martynov. Active platform stabilisation with a 6d seismometer, 2022.

[114] Y Aso, M Ando, S Otsuka, and K Tsubono. Active vibration isolation using a suspension point interferometer. *Journal of Physics: Conference Series*, 32(1):451, mar 2006.

[115] H. Kogelnik and T. Li. Laser beams and resonators. *Appl. Opt.*, 5(10):1550–1567, Oct 1966.

[116] D. Z. Anderson. Alignment of resonant optical cavities. *Appl. Opt.*, 23(17):2944–2949, Sep 1984.

[117] G. Mueller. Beam jitter coupling in advanced ligo. *Opt. Express*, 13(18):7118–7132, Sep 2005.

[118] P. Fritschel, N. Mavalvala, D. Shoemaker, D. Sigg, M. Zucker, and G. González. Alignment of an interferometric gravitational wave detector. *Appl. Opt.*, 37(28):6734–6747, Oct 1998.

[119] S. Braccini, C. Bradaschia, R. Del Fabbro, A. Di Virgilio, I. Ferrante, F. Fidecaro, R. Flaminio, A. Gennai, A. Giassi, A. Giazotto, G. Gorini, G. Losurdo, F. Palla, A. Pasqualetti, D. Passuello, R. Poggiani, G. Torelli, and Z. Zhang. Seismic vibrations mechanical filters for the gravitational waves detector virgo. *Review of Scientific Instruments*, 67(8):2899–2902, 1996.

[120] E. Coccia. Mechanical filter for the suspension of gravitational wave antennas. *Review of Scientific Instruments*, 53(2):148–153, 1982.

[121] L. Ju and D. G. Blair. Low resonant frequency cantilever spring vibration isolator for gravitational wave detectors. *Review of Scientific Instruments*, 65(11):3482–3488, 1994.

[122] C. Torrie. *Development of Suspensions for the GEO 600 Gravitational Wave Detector.* PhD thesis, University of Glasgow, 1999.

[123] K. Numata and K. Yamamoto. Cryogenics. In Gregory Harry, Timothy P. Bodiya, and Riccardo DeSalvo, editors, *Optical coatings and thermal noise in precision measurement*, chapter 8, pages 108–128. Cambridge University Press, Cambridge, 2012.

[124] T. Nakanishi, T. Tsuchiyama, H. Mitsuyasu, Y. Iwamoto, and S. Takaki. Effect of partial solution nitriding on mechanical properties and corrosion resistance in a type 316l austenitic stainless steel plate. *Materials Science and Engineering: A*, 460-461:186–194, 2007.

[125] R. A. Lula. *Stainless Steel*, pages 37–39. American Society for Metals, 1986.

[126] E. D. Marquardt, J. P. Le, and R. Radebaugh. *Cryogenic Material Properties Database*, pages 681–687. Springer US, Boston, MA, 2002.

[127] R. W. Powell, C. Y. Ho, and P. E. Liley. Thermal conductivity of selected materials. Technical report, U.S. Department of Commerce, National Bureau of Standards, 1968.

[128] K. Numata, S. Otsuka, M. Ando, and K. Tsubono. Intrinsic losses in various kinds of fused silica. *Classical and Quantum Gravity*, 19(7):1697–1702, mar 2002.

[129] A. Schroeter, R. Nawrodt, R. Schnabel, S. Reid, I. Martin, S. Rowan, C. Schwarz, T. Koettig, R. Neubert, M. Thürk, W. Vodel, A. Tünnermann, K. Danzmann, and P. Seidel. On the mechanical quality factors of cryogenic test masses from fused silica and crystalline quartz, 2007.

[130] S. Hild, S. Chelkowski, A. Freise, J. Franc, N. Morgado, R. Flaminio, and R. DeSalvo. A xylophone configuration for a third-generation gravitational wave detector. *Classical and Quantum Gravity*, 27(1):015003, dec 2009.

[131] K. G. Lyon, G. L. Salinger, C. A. Swenson, and G. K. White. Linear thermal expansion measurements on silicon from 6 to 340 k. *Journal of Applied Physics*, 48(3):865–868, 1977.

[132] G. Ciani, M. A. Arain, S. M. Aston, D. Feldbaum, P. Fulda, J. Gleason, M. Heintze, R. M. Martin, C. L. Mueller, D. M. Nanda Kumar, A. Pele, D. H. Reitze, P. Sainathan, D. B. Tanner, L. F. Williams, and G. Mueller. Small optic suspensions for advanced ligo input optics and other precision optical experiments. *Review of Scientific Instruments*, 87(11):114504, 2016.

[133] B. Abbott, R. Abbott, R. Adhikari, et al. Detector description and performance for the first coincidence observations between ligo and geo. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 517(1):154–179, 2004.

[134] I. W. Martin, R. Nawrodt, K. Craig, C. Schwarz, R. Bassiri, G. Harry, J. Hough, S. Penn, S. Reid, R. Robie, and S. Rowan. Low temperature mechanical dissipation of an ion-beam sputtered silica film. *Classical and Quantum Gravity*, 31(3):035019, jan 2014.

[135] X. Liu, T. H. Metcalf, Q. Wang, and D. M. Photiadis. Elastic properties of several silicon nitride films. *MRS Proceedings*, 989:0989–A22–01, 2007.

[136] G. D. Cole, W. Zhang, M. J. Martin, J. Ye, and M. Aspelmeyer. Tenfold reduction of brownian noise in high-reflectivity optical coatings. *Nature Photonics*, 7(8):644–650, August 2013.

[137] S. D. Penn, M. M. Kinley-Hanlon, I. A. O. MacMillan, P. Heu, D. Follman, C. Deutsch, G. D. Cole, and G. M. Harry. Mechanical ringdown studies of large-area substrate-transferred gaas/algaas crystalline coatings. *J. Opt. Soc. Am. B*, 36(4):C15–C21, Apr 2019.

[138] M. J. Lawrence, B. Willke, M. E. Husman, E. K. Gustafson, and R. L. Byer. Dynamic response of a fabry–perot interferometer. *J. Opt. Soc. Am. B*, 16(4):523–532, Apr 1999.

[139] X. Li, J. Smetana, A. S. Ubhi, J. Bentley, Y. Chen, Y. Ma, H. Miao, and D. Martynov. Enhancing interferometer sensitivity without sacrificing bandwidth and stability: Beyond single-mode and resolved-sideband approximation. *Phys. Rev. D*, 103:122001, Jun 2021.

[140] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Phys. Rev. X*, 9:031040, Sep 2019.

[141] R. Abbott, T. D. Abbott, S. Abraham, et al. Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Phys. Rev. X*, 11:021053, Jun 2021.

[142] I. Kamaretsos, M. Hannam, S. Husa, and B. S. Sathyaprakash. Black-hole hair loss: Learning about binary progenitors from ringdown signals. *Phys. Rev. D*, 85:024018, Jan 2012.

[143] I. Kamaretsos, M. Hannam, and B. S. Sathyaprakash. Is black-hole ringdown a memory of its progenitor? *Phys. Rev. Lett.*, 109:141102, Oct 2012.

[144] A. Taracchini, A. Buonanno, G. Khanna, and S. A. Hughes. Small mass plunging into a kerr black hole: Anatomy of the inspiral-merger-ringdown waveforms. *Phys. Rev. D*, 90:084025, Oct 2014.

[145] B. P. Abbott, R. Abbott, T. D. Abbott, et al. First targeted search for gravitational-wave bursts from core-collapse supernovae in data of first-generation laser interferometer detectors. *Phys. Rev. D*, 94:102001, Nov 2016.

[146] S. A. Hughes, A. Apte, G. Khanna, and H. Lim. Learning about black hole binaries from their ringdown spectra. *Phys. Rev. Lett.*, 123:161101, Oct 2019.

[147] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.

[148] H. Miao, H. Yang, and D. Martynov. Towards the design of gravitational-wave detectors for probing neutron-star physics. *Phys. Rev. D*, 98:044044, Aug 2018.

[149] D. Martynov, H. Miao, H. Yang, F. H. Vivanco, E. Thrane, R. Smith, P. Lasky, W. E. East, R. Adhikari, A. Bauswein, A. Brooks, Y. Chen, T. Corbitt, A. Freise, H. Grote, Yu. Levin, C. Zhao, and A. Vecchio. Exploring the sensitivity of gravitational wave detectors to neutron star physics. *Phys. Rev. D*, 99:102004, May 2019.

[150] C. D. Ott. The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, 26(6):063001, feb 2009.

[151] K. Kotake. Multiple physical elements to determine the gravitational-wave signatures of core-collapse supernovae. *Comptes Rendus Physique*, 14(4):318 – 351, 2013. Gravitational waves / Ondes gravitationnelles.

[152] D. Radice, V. Morozova, A. Burrows, D. Vartanyan, and H. Nagakura. Characterizing the gravitational wave signal from core-collapse supernovae. *The Astrophysical Journal*, 876(1):L9, apr 2019.

[153] M. Punturo, M. Abernathy, F. Acernese, et al. The einstein telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, sep 2010.

[154] M. Maggiore, C. Van Den Broeck, N. Bartolo, et al. Science case for the einstein telescope. *Journal of Cosmology and Astroparticle Physics*, 2020(03):050–050, mar 2020.

[155] D. Reitze, R. X. Adhikari, S. Ballmer, et al. Cosmic explorer: The u.s. contribution to gravitational-wave astronomy beyond ligo, 2019.

[156] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Exploring the sensitivity of next generation gravitational wave detectors. *Classical and Quantum Gravity*, 34(4):044001, jan 2017.

[157] W. DeRocco and A. Hook. Axion interferometry. *Phys. Rev. D*, 98:035021, Aug 2018.

[158] I. Obata, T. Fujita, and Y. Michimura. Optical ring cavity search for axion dark matter. *Phys. Rev. Lett.*, 121:161301, Oct 2018.

[159] H. Liu, B. D. Elwood, M. Evans, and J. Thaler. Searching for axion dark matter with birefringent cavities. *Phys. Rev. D*, 100:023548, Jul 2019.

[160] D. Martynov and H. Miao. Quantum-enhanced interferometry for axion searches. *Phys. Rev. D*, 101:095034, May 2020.

[161] V. B. Braginsky, M. L. Gorodetsky, F. Ya. Khalili, and K. S. Thorne. Energetic quantum limit in large-scale interferometers. In *AIP Conference Proceedings*, volume 523, pages 180–190. AIP, 2000.

[162] M. Tsang, H. M. Wiseman, and C. M. Caves. Fundamental quantum limit to waveform estimation. *Phys. Rev. Lett.*, 106:090401, Mar 2011.

[163] H. Miao, Y. Ma, C. Zhao, and Y. Chen. Enhancing the bandwidth of gravitational-wave detectors with unstable optomechanical filters. *Physical review letters*, 115(21):211104, 2015.

[164] J. Bentley, P. Jones, D. Martynov, A. Freise, and H. Miao. Converting the signal-recycling cavity into an unstable optomechanical filter to enhance the detection bandwidth of gravitational-wave detectors. *Phys. Rev. D*, 99:102001, May 2019.

[165] X. Li, M. Goryachev, Y. Ma, M. E. Tobar, C. Zhao, R. X. Adhikari, and Y. Chen. Broadband sensitivity improvement via coherent quantum feedback with pt symmetry, 2020.

[166] Y. Chen. Macroscopic quantum mechanics: theory and experimental concepts of optomechanics. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 46(10):104001, may 2013.

[167] B. L. Schumaker and C. M. Caves. New formalism for two-photon quantum optics. ii. mathematical foundation and compact notation. *Phys. Rev. A*, 31:3093–3111, May 1985.

[168] K. J. Åström and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers.* Princeton University Press, 2010.

[169] J. L. Cieśliński and B. Ratkiewicz. On simulations of the classical harmonic oscillator equation by difference equations. *Advances in Difference Equations*, May 2006.

[170] T. Zhang, J. Smetana, Y. Chen, J. Bentley, D. Martynov, H. Miao, W. E. East, and H. Yang. Toward observing neutron star collapse with gravitational wave detectors. *Phys. Rev. D*, 103:044063, Feb 2021.

[171] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817a. *The Astrophysical Journal*, 848(2):L13, oct 2017.

[172] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Multi-messenger observations of a binary neutron star merger. *The Astrophysical Journal*, 848(2):L12, oct 2017.

[173] B. P. Abbott, R. Abbott, T. D. Abbott, et al. A gravitational-wave standard siren measurement of the hubble constant. *Nature*, 551(7678):85–88, 11 2017.

[174] B. F. Schutz. Determining the hubble constant from gravitational wave observations. *Nature*, 323(6086):310–311, 09 1986.

[175] D. E. Holz and S. A. Hughes. Using gravitational-wave standard sirens. *The Astrophysical Journal*, 629(1):15–22, aug 2005.

[176] N. Dalal, D. E. Holz, S. A. Hughes, and B. Jain. Short grb and binary black hole standard sirens as a probe of dark energy. *Phys. Rev. D*, 74:063006, Sep 2006.

[177] S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers. EXPLORING SHORT GAMMA-RAY BURSTS AS GRAVITATIONAL-WAVE STANDARD SIRENS. *The Astrophysical Journal*, 725(1):496–514, nov 2010.

[178] S. Nissanke, D. E. Holz, N. Dalal, S. A. Hughes, J. L. Sievers, and C. M. Hirata. Determining the hubble constant from gravitational wave observations of merging compact binaries. *arXiv preprint arXiv:1307.2638*, 2013.

[179] E. E. Flanagan and T. Hinderer. Constraining neutron star tidal Love numbers with gravitational wave detectors. *Phys. Rev. D*, 77:021502, 2008.

[180] K. Chatziioannou. Neutron star tidal deformability and equation of state constraints. *Gen. Rel. Grav.*, 52(11):109, 2020.

[181] Z. Pan, Z. Lyu, B. Bonga, N. Ortiz, and H. Yang. Probing Crust Meltdown in Inspiraling Binary Neutron Stars. *Phys. Rev. Lett.*, 125(20):201102, 2020.

[182] E. Poisson. Gravitomagnetic tidal resonance in neutron-star binary inspirals. *Phys. Rev. D*, 101(10):104028, 2020.

[183] M. Agathos, F. Zappa, S. Bernuzzi, A. Perego, M. Breschi, and D. Radice. Inferring Prompt Black-Hole Formation in Neutron Star Mergers from Gravitational-Wave Data. *Phys. Rev. D*, 101(4):044006, 2020.

[184] D. Radice, S. Bernuzzi, and A. Perego. The dynamics of binary neutron star mergers and gw170817. *Annual Review of Nuclear and Particle Science*, 70(1):95–119, 2020.

[185] V. Paschalidis. General relativistic simulations of compact binary mergers as engines for short gamma-ray bursts. *Class. Quant. Grav.*, 34(8):084002, 2017.

[186] B.D. Metzger, E. Quataert, and T.A. Thompson. Short Duration Gamma-Ray Bursts with Extended Emission from Proto-Magnetar Spin-Down. *Mon. Not. Roy. Astron. Soc.*, 385:1455, 2008.

[187] N. Bucciantini, B. D. Metzger, T. A. Thompson, and E. Quataert. Short gamma-ray bursts with extended emission from magnetar birth: jet formation and collimation. *Monthly Notices of the Royal Astronomical Society*, 419(2):1537–1545, 2012.

[188] A. Rowlinson, P. T. O'Brien, B. D. Metzger, N. R. Tanvir, and A. J. Levan. Signatures of magnetar central engines in short GRB lightcurves. *Mon. Not. Roy. Astron. Soc.*, 430:1061, 2013.

[189] B. D. Metzger. Kilonovae. *Living Rev. Rel.*, 20(1):3, 2017.

[190] D. M. Siegel. GW170817–the first observed neutron star merger and its kilonova: implications for the astrophysical site of the r-process. *Eur. Phys. J. A*, 55(11):203, 2019.

[191] N. Sarin, P. D. Lasky, and G. Ashton. Gravitational waves or deconfined quarks: What causes the premature collapse of neutron stars born in short gamma-ray bursts? *Physical Review D*, 101(6):063021, 2020.

[192] M. Lucca and L. Sagunski. The lifetime of binary neutron star merger remnants, 2020.

[193] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Search for gravitational waves from a long-lived remnant of the binary neutron star merger GW170817. *The Astrophysical Journal*, 875(2):160, apr 2019.

[194] K. Ackley, V. B. Adya, P. Agrawal, et al. Neutron star extreme matter observatory: A kilohertz-band gravitational-wave detector in the global network, 2020.

[195] The ET Science Team. *Einstein gravitational wave Telescope conceptual design.* European Commission, Jun 2011.

[196] ET Steering Committee Editorial Team. *Einstein Telescope design report update 2020.* Einstein Telescope Collaboration, Sep 2020.

[197] W. E. East, V. Paschalidis, F. Pretorius, and A. Tsokaros. Binary neutron star mergers: Effects of spin and post-merger dynamics. *Physical Review D*, 100(12):124042, 2019.

[198] M. Lucca, L. Sagunski, F. Guercilena, and C. M. Fromm. Shedding light on the angular momentum evolution of binary neutron star merger remnants: a semi-analytic model, 2020.

[199] L. Baiotti and L. Rezzolla. Binary neutron star mergers: a review of einstein's richest laboratory. *Reports on Progress in Physics*, 80(9):096901, 2017.

[200] M. Breschi, S. Bernuzzi, F. Zappa, M. Agathos, A. Perego, D. Radice, and A. Nagar. Kilohertz gravitational waves from binary neutron star remnants: Time-domain model and constraints on extreme matter. *Physical Review D*, 100(10):104029, 2019.

[201] W. E. East, F. Pretorius, and B. C. Stephens. Hydrodynamics in full general relativity with conservative AMR. *Phys. Rev. D*, 85:124010, 2012.

[202] W. E. East and L. Lehner. Fate of a neutron star with an endoparasitic black hole and implications for dark matter. *Physical Review D*, 100(12):124026, 2019.

[203] N. Stergioulas and J. L. Friedman. Comparing models of rapidly rotating relativistic stars constructed by two numerical methods. *Astrophys. J.*, 444:306, 1995.

[204] L. Baiotti, I. Hawke, and L. Rezzolla. On the Gravitational radiation from the collapse of neutron stars to rotating black holes. *Class. Quant. Grav.*, 24:S187–S206, 2007.

[205] A. Bauswein and H.-T. Janka. Measuring neutron-star properties via gravitational waves from neutron-star mergers. *Physical review letters*, 108(1):011101, 2012.

[206] K. Hotokezaka, K. Kiuchi, K. Kyutoku, T. Muranushi, Y. Sekiguchi, M. Shibata, and K. Taniguchi. Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform. *Physical Review D*, 88(4):044026, 2013.

[207] L. Lehner, S. L. Liebling, C. Palenzuela, O. L. Caballero, E. O'Connor, M. Anderson, and D. Neilsen. Unequal mass binary neutron star mergers and multimessenger signals. *Classical and Quantum Gravity*, 33(18):184002, 2016.

[208] H. Yang, V. Paschalidis, K. Yagi, L. Lehner, F. Pretorius, and N. Yunes. Gravitational wave spectroscopy of binary neutron star merger remnants with mode stacking. *Physical Review D*, 97(2):024049, 2018.

[209] K. Kiuchi, Y. Sekiguchi, M. Shibata, and K. Taniguchi. Long-term general relativistic simulation of binary neutron stars collapsing to a black hole. *Physical Review D*, 80(6):064037, 2009.

[210] L. Rezzolla, L. Baiotti, B. Giacomazzo, D. Link, and J. A. Font. Accurate evolutions of unequal-mass neutron-star binaries: properties of the torus and short grb engines. *Classical and Quantum Gravity*, 27(11):114105, 2010.

[211] S. Bernuzzi, T. Dietrich, W. Tichy, and B. Bruegmann. Mergers of binary neutron stars with realistic spin. *Physical Review D*, 89(10):104021, 2014.

[212] W. Kastaun, F. Galeazzi, D. Alic, L. Rezzolla, and J. A. Font. Black hole from merging binary neutron stars: How fast can it spin? *Physical Review D*, 88(2):021501, 2013.

[213] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, M. Shibata, and K. Taniguchi. Dynamical mass ejection from the merger of asymmetric binary neutron stars: Radiation-hydrodynamics study in general relativity. *Physical Review D*, 93(12):124046, 2016.

[214] A. Bauswein, T. W. Baumgarte, and H.-T. Janka. Prompt merger collapse and the maximum mass of neutron stars. *Physical review letters*, 111(13):131101, 2013.

[215] L. Rezzolla, E. R. Most, and L. R. Weih. Using gravitational-wave observations and quasi-universal relations to constrain the maximum mass of neutron stars. *The Astrophysical Journal Letters*, 852(2):L25, 2018.

[216] H. Yang, D. A. Nichols, F. Zhang, A. Zimmerman, Z. Zhang, and Y. Chen. Quasinormal-mode spectrum of kerr black holes and its geometric interpretation. *Physical review D*, 86(10):104006, 2012.

[217] L. Baiotti, B. Giacomazzo, and L. Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: Prompt and delayed collapse to a black hole. *Physical Review D*, 78(8):084033, 2008.

[218] R. Essick, S. Vitale, and M. Evans. Frequency-dependent responses in third generation gravitational-wave detectors. *Phys. Rev. D*, 96:084004, Oct 2017.

[219] H. Miao, N. D. Smith, and M. Evans. Quantum limit for laser interferometric gravitational-wave detectors from optical dissipation. *Phys. Rev. X*, 9:011053, Mar 2019.

[220] T. Zhang, J. Bentley, and H. Miao. A broadband signal recycling scheme for saturating the quantum limit from optical losses, 2020.

[221] D. V. Martynov, V. V. Frolov, S. Kandhasamy, et al. Quantum correlation measurements in interferometric gravitational-wave detectors. *Phys. Rev. A*, 95:043831, Apr 2017.

[222] The LIGO Scientific Collaboration. GWINC-the gravitational wave interferometer noise calculator, 2020.

[223] Z. Michael and S. Whitcomb. Measurement of optical path fluctuations due to residual gas in the ligo 40 meter interferometer. *LIGO-DCC*, Sep 1997.

[224] M. L. Gorodetsky. Thermal noises and noise compensation in high-reflection multilayer coating. *Physics Letters A*, 372(46):6813 – 6822, 2008.

[225] R. F. Greene and H. B. Callen. On the formalism of thermodynamic fluctuation theory. *Phys. Rev.*, 83:1231–1235, Sep 1951.

[226] W. Yam, S. Gras, and M. Evans. Multimaterial coatings with reduced thermal noise. *Phys. Rev. D*, 91:042002, Feb 2015.

[227] S. Gras, H. Yu, W. Yam, D. Martynov, and M. Evans. Audio-band coating thermal noise measurement for advanced ligo with a multimode optical resonator. *Phys. Rev. D*, 95:022001, Jan 2017.

[228] J. Franc, N. Morgado, R. Flaminio, R. Nawrodt, I. Martin, L. Cunningham, A. Cumming, S. Rowan, and J. Hough. Mirror thermal noise in laser interferometer gravitational wave detectors operating at room and cryogenic temperature, 2009.

[229] J. R. Hutchinson. Vibrations of Solid Cylinders. *Journal of Applied Mechanics*, 47(4):901–907, 12 1980.

[230] H. Yang, K. Yagi, J. Blackman, L. Lehner, V. Paschalidis, F. Pretorius, and N. Yunes. Black hole spectroscopy with coherent mode stacking. *Physical review letters*, 118(16):161101, 2017.

[231] Jiri Smetana, Rebecca Walters, Sophie Bauchinger, Amit Singh Ubhi, Sam Cooper, David Hoyland, Richard Abbott, Christoph Baune, Peter Fritchel, Oliver Gerberding, Semjon Köhnke, Haixing Miao, Sebastian Rode, and Denis Martynov. Compact michelson interferometers with subpicometer sensitivity. *Phys. Rev. Applied*, 18:034040, Sep 2022.

[232] R. Magee, D. Chatterjee, L. P. Singer, et al. First demonstration of early warning gravitational-wave alerts. *The Astrophysical Journal Letters*, 910(2):L21, apr 2021.

[233] M. Branchesi. Multi-messenger astronomy: gravitational waves, neutrinos, photons, and cosmic rays. *Journal of Physics: Conference Series*, 718:022004, may 2016.

[234] P. Amaro-Seoane, J. R. Gair, M. Freitag, et al. Intermediate and extreme mass-ratio inspirals—astrophysics, science applications and detection using LISA. *Classical and Quantum Gravity*, 24(17):R113–R169, aug 2007.

[235] K. Venkateswara, C. A. Hagedorn, M. D. Turner, T. Arp, and J. H. Gundlach. A high-precision mechanical absolute-rotation sensor. *Review of Scientific Instruments*, 85(1):015005, 2014.

[236] M. P. Ross, K. Venkateswara, C. A. Hagedorn, C. J. Leupold, P. W. F. Forsyth, J. D. Wegner, E. A. Shaw, J. G. Lee, and J. H. Gundlach. A low-frequency torsion pendulum with interferometric readout. *Review of Scientific Instruments*, 92(5):054502, 2021.

[237] A. Staley, D. Martynov, R. Abbott, et al. Achieving resonance in the advanced LIGO gravitational-wave interferometer. *Classical and Quantum Gravity*, 31(24):245010, nov 2014.

[238] E. A. Shaw, M. P. Ross, C. A. Hagedorn, E. G. Adelberger, and J. H. Gundlach. Torsion-balance search for ultralow-mass bosonic dark matter. *Phys. Rev. D*, 105:042007, Feb 2022.

[239] S. Thomas, C. Bredendiek, and N. Pohl. A sige-based 240-ghz fmcw radar system for high-resolution measurements. *IEEE Transactions on Microwave Theory and Techniques*, 67(11):4599–4609, 2019.

[240] W. Yu, P. Pfeiffer, A. Morsali, J. Yang, and J. Fontaine. Comb-calibrated frequency sweeping interferometry for absolute distance and vibration measurement. *Opt. Lett.*, 44(20):5069–5072, Oct 2019.

[241] M. Sosin, H. Mainaud-Durand, V. Rude, and J. Rutkowski. Frequency sweeping interferometry for robust and reliable distance measurements in harsh accelerator environment. In Erik Novak and James D. Trolinger, editors, *Applied Optical Metrology III*, volume 11102, pages 145 – 161. International Society for Optics and Photonics, SPIE, 2019.

[242] J. Watchi, S. Cooper, B. Ding, C. M. Mow-Lowry, and C. Collette. Contributed review: A review of compact interferometers. *Review of Scientific Instruments*, 89(12):121501, 2018.

[243] R. M. de la Rue, R. F. Humphryes, I. M. Mason, and E. A. Ash. Acoustic-surface-wave amplitude and phase measurements using laser probes. *Proceedings of the Institution of Electrical Engineers*, 119:117–126(9), February 1972.

[244] C. Weichert, P. Köchert, R. Köning, J. Flügge, B. Andreas, U. Kuetgens, and A. Yacoot. A heterodyne interferometer with periodic nonlinearities smaller than ±10 pm. *Measurement Science and Technology*, 23(9):094005, jul 2012.

[245] M. Armano, H. Audley, G. Auger, et al. Sub-femto-g free fall for space-based gravitational wave observatories: Lisa pathfinder results. *Phys. Rev. Lett.*, 116:231101, Jun 2016.

[246] D. Royer and E. Dieulesaint. Optical detection of sub-angstrom transient mechanical displacements. *IEEE 1986 Ultrasonics Symposium*, pages 527–530, 1986.

[247] S. J. Cooper, C. J. Collins, L. Prokhorov, J. Warner, D. Hoyland, and C. M. Mow-Lowry. Interferometric sensing of a commercial geophone. *Classical and Quantum Gravity*, 39(7):075023, mar 2022.

[248] K.-S. Isleif, G. Heinzel, M. Mehmet, and O. Gerberding. Compact multifringe interferometry with subpicometer precision. *Physical Review Applied*, 12(3), Sep 2019.

[249] G. Heinzel, F. Guzmán Cervantes, A. F. García Marín, J. Kullmann, W. Feng, and K. Danzmann. Deep phase modulation interferometry. *Opt. Express*, 18(18):19076–19086, Aug 2010.

[250] O. Gerberding, K.-S. Isleif, M. Mehmet, K. Danzmann, and G. Heinzel. Laser-frequency stabilization via a quasimonolithic mach-zehnder interferometer with arms of unequal length and balanced dc readout. *Phys. Rev. Applied*, 7:024027, Feb 2017.

[251] O. Gerberding and K.-S. Isleif. Ghost beam suppression in deep frequency modulation interferometry for compact on-axis optical heads. *Sensors*, 21(5), 2021.

[252] L. R. Watkins and M. J. Collett. Ellipse fitting for interferometry. part 2: experimental realization. *Appl. Opt.*, 53(32):7697–7703, Nov 2014.

[253] M. J. Collett and L. R. Watkins. Ellipse fitting for interferometry. part 3: dynamic method. *J. Opt. Soc. Am. A*, 32(3):491–496, Mar 2015.

[254] K.-S. Isleif, O. Gerberding, T. S. Schwarze, M. Mehmet, G. Heinzel, and F. Guzmán Cervantes. Experimental demonstration of deep frequency modulation interferometry. *Opt. Express*, 24(2):1676–1684, Jan 2016.

[255] R. Fleddermann, C. Diekmann, F. Steier, M. Tröbs, G. Heinzel, and K. Danzmann. Sub-pm$\sqrt{\text{hz}}^{-1}$ non-reciprocal noise in the LISA backlink fiber. *Classical and Quantum Gravity*, 35(7):075007, feb 2018.