Development of Optical Suspensions and Laser Technology for a Cryogenic Interferometer Prototype

Victoria Graham

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

SCHOOL OF PHYSICS & ASTRONOMY

COLLEGE OF SCIENCE & ENGINEERING

University of Glasgow

2024

Acknowledgements

This thesis is dedicated to the many people whose support has contributed to its completion. In the main, this thesis should then be dedicated to Andrew Spencer and Bryan Barr. You have both spent countless hours explaining to me every concept I know about interferometry, experimental design, optical suspensions, thesis writing techniques, and more. Many of these points took several attempts to sink in, but nonetheless these discussions were always met with kind humour and patience.

Hours of discussion and support have also come from the exceptional experimental group at the IGR. I would like to thank in particular the contributions from Giles Hammond, Alan Cumming, Graeme Eddolls, Mark Barton, Thejas Seetharamu, Russell Jones, Jennifer Docherty, and Ardiana Nela. Thanks should also be extended to the members of the wider IGR research group, with special mentions owed to Daniel Williams, Rachel Gray, Joe Bayley, and Jennifer Wright. It has been a pleasure working with you all.

To the members of room 253: Leigh Smith, Jess Irwin, Ross Johnston, Federico Stachurski, Thejas Seetharamu, Storm Colloms, Naren Nagarajan, Eungwang Harry Seo, Josh Sharkey. You have turned years of work and months of writing into a truly enjoyable experience. If they would let me keep writing so I could stay in 253, I would.

To the Centre for Gravitational Astrophysics at the Australian National University, thank you for making Canberra feel like a home to me for the short time that I was there. I am especially grateful to Johannes Eichholz for his guidance and supervision during this period.

Finally, I would like to pay recognition to the invaluable contributions of my Mum, who had to endure near daily updates on thesis writing, my Dad, who inspired me to pursue scientific research, and Zoe, who, surprisingly, can provide a lot of wisdom when needed.

Contents

List of Tables

List of Figures

Preface

This thesis is an account of work carried out at the Institute for Gravitational Research (IGR) at the University of Glasgow from October 2020 to November 2024 towards the development of the Glasgow Cryogenic Interferometer Facility. This facility will support research into crucial technologies for cryogenic gravitational wave detectors such as silicon suspended optics, and the development of near-infrared lasers. Also included is the work conducted during a six month visit to the Centre for Gravitational Astrophysics at the Australian National University from January to July 2023. It is an account of work done by the author, except where explicit references are made to the contributions of others.

Chapter 1 provides an introduction to gravitational wave (GW) astrophysics which includes GW sources and the main scientific results from current detections. Interferometry is outlined as a method of gravitational wave detection. The current generation of gravitational wave detectors, their limiting noise sources, and future plans for third generation detectors is discussed. The Glasgow Cryogenic Interferometer Facility (GCIF) is introduced as a prototype facility which aims to investigate monolithic silicon suspended optics at cryogenic temperatures within a double-cavity interferometer, and develop laser technologies and noise-reduction schemes at 1550 nm and 2000 nm wavelengths.

Chapter 2 gives a deeper introduction to suspension dynamics and thermal displacement noise which is necessary for the understanding of Chapters 3 and 4. The suspension mechanical loss contribution from material bulk loss, surface loss, thermoelastic loss, and jointing loss is shown to contribute to suspension thermal noise. The design of a low noise suspension is possible through reducing energy dissipation in high-loss regions, creating a high-Q system which contains thermal noise in suspension resonances, and through use of low-noise suspension materials. Silicon and sapphire are discussed as the material choices for cryogenic interferometer detectors.

The explanations in this chapter were developed by the author with the support and advice of Dr. Alan Cumming, Dr. Graeme Eddolls, and Prof. Giles Hammond.

Chapter 3 discusses the development of two room temperature double suspensions which will be used in the Glasgow Cryogenic Interferometry Facility for beam steering and as 10 m cavity optics. This work characterised and developed two pre-existing suspension designs, using a dynamical model to calculate resonant frequencies and frequency response functions which were compared to experimental noise requirements. Modifications were designed to ensure the suspensions met GCIF length noise requirements through a reduction in the vertical seismic motion coupling to the optic. A characterisation methodology was developed for the 300 g double suspension, which allowed measurements of force-displacement and vertical-vertical transfer functions to be taken both pre- and post-modification. This work included an investigation into the origin and mitigation of unexpected resonances observed during the characterisation. The optic length noise was then calculated for the unmodified and modified suspensions to demonstrate that the modification process would meet the experimental noise requirements in the frequency band of interest.

The auxiliary suspension was originally designed by Dr. Jan-Simon Hennig and Mr. Russell Jones, while the 300g double suspension was originally designed by Dr. Neil Gordon and Mr. Russell Jones. An existing Mathematica model was used as the basis of the suspension models in this body of work. This original model was developed by Dr. Mark Barton and was adapted for use for the GCIF suspensions by the author with the help of Dr. Barton. The experimental procedures in this chapter were conducted with the help and supervision of Dr. Bryan Barr with advice and input from Dr. Andrew Spencer. Dr. Kristian Anastasiou, Mr. Steven O'Shea and Dr. Abhinav Prasad designed and built the custom 5 V biased unipolar driver which allowed the necessary drive signal for the measurement of the suspension vertical-vertical transfer function. The suspension modifications were designed by the author with early input from Dr. Eric Oelker and later discussions with Dr. Andrew Spencer, Dr. Bryan Barr, Prof. Giles Hammond, and with technical advice from Mr. Russell Jones. The author would also like to thank Dr. Iain Martin and Mr. Ross Johnston for support with the calculation of coating thermal noise which was used as the experimental noise requirement throughout this chapter.

Chapter 4 outlines the design process for a cryogenic triple-stage suspension with 1 kg silicon optics, suitable for use in the GCIF cryogenic test cavity. The cryogenic suspension was designed to meet seismic noise, heat extraction, and suspension displacement thermal noise requirements at both the 123 K and 18 K operational temperatures. The initial suspension design was based on experimental sensitivity and space requirements, material availability, and ease of manufacture, with further development performed though an investigation of the number of blade spring stages necessary for the required level of seismic isolation. The stability of a compact nested pendulum platform for further seismic isolation was considered. In the heat extraction modelling section of the chapter, two test mass cooling schemes were devised. The first was for the radiatively dominated 123 K phase of operation which utilises a heat shield to hold the test mass at 123 K with minimal thermal gradients across the suspension. The second scheme was for the conductively dominated regime of 18 K, where a chain of heat links were modelled to cool the suspension reaction chain and the upper two masses of the test chain. It was shown that the test mass could be conductively cooled through crystalline fibres to the intermediate mass while accounting for bonding and phonon-boundary scattering effects, a strategy which maintains low seismic and thermal noise of the suspension. Finally, test mass displacement thermal noise calculations were performed for the temperature cases of 123 K and 18 K which compared both silicon and sapphire suspension fibres of different jointing scenarios. These results showed the domination of joint loss as a contributor to suspension thermal noise and demonstrated that both silicon and sapphire fibres would meet the GCIF experimental sensitivity requirements.

The suspension was designed by the author with advice and support at the early stages from Dr. Eric Oelker, and later from Prof. Giles Hammond, Dr. Andrew Spencer, Dr. Alan Cumming, Dr. Graeme Eddolls, and Dr. Mark Barton. Seismic data was provided by Leiden Cryogenics and was used by the author to calculate the optic length noise coupling from a cryocooler-dominated seismic noise floor.

Chapter 5 discusses the characterisation and stabilisation of three 2 µm external cavity diode lasers (ECDLs), covering work performed at both the University of Glasgow and the Australian National University. Measurements of laser output power, intensity noise, frequency noise, and the response to modulation were presented for the characterisation process. The first intensity and frequency stabilisation measurements were demonstrated for a 2 µm ECDL, implemented via servo loops which fed back to the diode injection current. An intensity noise reduction of 10× was demonstrated at the University of Glasgow across 10 Hz – 1 kHz (ignoring the large suppression of noise peaks), and a frequency noise suppression of 200× was demonstrated at the Australian National University between 10 Hz – 1 kHz.

The ECDL design used throughout this work was developed by Dr. Johannes Eichholz and Dr. Disha Kapasi. The low noise servo electronics used at the University of Glasgow were designed by Dr. Eric Oelker based upon a design used in the LIGO Collaboration (developed by Dr. Daniel Sigg and Dr. Paul Schwinberg), which was built and tested by Dr. Oelker and the author. The experimental work in this chapter was conducted by the author with the help and supervision of Dr. Bryan Barr and Dr. Andrew Spencer at the University of Glasgow, and Dr. Johannes Eichholz while at the Australian National University. The author would also like to thank Dr. Stephen Webster for advice on the modelling and design of optical systems.

Chapter 6 provides an overview of the work carried out in the earlier chapters, containing the main results and conclusions from each stage of work. This chapter collates length noise estimates from previous chapters to provide a noise budget for the GCIF.

Declaration

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

Chapter 1

GRAVITATIONAL WAVE ASTRONOMY AND INTERFEROMETRIC DETECTION

1.1 The Origin of Gravitational Waves

For most of history astronomers have relied on observations of photons, neutrinos, and cosmic rays to learn more about our universe. The discovery of gravitational waves (GW) in 2015 heralded a new era of astronomy, allowing scientists to see our surroundings in an entirely new way [2]. Already, gravitational waves have observed the merger of black holes and neutron stars, allowed testing of general relativity, and provided a new measurement of the Hubble constant with only 9 years of data [43, 45, 46]. With more sensitive detectors it is hoped that gravitational waves will gather observations of the early universe, test alternate theories of gravity, and even explore the dimensions of the Universe [70, 182].

Albert Einstein first predicted the existence of gravitational waves shortly after his seminal 1915 publication on the theory of general relativity [75]. In this 1915 paper he proposed a four-dimensional spacetime in which the presence of matter curves the spacetime around it, with this curvature determining the path taken by the matter. It is perhaps not a far leap to imagine the acceleration of masses creating distortions which propagate though the gravitational field. Indeed, it only took Einstein until 1916 to publish his paper 'Approximative Integration of the Field Equations of Gravitation' in which he writes 'we shall investigate gravitational waves and how they originate' [76].

In the years since Einstein's predictions astronomers have done just that. Gravitational waves are known to originate from any acceleration of mass without spherical symmetry, and in their lowest order emit quadrupolar radiation (a monopole or dipole GW mode would violate the laws of mass conservation, and momentum conservation respectively) [166]. This quadrupolar radiation has two possible polarisations at 45° to one anther: h_+ and $h_\times$. The effect of these orthogonal polarisations on a ring of particles can be seen in Figure 1.1, with the GW travelling out of the page.

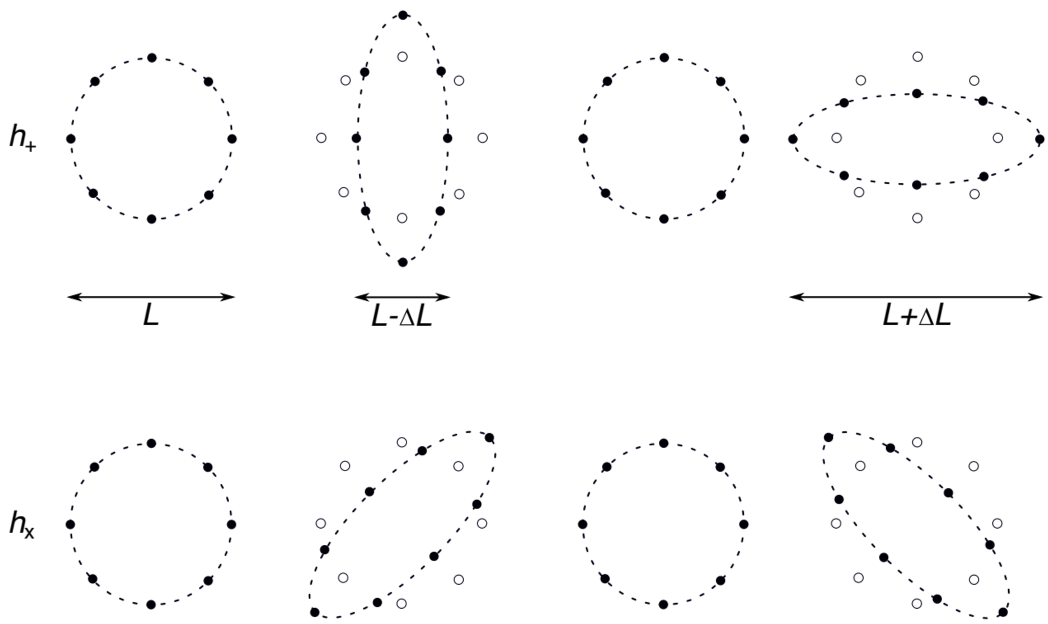

Figure 1.1: A gravitational wave's effect on a ring of particles showing the change in separation between the particles for the two possible GW polarisations, h_+ and $h_\times$ [150]. The GW is travelling in a direction normal to the page.

Figure 1.1 shows that the strength or strain, h, of the wave causes a change in separation, ΔL, between particles of a known distance, L. We can therefore express the strain of a GW signal as

$$h = \frac{2\,\Delta L}{L}. \tag{1.1}$$

Measurements of this ΔL are what allow the detection of gravitational waves [136].

Gravitational waves originate from a wide variety of sources. These waves will be produced by terrestrial objects such as cars or tennis balls, but due to the rigidity of spacetime masses as small as these will produce only miniscule gravitational disturbances which no human-built machine could hope to detect. In order to look for high-mass wave sources we must look beyond earth. Gravitational wave researchers have classified four different categories of gravitational radiation: continuous, burst, stochastic, and compact binary inspiral. These are differentiated based on their origins and signal properties.

1.1.1 Continuous Waves

Continuous gravitational waves originate from a rapidly rotating source which does not display spherical symmetry around its rotation axis, such as a spinning neutron star with surface irregularities or a non-inspiralling compact-object binary system [87]. If the source is rotating at a constant rate the gravitational wave emitted will be both continuous and sinusoidal over a short time period, and if observed over a long time period the rotating object will lose energy through emission of gravitational and electromagnetic energy such that rotation slows, changing the observed GW frequency [126]. Additionally, effective frequency changes occur from the relative motion of the Earth and source object. Radio and X-Ray observations of pulsars place the expected gravitational wave frequency between 1 – 10 Hz for the majority of known pulsar sources, which lies outside the sensitivity band of current GW detectors [183]. These type of signals are weak in relation to other gravitational wave sources and have not yet been detected, but the likelihood of a continuous wave detection increases with the use of third generation detectors [87, 183].

1.1.2 Burst Sources

A burst gravitational wave is an umbrella term for a gravitational wave coming from an unidentified source, with signals expected to be short in duration with an ill-defined waveform. Astronomical events which may produce a burst signal include core-collapse supernovae or gamma-ray bursts [96]. Searches for burst signals are performed by examining detections of excess energy which are coincident between multiple detectors and have a similar waveform [16]. Gravitational wave observations of burst signals may coincide with multi-messenger observations such as electromagnetic waves or neutrinos which would further help to verify such a detection [96].

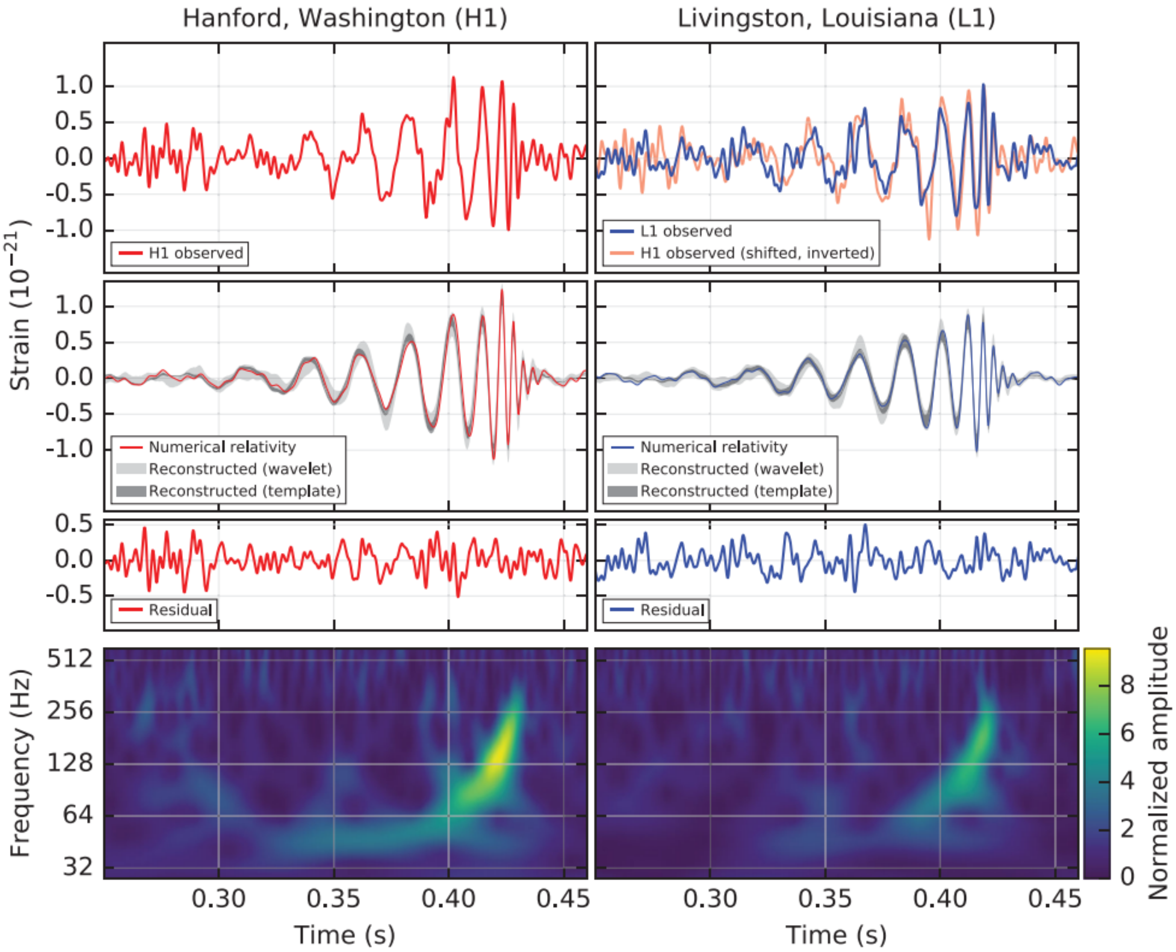

Figure 1.2: Gravitational wave event GW150914 of a binary black hole merger captured by the aLIGO detectors. Shown are the strain signal, waveforms, residual noise, and a time-frequency representation of the strain data [2].

1.1.3 Stochastic Waves

Stochastic waves are the gravitational equivalent of the cosmic microwave background radiation: they originate from many different, random sources which combine to create a continuous and isotropic gravitational background [39, 51]. One predicted source of the stochastic background are gravitational waves emitted from the Big Bang. If a GW signal from the Big Bang can be observed this would give valuable insight into the formation of the very early Universe some 300,000 years before even the formation of the cosmic microwave background [51].

1.1.4 Compact-Binary Coalescence

Binary systems composed of compact objects (black holes or neutron stars) create chirped gravitational waves as they orbit one another. If this orbit is not stable the compact objects will inspiral, orbiting closer and closer until they merge into one larger compact object. This process of inspiral, merger and subsequent ringdown of the new object creates a distinctly shaped waveform of the type shown in Figure 1.2. The high masses of the compact objects and the distinct shape of the gravitational waveform make this type of wave stronger and more identifiable than the other gravitational wave classes, and as such it is the only category of wave which humans have successfully detected [2]. Figure 1.2 shows the signal detected from the first observed gravitational wave, GW150914. In this observation, data is observed by two detectors (H1 and L1) to reduce the chance of a false detection. The observed signals are compared to a bank of predicted waveforms and if there is a sufficient match at both detector sites a GW detection can be claimed to be made with sufficient statistical certainty.

1.1.4.1 Astrophysical Results from Current Observations

Over 200 detections have been made of compact-binary inspiral events since the first detection in 2015 [187, 189]. These detections have began a new era in astronomy, opening the door to previously unobservable signals and having a huge impact on our understanding of the Universe. The detection of GW150914 provided the first direct observation of a black hole, presenting confirmation of Einstein's theory of GR [2]. The observations of binary black hole systems allow black hole populations to be mapped out to cosmological distances. The mass and spin of these populations can be compared to redshift and metallicity, producing clues as to the origin and evolution of black holes and binary black hole systems [131].

The first observation of a binary neutron star inspiral (GW170817) had a profound impact on scientific output [45]. This event coincided with detection of a gamma-ray burst and therefore was the first 'multi-messenger' signal seen in the GW community. This coincident detection verified that light and gravity have the same speed to 15 decimal points, disproving or severely constraining a number of modified gravity theories [177]. This observation also revealed that the origin of many of the heaviest elements lie in neutron star collisions, and provided the first measurement of the tidal deformation which each star induces on its companion. The deformation is directly dependent on the neutron star equation of state, therefore giving indications towards the nature of these dense objects. Additionally, GW170817 enabled the first estimate of the Hubble constant using GW data. The absolute distance to the source's host galaxy can be determined from the GW signal alone, giving an independent distance estimate to the commonly-used 'cosmic distance ladder' of electromagnetic sources. The electromagnetic observations from the multi-messenger detection provided information on the host galaxy's recessional velocity. These results are combined to output an entirely independent calculation of the Hubble constant, which describes local expansion rate of the Universe and is of fundamental importance to cosmology [44, 167].

1.2 Interferometry for GW Detection

1.2.1 Resonant Bar Detectors

The first gravitational wave detector was built at the University of Maryland by Josef Weber. The 'Weber bar' was a large aluminium cylinder thought to resonate at a characteristic frequency when a gravitational wave passed through, acting as a GW antenna. This resonance would be detected via a grid of sensitive piezoelectric sensors mounted to the cylinder's surface [132]. To negate false positive detections from local noise sources two detectors were used, separated by a significant geographical distance, with a resonance coincident at both detectors resulting in a positive detection. The bar dimensions of the bar were chosen to coincide with resonance frequencies expected during supernova collapse [206].

With his device Weber reported many GW detections during 1967-1973, however the GW strength implied by his detections was higher than theoretically predicted [12, 207]. Future repeats of this experiment occurred between the 1960s – 1980s, with several groups improving sensitivities through cryogenic cooling [18, 141]. None of these experiments yielded a positive detection result and Weber's publications are now mainly discredited, with the Weber bar now thought to lack the sensitivity required for most GW sources [127]. Despite this, Weber's work was important in starting a new interest in GW detection: Nobel laureate Kip Thorne referred to him as 'the founding father of this field' [6]. Resonant bar detectors are still in existence today [12], however have largely been superseded by detection techniques which are not limited to narrow-band sensitivity.

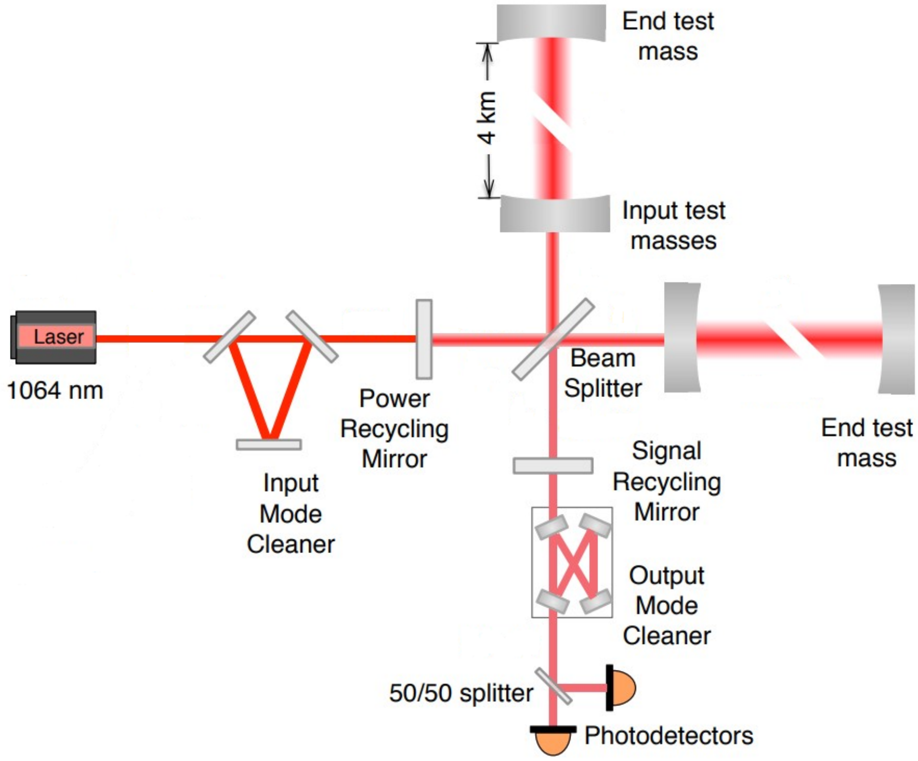

Figure 1.3: Diagram of the layout of an advanced LIGO-type interferometer showing the Michelson interferometer with 4km Fabry-Perot cavities, the power recycling mirror, the signal recycling mirror, and the mode cleaners [139].

1.2.2 Introduction to Interferometry

The interferometer was first developed as a tool for GW detection in the 1970s, with early prototypes being produced in MIT, Glasgow, and Garching by 1975 [128]. This design of detector would go on to make the first detection of gravitational waves in 2015 at the advanced LIGO facilities. Their basic design is based on the Michelson interferometer: laser light is incident upon a beam splitter and sent through two perpendicular arms of equal length. At the end of each arm is a free-hanging mirror, which reflects the laser light back along the arm. The light from the two paths then recombine at the beam splitter and the signal sensed by a photodetector. If no gravitational wave is present the two signals recombine with destructive interference (in reality the operating point of the interferometer is slightly detuned from this dark fringe to allow a small amount of carrier light through, giving the optimum sensitivity to the signal) and no output is seen from the photodetector; however the presence of a gravitational wave will produce phase shifts in the light travelling along the interferometers arms meaning the signal at the photodetector will no longer have complete destructive interference and an output can be measured. The strength of the GW will be proportional to the phase shift of the light and therefore also proportional to the output at the photodetector, allowing the amplitude of the waves to be measured against time [64].

1.2.2.1 Fabry-Perot Michelson Interferometer

Sensitivity to a GW signal is maximised for the Michelson interferometer when the gravitational wave period is twice the detector's light travel time, which gives maximum mirror displacement at the moment when the light is incident on the free-hanging mirror. For a gravitational wave signal from a typical black hole binary inspiral of 100 Hz frequency which is orientated perpendicular to the detector arms, the optimum interferometer arm length is therefore

$$L = \frac{5 \times 10^{-3}\, c}{2} = 750\,\text{km}. \tag{1.2}$$

This 750 km arm length is not practical for Earth-based detectors. Instead, GW interferometers employ Fabry-Perot cavities to increase light storage time within the arms: after the beam splitter in both arms a partially transmissive mirror (input test mass, or ITM) is added to form an optical cavity between this mirror and the end mirror (ETM). The cavity length is set as an integer multiple of $\frac{1}{2}\times$ the laser wavelength to provide constructive interference upon reflection of the laser light, therefore giving resonance of this light within the cavity. The cavity resonance increases the optical path length of the cavity while still allowing a more compact detector size, increasing the interaction time of a GW signal with the light field [64].

1.2.2.2 Dual-Recycled Fabry-Perot Michelson Interferometer

Additional mirrors can be added to further increase performance through signal and power recycling. If the detector is operated near the dark fringe at the photodetector, then most light will leave the interferometer towards the laser. The power recycling mirror is positioned between the beam splitter and laser ensuring that this light is not lost and is instead recycled back into the cavity, increasing the power. The signal recycling mirror is positioned between the beam splitter and photodetector. If a GW signal is present light will pass through the beam splitter and hit the signal recycling mirror - the additional cavity formed by the signal recycling mirror allows the signal to build up and increases the interaction of the signal with the GW [64].

The advanced LIGO detectors used the dual-recycled Fabry-Perot design to make the first observations of gravitational waves in September 2015, with the advanced upgrade being part of the 2nd generation of GW interferometers and shown in Figure 1.3. These detectors use additional technologies which build upon the dual-recycled Fabry-Perot design. The first of these is the input mode cleaner, formed from a triangular cavity placed before the power recycling mirror. This cavity cleans the spatial profile of the laser before the light enters the arm cavities and also acts to suppress input beam jitter and stabilise the laser frequency [139]. Between the signal recycling mirror and photodetector a bowtie-configuration of mirrors forms the output mode cleaner. This cavity filters out unwanted optical components from the signal, such as radio-frequency sidebands and higher-order optical modes, before detection at the photodetectors. The aLIGO detectors also utilise squeezed vacuum states which can improve sensitivity by as much as 35% in quantum noise limited frequency spans [195]. This technique uses the quantum entanglement of photons to decrease uncertainty in the frequency of the laser light while increasing the uncertainty around the laser power, or vice versa [72].

These interferometers are complex and precise instruments - to be able to sense a GW the detectors must be able to measure a change in arm length of 1/1000th the width of a proton [50]. To be able to achieve this extreme level of sensitivity the noise sources in these detectors must be well understood.

1.2.3 Noise Sources in Interferometric Detectors

This section will outline fundamental noise sources currently placing a limit on the sensitivities of 2nd generation ground-based detectors. The contribution of these noise sources in the advanced LIGO detectors are shown in Figure 1.4. Plotted in strain against frequency, the total of the noise sources creates the so-called 'sensitivity curve' of this detector which can be seen again in Figure 1.8 in comparison to other ground-based detectors. This noise curve shows the fundamental limits for the detector and provides a baseline noise requirement for technical noise sources, which are not covered in this work.

1.2.3.1 Seismic & Newtonian Noise

Seismic noise is dominant in the low frequency range of $1 - 10\,\mathrm{Hz}$. This noise is caused by ground motion which leads to movement of the test masses, with sources such as Earth tremors, anthropogenic sources, or ocean waves [26]. At a seismically quiet geographical location these sources may contribute a noise level of $10^{-7} \times f^{-2}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$, for frequency f [109]. Pendulum mirror suspensions are therefore utilised to isolate the test-masses from their surroundings by around 10 orders of magnitude [17].

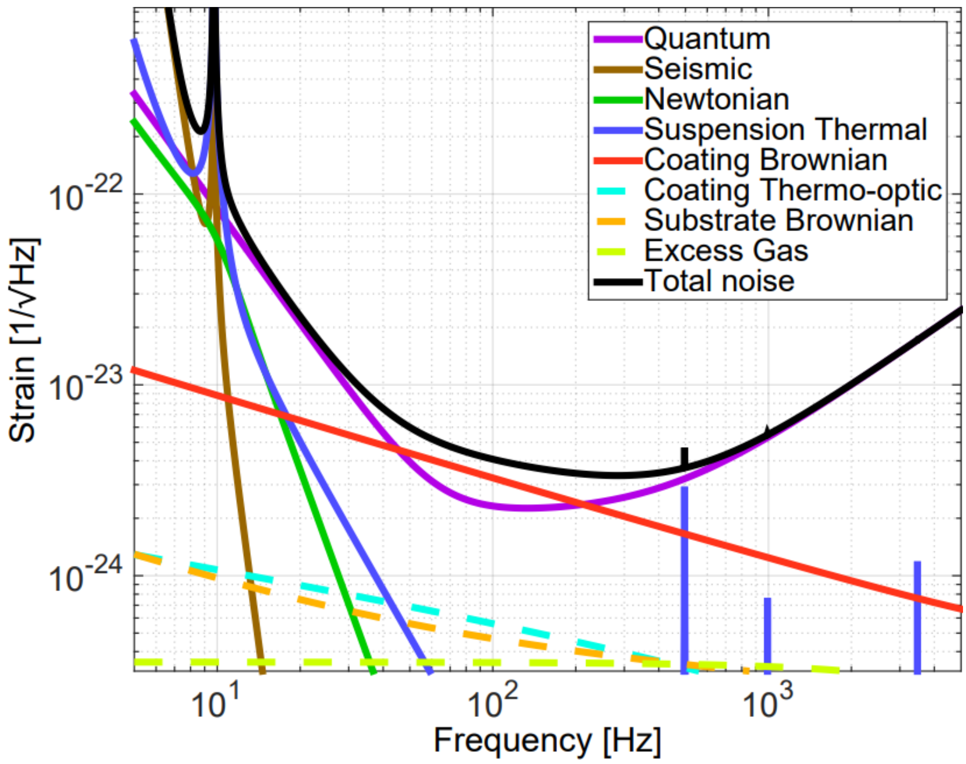

Figure 1.4: Advanced LIGO sensitivity curve showing main noise sources of seismic and Newtonian noise (brown and green), quantum noise (purple), and thermal noise (blue, red, dashed cyan, dashed orange) [22]. Thermo-optic noise is the combination of thermoelastic and thermorefractive noise. Excess gas noise (dashed yellow) arises from the interaction of residual gas particles with the test mass inside the vacuum tank and is an infrastructure limitation rather than a fundamental noise source [33].

Newtonian noise arises from fluctuating gravitational forces and is responsible for the lower frequency sensitivity limit [162]. Newtonian noise has a wide range of sources: nearby cars, people, and passing planes will all produce a gravitational force which acts on the test masses.

Alongside the isolation from pendulum mirror suspensions the detectors can be built underground to increase the distance from sources on the Earth's surface, therefore reducing local seismic and Newtonian noise. Additionally, the locality of these noise sources means the chances of seismic activity triggering an event in multiple detectors can be reduced by using two detectors placed geographically far apart from each other – the two advanced LIGO detectors are separated by over 3000 km. In order to detect signals of less than 10 Hz space-based detectors must be used [162].

1.2.3.2 Quantum Noise

There are two sources of quantum noise: photon shot noise and radiation pressure. Photon shot noise provides the upper frequency sensitivity limit in modern detectors, being the dominant noise source above a few hundred hertz. Photon shot noise occurs due to the discrete nature of photons: the number of photons arriving at the detector in a given time period will not be regularly spaced, but will arrive in a Poisson distribution given by the uncertainty in the phase noise of the laser beam. The effect of photon shot noise can be reduced by an increase in laser power [64].

Radiation pressure noise is caused by the transfer of momentum from photons to the suspended mirrors, causing slight movement of the mirror along the cavity axis. At the quantum limit this effect arises from the uncertainty in the amplitude of the laser beam. An increase in laser power reduces photon shot noise but increases radiation pressure noise, so laser power needs to be set as a compromise between these two effects [64, 99].

Since 2019 the advanced LIGO detectors have utilised light squeezing as a way of decreasing uncertainty in either the amplitude or phase domain at the cost of increasing uncertainty in the other domain [72]. The latest installed technology is frequency-dependent squeezing, first utilised in advanced LIGO in 2023. An additional optical cavity is used to transition the laser light from an amplitude-squeezed state to a phase-squeezed state above a chosen frequency. A reduction in frequency uncertainty is used at high frequencies to reduce the shot-noise limit, and a reduction in power uncertainty is utilised at low frequencies to reduce the effects of radiation pressure noise. This gives a total reduction of quantum noise in the detectors by decreasing noise in the experimental sensitivity band and increasing the noise outside the region of interest [49].

1.2.3.3 Thermal Noise

Thermal noise is dominant in the mid-frequency range of $40 - 400\,\mathrm{Hz}$ with contributions from thermoelastic, thermorefractive, and Brownian thermal noise [99]. The coatings, substrates, and suspension systems all contribute thermal noise as seen in Figure 1.4.

Thermorefractive noise is caused by the temperature-dependence of a material's refractive index. Thermal fluctuations in the optical substrates therefore leads to a change in refractive index, and as the beam passes through an optic (or into a coating) this change in refractive index introduces phase noise to the beam [73, 198]

Brownian noise arises due to the fluctuation-dissipation theorem which states that a measurable parameter of a dissipative system (here, the mechanical impedance of suspension materials, contributed to via mechanical losses) is subject to random thermal fluctuations known as Brownian thermal noise [120], discussed in Chapter 2.

The materials used in the mirror suspensions are chosen to have low thermal expansion coefficients however any thermal fluctuations can produce a slight expansion or contraction across the material surface - this is thermoelastic noise. For substrate or coating thermoelastic noise this is the expansion and contraction over the mirror or coating surface which impacts the phase stability of the laser beam. For a mirror suspension the thermoelastic noise contribution arises from expansion and contraction in the suspending fibre which moves the mirror position a small amount.

Suspension thermal noise from Brownian thermal noise and thermoelastic noise is described in depth in Chapter 2. Coating thermal noise is also important to consider here as the limiting noise source of current detector sensitivity in the mid-frequency band, as shown in Figure 1.4 [52]. There is currently a focused research effort in the GW community around the measurement and reduction of coating thermal noise. The University of Glasgow aims to take direct measurements of coating thermal noise at temperatures being considered for the operation of third generation detectors (see Sections 1.4.4 and 2.4). The test masses of the interferometer arms are surfaced with a highly-reflective coating to increase the power circulating in the cavity and reduce laser heating of the masses. These coatings must therefore display high reflectivity, low absorption, and low scatter. Alternating layers of dielectrics are typically used to manufacture coatings due to their durability, low absorption, and ability to be deposited over a large surface area via ion beam-spluttering [138]. The coatings in use at the advanced LIGO interferometer are formed of layers of titania-doped-tantala and amorphous silica [52]. The most significant source of coating thermal noise in current detectors is the Brownian thermal motion caused by friction in the mirror reflective coatings, with the friction arising from point defects, dislocations and grain boundaries within the material [52, 138]. The dominant contributor of Brownian dissipation in mirror coatings is the internal friction of tantala [100].

(a) Aerial view of LIGO Hanford site [129].

(b) Aerial view of the Virgo interferometer [204].

(c) Concept image of the KAGRA detector [42].

Figure 1.5: Aerial images and concept art showing the LIGO Hanford, Virgo, and KAGRA gravitational wave detector sites.

1.3 Current Interferometer Detectors

The first generation of gravitational wave detectors were the initial LIGO, Virgo, and GEO600 detectors, operational in the early 2000s. Despite being a groundbreaking proof-of-concept these sites did not make any detections of gravitational waves, and went through significant upgrades to increase sensitivity in the 2010s. Post-upgrade the LIGO and Virgo detectors reached ‘advanced’ status and began the second generation (2G) of gravitational wave detectors. This grouping is made up of advanced LIGO in the USA, advanced Virgo in Italy, KAGRA in Japan, GEO High Frequency in Germany, and has LIGO-India predicted for commission in the 2030s. A comparison of the 2G LIGO-Virgo-KAGRA (LVK) detectors is shown in Figure 1.5, with a detailed description of the suspension payload for each of the LVK detectors is given in Section 2.1.1. This is the generation which was sensitive enough to make the first GW detection, and has since observed many binary inspiral events between black holes and neutron stars.

The Gravitational European Observatory 600 (GEO600) operates from just outside Hannover as a joint collaboration between Germany and the UK. With 600 m length cavity arms this detector has a lower sensitivity than the other 2nd generation detectors with infrastructure limitations preventing the detector from reaching detection sensitivity. Despite this, it has been invaluable as a test facility and was the first demonstration of many of the technologies used for the first gravitational wave detection [92, 130].

The Laser Interferometer Gravitational-Wave Observatories (LIGO) are a GW detector pair situated in the United States. Each detector is an L-shaped interferometer with 4 km length arms, operational at room-temperature with fused-silica cavity optics. The detectors are situated far apart with one site in Hanford, Washington and the second in

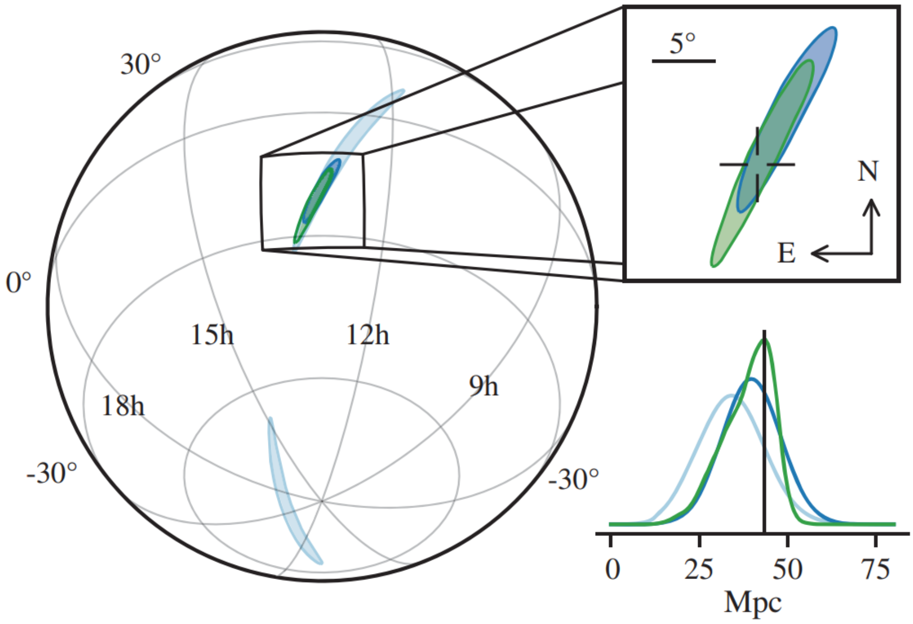

Figure 1.6: Sky localisation for source GW170817 from a rapid localisation algorithm using data from the two aLIGO detectors (light blue), then the aLIGO detectors and Virgo (dark blue). The green contours show the localisation of a higher-latency algorithm which later narrowed the position further [45].

Livingston, Louisiana [3]. This aids triangulation of GW source location and ensures that each detector is subject to independent seismic noise. The LIGO detectors began their GW search in 2002 but did not make any detections until after their 2008-2015 upgrade to advanced LIGO (aLIGO) [47]. The aLIGO detectors made the historic first observation of GW waves in 2015 by detection of a binary black-hole merger [2].

The Virgo gravitational wave detector situated near Pisa, Italy, is a room-temperature L-shaped interferometer detector, which has arm lengths of 3 km and fused-silica cavity optics [37]. This detector began operation in 2007 but, similarly to LIGO, did not make any detections until the upgrade to advanced Virgo (adV) status in 2017 at which point the detector made joint observations with LIGO as part of the LIGO-Virgo Collaboration (LVC) [5]. With three detectors widely spaced across the globe, the LVC could triangulate GW source locations to a much greater accuracy than with the two LIGO detectors alone. Figure 1.6 shows the improvement of sky localisation for binary neutron star merger GW170817 where the light blue bands display the localisation from the two aLIGO detectors alone, and the dark blue from the same localisation algorithm including the Virgo interferometer and placing much greater constraints on sky location [45].

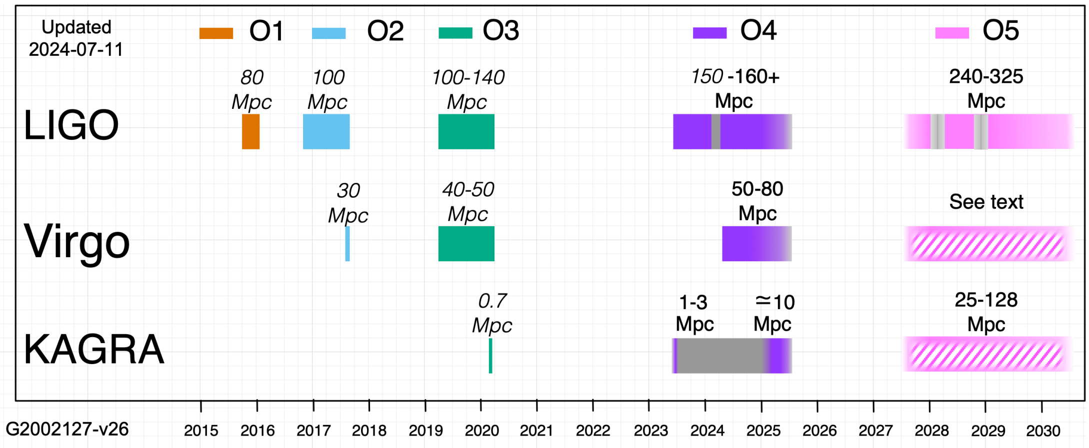

Figure 1.7: The planned observing schedule for the LVK detectors from 2015 - 2030. At the time of writing the O4 observing run is ongoing, and the Virgo Collaboration is reassessing their entry date into O5 [190]. The grey areas represent a break in observing for maintenance or commissioning work. For each observing run, the binary neutron star range is given for single-detector SNR threshold of 8 in megaparsecs (Mpc).

The fifth detector to join operations was the Japanese Kamioke Gravitational-Wave (KAGRA) detector. KAGRA is also an L-shaped interferometer with 3 km length arm cavities, but has several unique properties compared to the LVC detectors. KAGRA is the first detector to operate at cryogenic temperatures, to use sapphire cavity optics, and the first detector situated underground to reduce seismic and Newtonian noise sources [42]. These advancements mean it is considered by some to lie between the 2nd and 3rd generations of GW detectors. KAGRA's first observation run began in 2020 where it joined the LIGO and Virgo detectors to form the LIGO-Virgo-KAGRA Collaboration (LVK), although KAGRA did not have a sufficient sensitivity to make observations in this run.

1.3.1 Future Upgrades to Current Detectors

The observational periods of the LVK detectors are split into 'observing runs', with co-ordinated down time between each run for commissioning or construction. Small modifications can be made to detectors between observing runs to include new technologies and improve detector sensitivity without a complete re-build of the facilities. The observing schedule for O1 to O5 is shown in Figure 1.7. The first observing run, O1, was only four months long with just the two aLIGO detectors online, but resulted in the first three detections of binary black hole mergers. The second observing run, O2, lasted almost ten months for aLIGO with advanced Virgo joining near the end of the run. Eight detections were made over the course of this run, including the first three-detector observation of the GW170817 binary neutron star merger. The third observing run was the first to include the four detectors of the LVK. Improvements in detector performance resulted in 79 detections in almost the same time period as O2. Observing run 4 is ongoing at the time of writing with planned completion in June 2025. The final observing period of the current detectors is expected to run between 2027 – 2030 whereafter the detectors will undergo a further stage of upgrades, with proposed plans described below [190, 188].

Works have already began for the next upgrade to the aLIGO detectors, named LIGO A+. These works will use the existing site infrastructure while improving sensitivity to binary neutron star inspiral range. This upgrade will include improve frequency dependent squeezing using a 300 m filter cavity, a balanced homodyne readout scheme, and lower loss test mass coatings [48]. The implementation of these upgrades are to begin in O4, with the completed A+ expected to be ready for the beginning of O5.

LIGO A# is a proposed 2030s upgrade to the aLIGO detectors. This room-temperature upgrade plans to incorporate heavy fused-silica optics, increasing the cavity test masses to 100 kg, alongside new active isolation technologies. The stress on the fused-silica suspension fibres will be increased, as well as an increase in cavity laser power and the inclusion of new optical coatings with a lower thermal noise. This improves the sensitivity of the detector using existing infrastructure and well-known technologies [96].

LIGO Voyager is a proposal to reach the infrastructure sensitivity limit at the current LIGO detector sites. LIGO Voyager plans to use the existing sites with 4 km arm lengths and upgrade these through the installation of 200 kg silicon test masses with amorphous silicon coatings, and through the cooling of the cavity optics to a cryogenic temperature of 123 K. For compatibility with silicon optics a laser wavelength of 2000 nm is proposed. Newtonian noise would be reduced using a sensing array of seismometers and active noise reduction techniques [8].

A concept study exists for a proposed upgrade to the adV interferometer, named Virgo nEXT. This upgrade is proposed to begin installation in 2029 after the end of the planned final observing run of the current 2G detectors. Frequency-dependent squeezing and an increase in laser power aims to improve quantum-noise limited sensitivity. The use of new mirror coatings and an increase in test mass size reduces some thermal noise contributions. Improving the site vacuum system reduces excess gas noise. Finally there is planned sensing and filtering of Newtonian noise to improve sensitivity between 10 – 20 Hz [53].

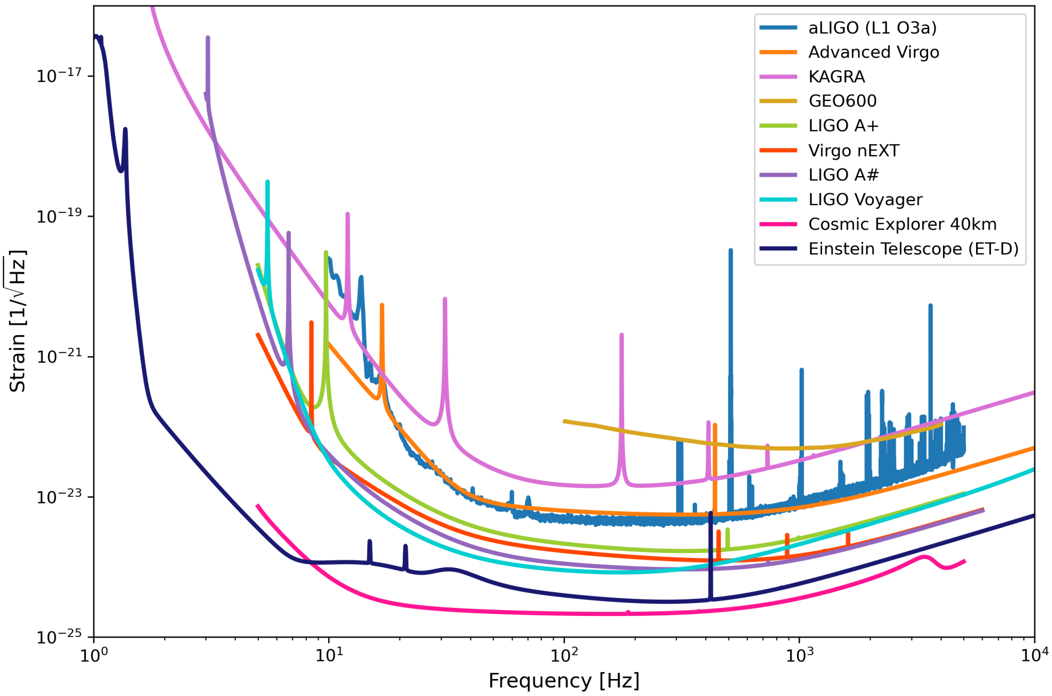

Figure 1.8: Strain sensitivities of current and proposed ground-based gravitational wave detectors. ET-D includes the combined sensitivities of ET-LF and ET-HF.

There is no published plan for 2030s upgrades at the KAGRA detector with KAGRA's 10-year plan currently under review. KAGRA aims to use their existing knowledge of cryogenic interferometry and crystalline suspensions to aid research and development for third generation detectors [114].

1.4 Third Generation Interferometer Detectors

1.4.1 Planned Detectors

The third generation (3G) plans to build on the success of the current GW detectors to further increase both sensitivity and bandwidth. Einstein Telescope (ET) and Cosmic Explorer (CE) are two third generation detectors which will each use new technologies to gain a 10× improvement in sensitivity over aLIGO across all frequencies, as seen in Figure 1.8.

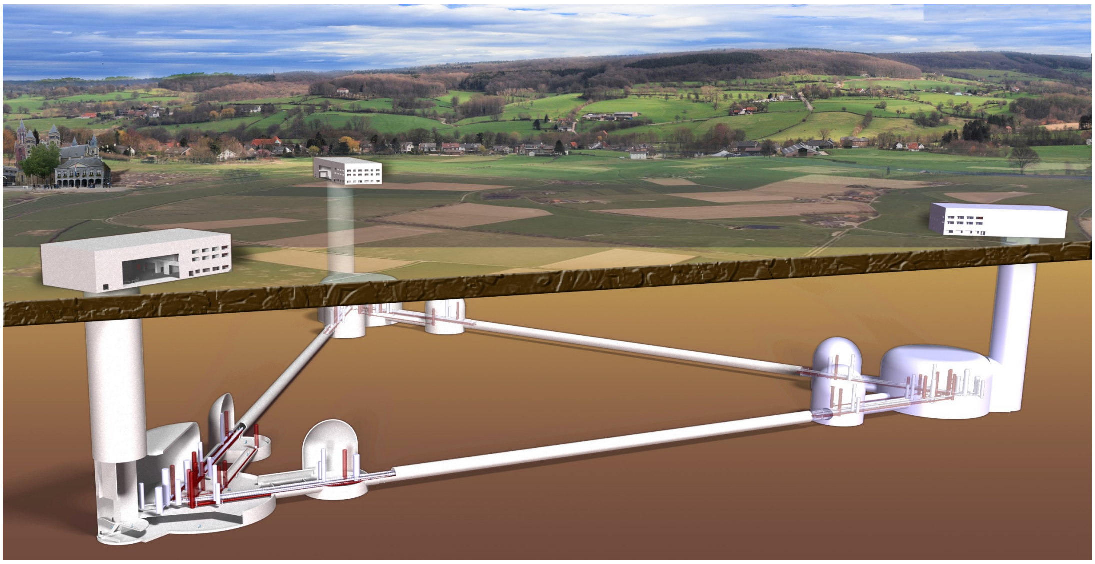

Figure 1.9: An artist's impression of the Einstein Telescope layout, showing the 10 km arms in a triangular configuration with three end stations [40].

The Einstein Telescope is a planned European detector aiming to begin construction in the 2030s. In place of the L-configuration of other GW interferometers, Einstein Telescope will contain three nested detectors in a triangle configuration, with arms lengths of 10 km. An artistic representation is shown in Figure 1.9. Each detector will contain two interferometers, with one interferometer optimised for detection of low-frequency GWs (ET-LF) and the other optimised for detection of high-frequency GWs (ET-HF). The low-frequency detector is planned for cryogenic operation with long silicon suspensions, whereas the high-frequency detector will utilise room temperature heavy fused-silica test masses with a high cavity power [157]. The entire ET facility will be built underground to vastly reduce the seismic noise background and allow detection of lower frequency GWs [182].

Cosmic Explorer is a next generation design concept by the United States. CE will have two detectors which will retain the L-shape and basic design of aLIGO and LIGO A# but push the arm lengths up to 40 km (or 20 km at the second site) [156]. This increases the amplitude of detected signals without much increase in detector noise [158]. Future iterations of CE would look to adopt cryogenic technology to further increase sensitivity [71].

The Laser Interferometer Space Antenna (LISA) is a planned interferometric GW detector operated by the European Space Agency. Unlike the other 3G detectors LISA will be space-based, placed in an Earth-trailing orbit with arm lengths of 2.6 million kilometres. The LISA detector would allow measurement of low-frequency GW signals which could not be detected by ground-based detectors, such as signals from binary compact objects

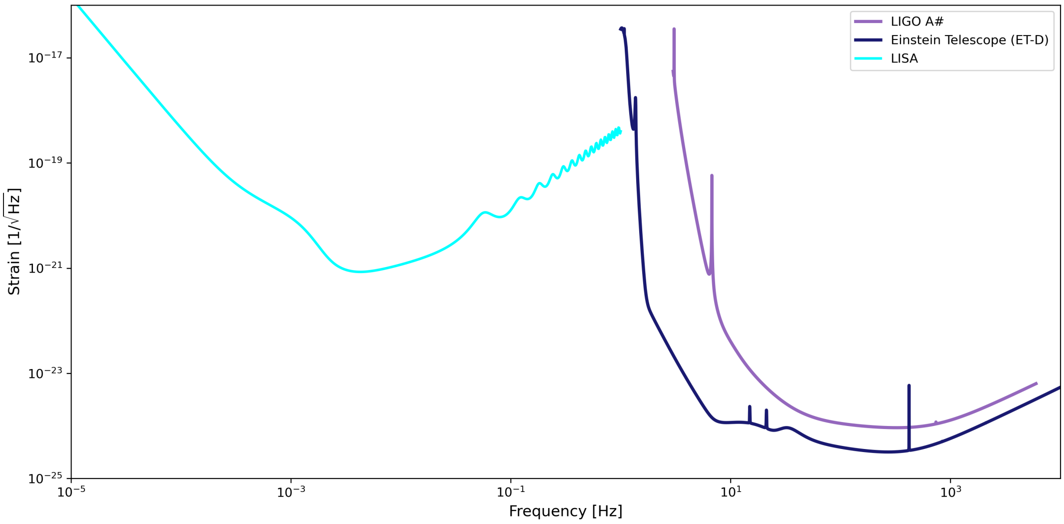

Figure 1.10: Strain sensitivity of the LISA detector in comparison to the ET and LIGO A# curves from Figure 1.8, which may be in operation on a similar timescale.

within the Milky Way, supermassive binary black holes from distant galaxies, and compact binary inspirals with extreme mass-ratios [11]. The strain sensitivity curve for the LISA detector is shown in Figure 1.10 compared to two ground-based detectors which may be in operation around a similar time-frame. LISA has a planned launch date of 2035 [180].

1.4.2 Cryogenic Technology

One of the main technological advances in GW detection will be the cooling of optics down to cryogenic temperatures. First demonstrated in Japan's KAGRA facility, cryogenic cooling has already demonstrated a significant decrease in thermal noise in the mid-low frequency band by reducing both thermoelastic and Brownian noise [38]. The fused-silica optics in use at room temperature facilities are unsuitable for use at low temperatures due to silica's broad dissipation peak at 40 - 60 K [138]. Investigation of alternative materials is an active area of research and is vital to the success of the 3rd generation facilities. Both sapphire and silicon are being explored for use in the test masses, with mono-crystalline sapphire in use at KAGRA and silicon likely to be used for ET. With the main optic substrate being changed, this has a knock-on effect to other interferometer components such as laser wavelength, test mass coatings, and suspension bonding techniques. Crystalline suspensions and suspension bonding techniques are discussed in Chapter 2. The study of a 2000 nm laser (a wavelength suitable for use with silicon optics) is covered in Chapter 5.

1.4.3 Next Generation Interferometer Prototypes

Prototype interferometer facilities are crucial for the research and development necessary to advance current detectors and start the build of the next generation. These facilities allow proof-of-concept to be shown for new ideas, and allow new experimental procedures to be safely tested without interference to detector operation or observing periods, all with a quick turnaround time that a large-scale detector could not accommodate.

The Albert Einstein Institute houses a 10 m L-shaped interferometer prototype in Hannover, Germany. This facility aims to be quantum-noise limited across 50 – 500 Hz to investigate potential methods of overcoming the standard quantum limit, looking into both radiation pressure noise and shot noise [91].

A 40 m L-shaped interferometer housed at the California Institute of Technology is used as the control prototype for the aLIGO detectors. In the past it has been used to test optical configurations and controls, and develop interferometer diagnostic techniques [9]. The next stage for the 40 m prototype is the Mariner project which aims to prototype technologies required for the LIGO Voyager upgrade such as cryogenic silicon test masses and a laser of 2000 nm wavelength [10, 81].

The University of Western Australia have installed the 'Gingin' interferometer prototype near Perth on Australia's west coast. This L-shaped interferometer operates at room-temperature with 80 m arms and specialises in high-cavity power investigations. Currently it is the only prototype facility with a mirror-mass to laser power ratio comparable to the kilometer-scale GW detectors [208].

ET Pathfinder is a joint research and development facility built in Maastricht as a collaboration between Belgium, Germany, and the Netherlands. This prototype interferometer is L-shaped with 10 m length arms. The facility is built specifically to target demonstration of technologies needed for the planned ET detector such as low-noise cryogenic cooling of test masses to 120 K and 15 K, investigation into silicon as a test mass material, laser stabilisation and components at 1550 nm – 2000 nm wavelengths, and advanced quantum noise reduction techniques [184, 185].

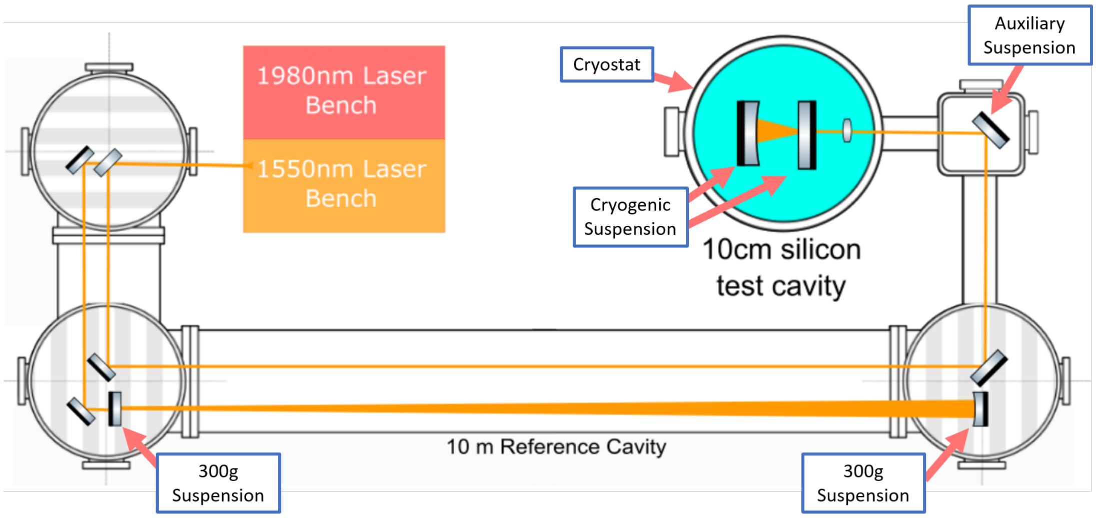

Figure 1.11: GCIF layout showing the 10 m room temperature reference cavity for laser stabilisation and the 10 cm cryogenic test cavity for silicon suspended optics [173]. The input laser optics and 10 m cavity optics sit in vacuum chambers attached to the beam tubes. After the 10 m arm a small square vacuum chamber houses the mirror used to steer the beam into the cryogenic cavity. The 10 cm test cavity will sit inside the cryostat shown in Figure 1.12. The position of the suspensions developed within this work are marked.

1.4.4 Glasgow Cryogenic Interferometer Facility

Investigations of interferometric detection of gravitational waves began early in Glasgow, with Ronald Drever's formation of a gravitational-wave research group in 1970 and the first operations of an interferometer prototype in the 1980s [86, 105]. In the early 2000s the facility was moved into a clean room laboratory with an upgraded layout featuring two 10 m length arms and one 5 m length arm set up in parallel, with nine 1 m-diameter vacuum tanks, and a control and data acquisition system akin to the one used at aLIGO. This facility was in operation between 2003 – 2024. Experiments conducted over this time period included speed meter interferometry [89], an investigation into optical springs to reduce the effect of radiation pressure noise [88, 135, 213], diffractive optics for GW detectors [21], an experiment for the direct measurement of coating thermal noise at room temperature, and an investigation into the control scheme for a dual recycled Michelson interferometer [20, 107, 176].

In 2024 the Glasgow 10 m prototype was upgraded to the 'Glasgow Cryogenic Interferometery Facility', or GCIF: a double cavity cryogenic interferometer prototype. This laboratory will allow the testing of future technologies required for 3rd generation detectors such as fully crystalline suspended optics, materials research, 1550 – 2000 nm laser technology, and cryogenics research. The GCIF will have one 10 m room-temperature reference cavity, leading to a shorter 10 cm cryogenic cavity with suspended silicon optics. The laboratory

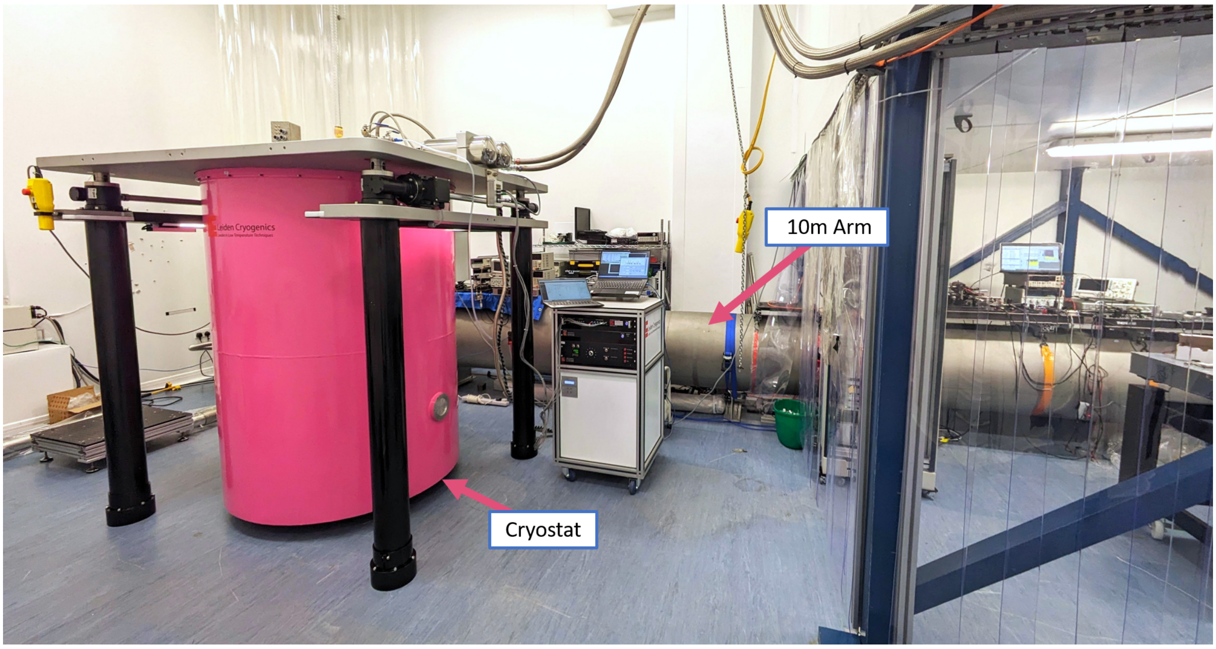

Figure 1.12: Photograph of the newly-installed cryostat in the GCIF. The cryostat, pink, is capable of cooling to 7 K and will house the laboratory's cryogenic test cavity. The 10 m room temperature beam tube can be seen against the back wall of the laboratory. This room-temperature cavity will be used for pre-stabilisation of the laser beam.

layout is shown in Figure 1.11, with the cryostat and 10 m beam tube shown in Figure 1.12. The facility plans to operate at both 1550 nm and 2000 nm, allowing research into the viability of both wavelengths for use at future detectors and for the development of a prototype lock acquisition scheme with silicon test masses. Once completed, this facility plans to operate as a multi-experiment facility with a quick turnaround time. Planned future investigations include direct measurements of cryogenic coating thermal noise, a study of the base noise performance of silicon suspensions, optimisation of low-noise cooling and control at cryogenic temperatures, and the study of ice formation on silicon optics.

1.5 Conclusion

This chapter provides the scientific background necessary to understand the origin of gravitational wave events, the interferometric detection of these events, and the infrastructure landscape of existing and future detectors. Interferometer prototypes are invaluable in facilitating the investigations and proof-of-concept of future 3G technologies. The Glasgow Cryogenic Interferometer Facility is one such prototype which plans to investigate monolithic silicon suspended optics at cryogenic temperatures within a double-cavity interferometer, and the development of laser technologies and noise-reduction schemes at 1550 nm and 2000 nm wavelengths. This thesis describes the groundwork investigations for the 2024 remodelling of the Glasgow facility which includes both modelling and table-top investigations into suspension design and laser optics.

Chapter 2 gives an introduction on suspension dynamics and the calculation of thermal displacement noise which is necessary for a full understanding of suspension modelling and design.

Chapter 3 describes the modelling, characterisation, and modification of two designs of optical suspension for use in the GCIF. These suspensions will be used as cavity optics for the room-temperature 10 m pre-stabilisation cavity, and for beam steering.

Chapter 4 outlines the design concept of a cryogenic, low-noise crystalline suspension for use as the cavity optic for the GCIF 10 cm test cavity. This suspension will feature a monolithic lower stage with silicon test masses and crystalline fibres. This chapter outlines a design which meets the seismic isolation and thermal noise requirements required for a direct measurement of coating thermal noise, and investigates the cooling scheme of such a suspension at both 18 K and 123 K.

Chapter 5 discusses the build and characterisation of 2000 nm external-cavity diode lasers for use as the GCIF input beam and for investigations into 2000 nm technologies. A control system was designed and implemented to demonstrate reduction of intensity and frequency noise through feedback to the diode current.

Chapter 6 provides a discussion of the work contained in this thesis with an emphasis on the final design concepts for the GCIF suspensions, the demonstration of the build and noise reduction in a 2000 nm external-cavity diode laser, and recommendations for further investigations for future stages of the GCIF upgrade.

Chapter 2

SUSPENSION DYNAMICS AND THERMAL DISPLACEMENT NOISE

2.1 An Introduction to Suspension Design

Cavity mirrors in gravitational wave detectors are suspended as pendulums to reduce seismic noise coupling to the mirror through natural ground motion, anthropogenic noise, and mechanical vibrations. Multiple pendulum stages can be stacked to further reduce this coupling [3, 155]. Consider a simple pendulum consisting of a rigid mass at the end of an elastic wire. The transfer function of the amplitude of motion of this pendulum from some ground motion, x_{g}, to mass motion, x_{m}, may be expressed as

$$\frac{x_{\mathrm{m}}}{x_{\mathrm{g}}} = \frac{\omega_0^2}{\sqrt{(\omega_0^2 - \omega^2)^2 + (\gamma\,\omega)^2}} \tag{2.1}$$

with phase, θ, given by

$$\tan(\theta) = \frac{\gamma\omega}{\omega_0^2 - \omega^2} \tag{2.2}$$

[79, 125]. The pendulum resonant frequency is represented as ω_0, with ω the angular frequency of ground motion. γ describes the damping of pendulum resonances – low damping is assumed for the suspensions in this work to decrease off-resonance thermal noise as expressed in Figure 2.1. For ω below the suspension resonance $x_{\mathrm{m}}/x_{\mathrm{g}}$ is unity. There is a resonance peak as ω approaches ω_0 of width dictated by the damping factor. As ω increases above the suspension resonance the ground motion transferring to mass motion decreases with frequency, in the limit of low damping, as

$$\frac{x_{\mathrm{m}}}{x_{\mathrm{g}}} \propto \frac{\omega_0^2}{\omega^2}. \tag{2.3}$$

In the upper stages the masses and suspension wires are typically manufactured from metal for robustness and ease of manufacture, whereas the penultimate mass, fibres, and test mass (TM) of the 'monolithic', lower stage are glass or crystalline to reduce thermal noise around the optic, such as can be seen in Figure 2.2 [155].

Thermal noise arises from random thermal fluctuations inside the suspension materials which cause energy dissipation through vibrations of the material's atomic lattice [31, 58]. The thermal noise of the suspension depends both on the loss of the suspension materials and on the energy distribution in the region - for example, a region of the suspension manufactured from a low loss material but with high energy dissipation may contribute more thermal noise to the system than a region with a high-loss material and minimal energy dissipation. The energy distribution across the suspension chain is investigated through the study of resonant modes of the suspension system. In a resonating system the quality factor, Q, is the ratio between the maximum energy stored in the system and the energy lost in one cycle of oscillation [93]. In a high-Q physical system this energy dissipation is contained in the natural resonant modes of the system whereas in a low-Q system this energy is dissipated off-resonance, with a comparison of the effect on thermal noise depicted in Figure 2.1. In order to reduce the thermal noise of the suspension chain in the experimental sensitivity band it is therefore advantageous to build a high-Q system to contain the thermal energy in resonance, then design the suspension such that these resonance peaks lie outside the chosen sensitivity band, or can be actively damped.

2.1.1 Suspensions of the Current Generation of Detectors

The specifics of the suspension designs used in GW detection vary between detector sites. Common across all designs are the metal upper mass stages, with metal wires attaching to the mass which at their top end attach to triangular blade springs to provide vertical seismic isolation to the suspension chain [155]. See Chapter 4 for a full discussion of cantilever blade springs. Multiple metal mass stages can be used in the suspension chain. Attached below the metal stages are the kg-scale glass or crystalline monolithic suspension masses. The cavity test masses are cylindrical optics which feature a multi-layer optical coating on their front face to maximise the light resonating in the arm cavities [174]. The curved sides of the optics are cut and polished to form flats onto which 'ears' can be bonded to form the connection between test masses and fibres. In the upper suspension stages stainless steel wires form the attachments between masses, whereas fused-silica or crystalline fibres form this attachment in the monolithic stage. A reaction chain or cage is suspended next to the main chain allowing for positional control of each mass.

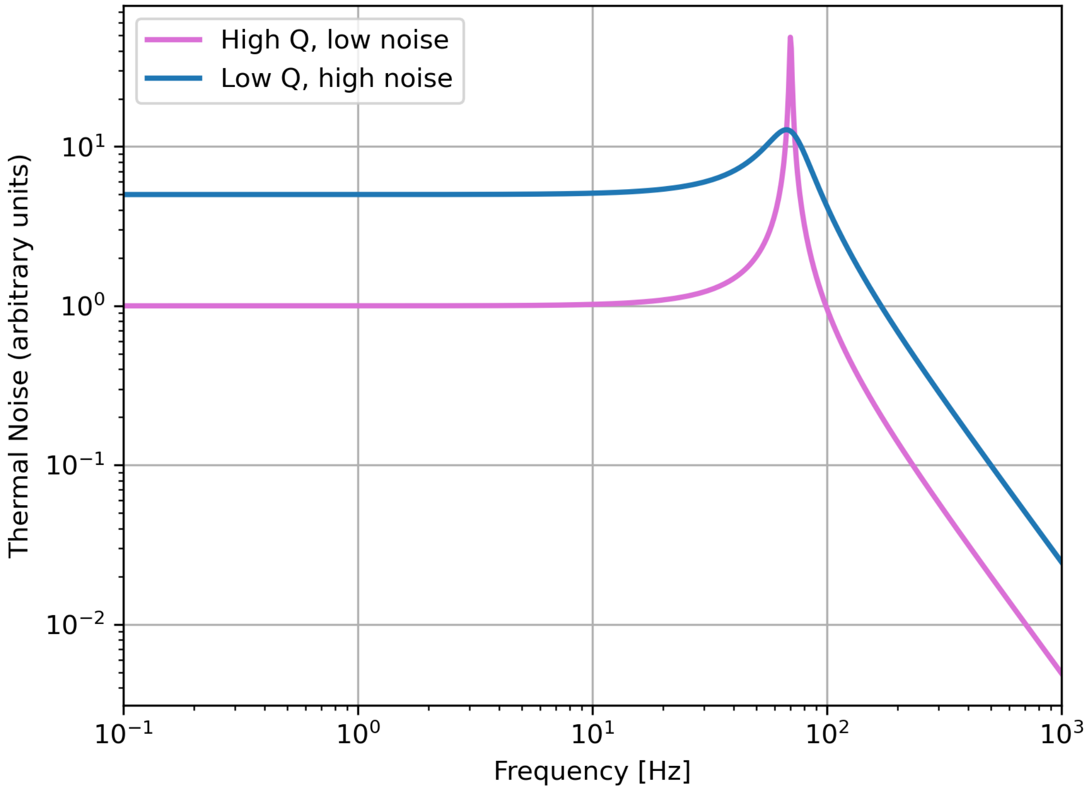

Figure 2.1: Comparison of the thermal noise (in arbitrary units) of a high-Q system (pink) where the thermal energy is contained in resonance, and a low-Q system (blue) which has a higher off-resonance thermal noise.

The actuators are required for damping and operational control of the detector, and the reaction chain arrangement mitigates coupling of seismic and vibrational noise into the mirror positions through the actuators. The suspension designs of the current generation of GW detectors are outlined in the following.

2.1.1.1 Advanced LIGO Suspension

The aLIGO 'quad-suspension' system shown in Figure 2.2 is adapted from the GEO600 triple suspension design, and is a four-stage suspension with two metal upper masses and a fused-silica monolithic stage of two masses [153, 160]. This suspension is capable of isolating seismic motion reaching the test mass by 10 orders of magnitude at 10 Hz [17]. Steel blades springs are used at the suspension mounting point and at the two upper masses to provide three stages of vertical isolation to the chain. The fused-silica ears have a welded horn design and are attached to the masses via hydroxide catalysis bonding (see Section 2.3.1.4), with the fibres welded to the ears as shown in Figure 2.3(a). The

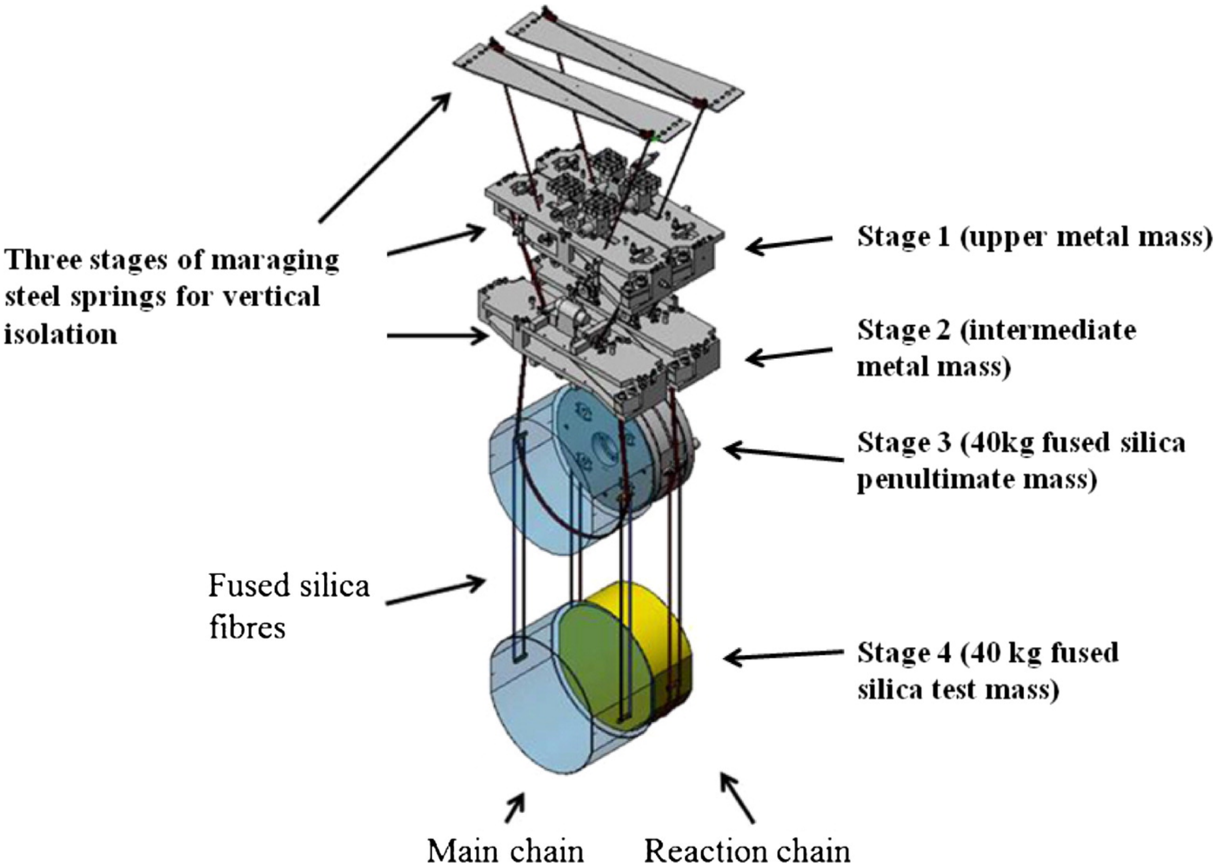

Figure 2.2: Diagram of the advanced LIGO quadruple pendulum suspension showing the main chain and reaction chain. The main chain is comprised of two metal upper masses and the monolithic stage of two 40kg fused-silica masses which are joined by four fused-silica fibres [79].

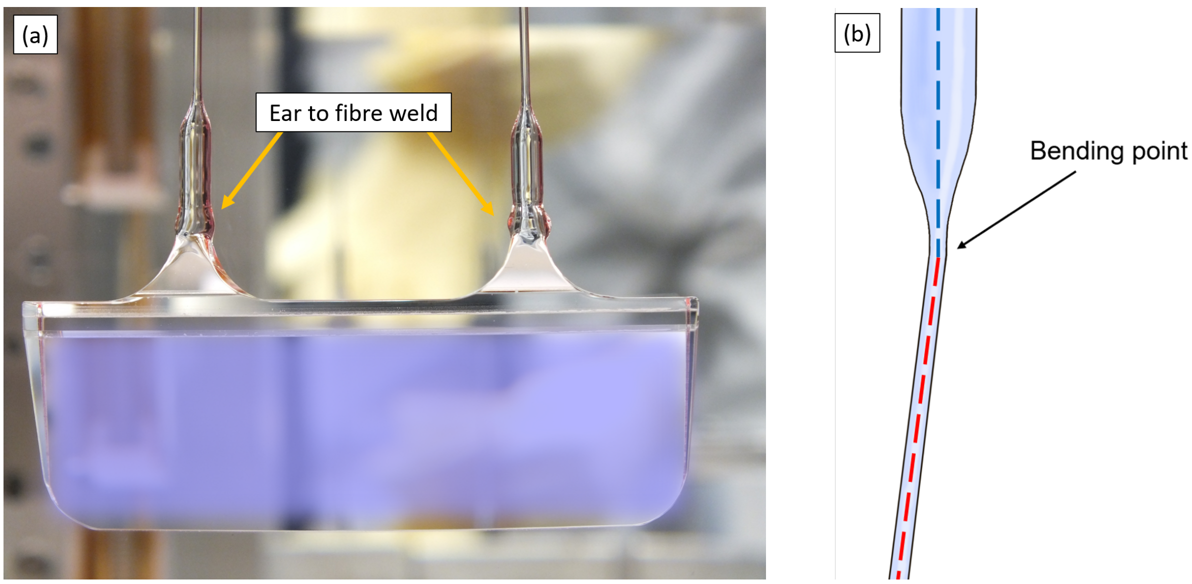

Figure 2.3: (a) Side view of the aLIGO fused-silica ears showing the fused-silica fibres which have been attached via welding, and the HC bonded region highlighted in purple [124]. (b) Illustration of the bending point of a fibre at the intersection of the vertical axis (blue) to the normal to the top surface of the mass (red).

reaction masses are hung as a reaction chain, sitting 'behind' the main chain masses in the cavity. The upper two reaction masses are steel and allow for the adjustment of the main chain upper masses via coil-magnet actuators, while the test masses are controlled through gentler electrostatic actuation [47]. The entire quad-suspension is mounted on a triple-stage seismic isolation platform which gives three orders-of-magnitude of seismic isolation at 1 Hz [140]. Position and vibration sensors are used to measure the platform motion, which is then reduced by active damping.

2.1.1.2 Advanced Virgo Suspension

The adV detector in Italy has a mechanical structure built for test mass seismic isolation called the 'superattenuator'. This is formed of an inverted pendulum stage, a chain of passive mechanical filters, and the suspension payload [4]. The payload comprises of a penultimate mass known as the 'marionette', the test mass, and the reaction cage which lies around the marionette and test mass. These stages are labelled in Figure 2.4.

The inverted pendulum stage has three legs which are each connected to the ground via a flexible joint, providing a pre-isolation stage for low-frequency (40 mHz) horizontal seismic motion [4]. The 'Top Ring' mounting point of the suspension chain sits at the top of the inverted pendulum legs and houses the first of the filters. Also housed on the Top Ring are sensors and actuators for active reduction of low-frequency seismic noise and damping of superattenuator resonances. Below the Top Ring are a series of five further filters where blade springs and anti-springs are utilised to decrease the vertical seismic motion reaching the payload. The final filter stage is extended to sit around the marionette and test mass. Coil-magnet actuators are installed on the top-stage, marionette, and mirror to actively damp filter chain resonances between 200 mHz – 2 Hz, and to maintain cavity lock via control of the mirror position [4, 26].

The monolithic suspension stage of the 42 kg fused-silica mirror is shown in Figure 2.5. As in the aLIGO suspension, ears are bonded to flats on the mirror, however a nailhead design has been chosen for adV rather than welded horns. The ends of each fibre are bonded to a fused-silica anchor with one end clamped to the marionette and the other silicate bonded to the ear to form the monolithic suspension stage [13].

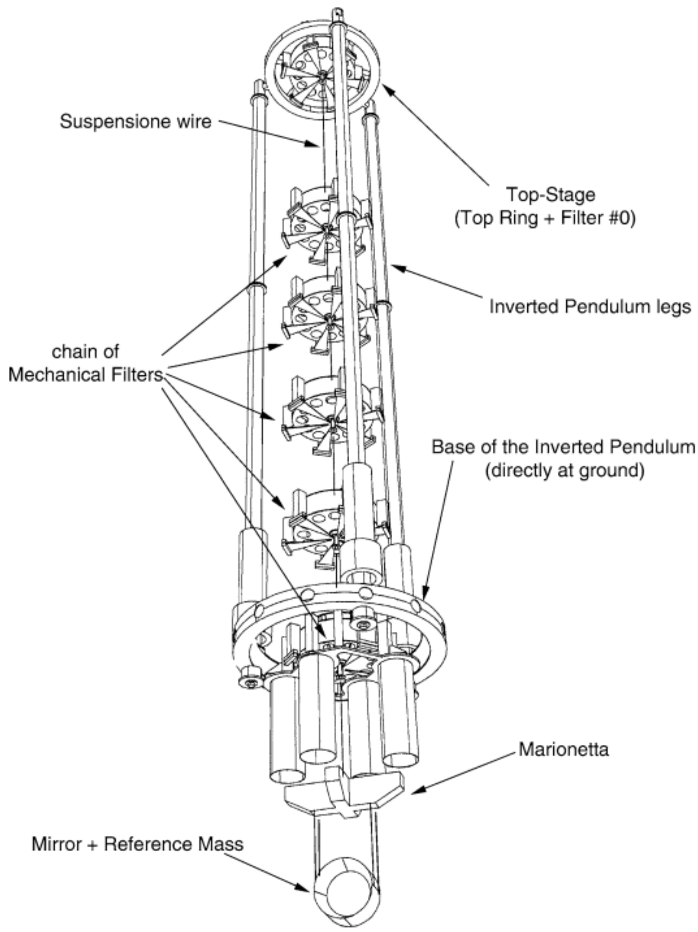

Figure 2.4: Diagram of the Virgo superattenuator showing the inverted pendulum stage, mechanical filters, the marionette, and the optic. The advanced Virgo upgrade uses an actuation cage around the marionette and mirror which replaces the reference mass shown in the image [66].

2.1.1.3 KAGRA Suspension

KAGRA is the only GW detector currently operating at cryogenic temperatures. As such, the design slightly differs from the room-temperature aLIGO and adV suspensions. The entire KAGRA suspension chain is nine stages long and 13.5 m in height with the upper five stages forming the 'Type-A' tower, operational at room temperature, and the lower four stages forming the cryogenic payload which is cooled to 20 K [14]. These stages can be seen in Figure 2.6. Similarly to the Virgo design, an inverted pendulum is used for pre-isolation from low-frequency horizontal seismic motion. Geometric anti-spring filters are installed to provide vertical seismic isolation at each stage of the Type-A tower. Geometric anti-springs are formed from triangular blade springs radially spaced around a ring at their base and attached to a common mass at their tip, creating a spherically symmetric, compact isolation system [123].

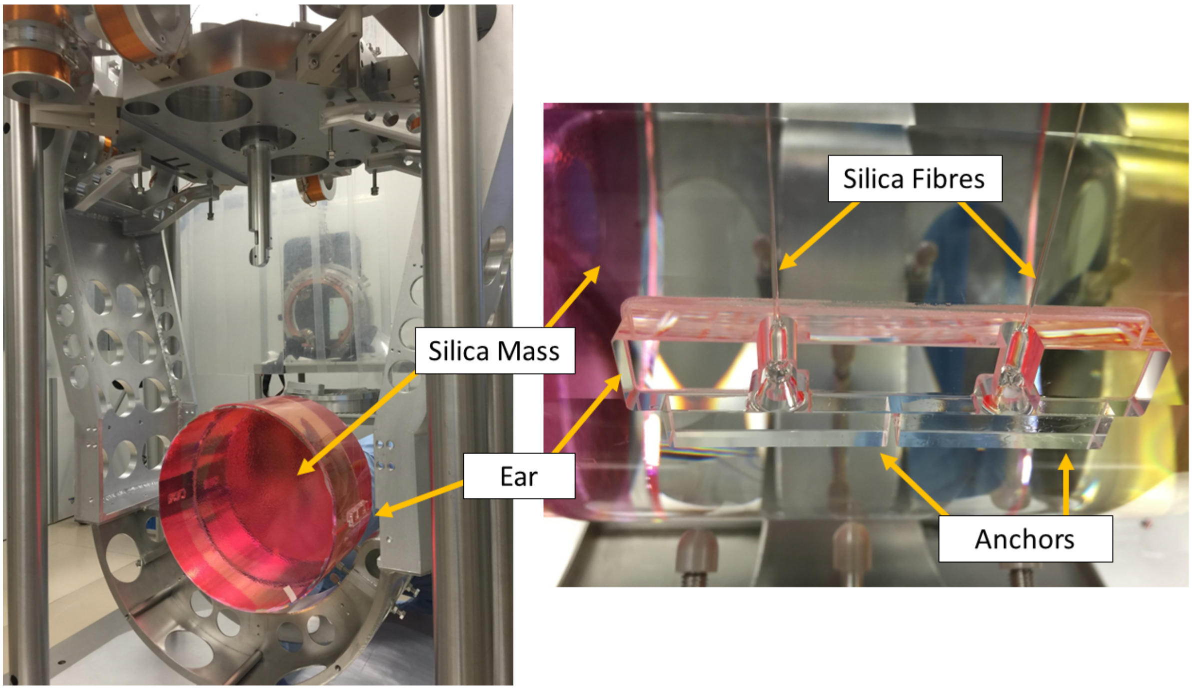

Figure 2.5: adV monolithic suspension showing a) mirror payload and b) details of the nailhead ear design [54].

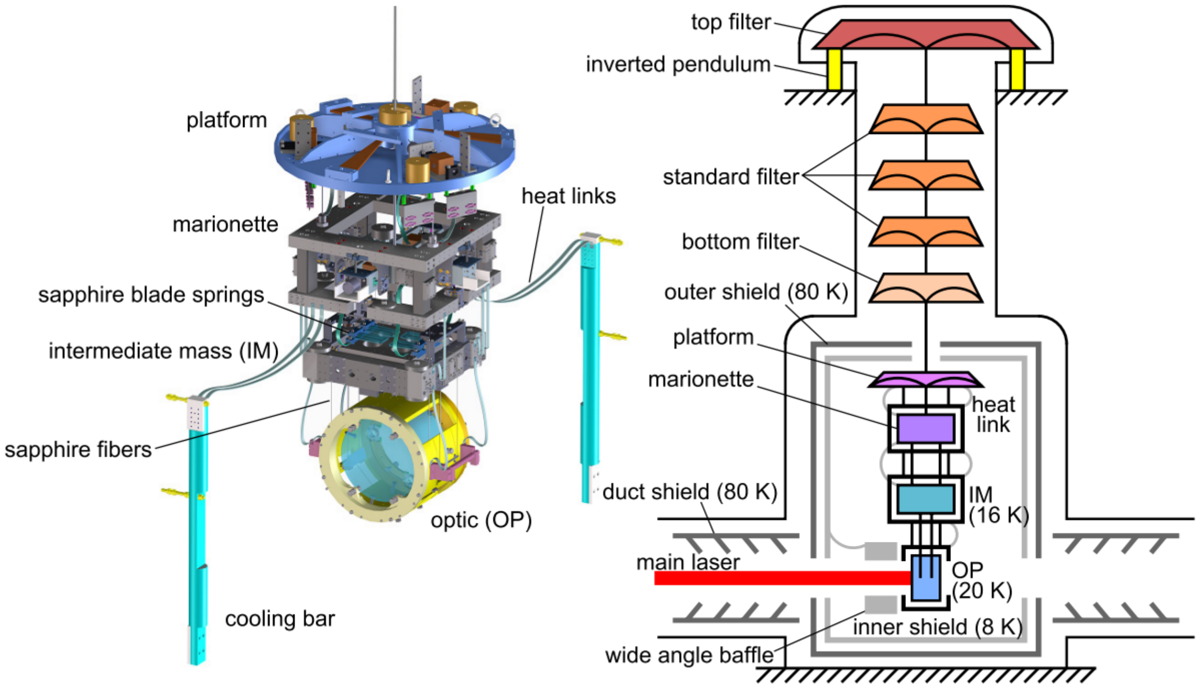

Figure 2.6: Diagram of the KAGRA suspension showing the cryogenic payload on the left and the Type A tower with attached cryogenic payload on the right [14].

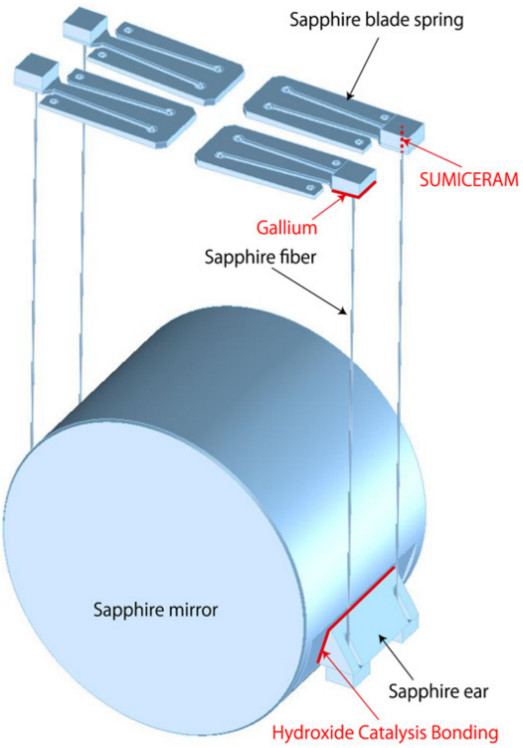

Figure 2.7: Schematic view of the KAGRA sapphire suspension showing the bonding between the sapphire mirror, ears, fibres, and blade springs [197].

The cryogenic payload is attached to the final filter of the Type-A tower by a single steel wire. The four stages of the cryogenic payload are made up of a stainless steel platform, marionette, intermediate mass, then a 23 kg sapphire test mass as shown on the left of Figure 2.6. The monolithic stage is shown in Figure 2.7 and is formed of a sapphire mirror with sapphire ears bonded to flats on each mirror side. Four sapphire fibres attach to the ears in a similar design to the Virgo suspension. At their top end, the fibres are bonded to sapphire blade springs which connect to the intermediate mass stage [14]. A reaction chain is hung from the platform stage for control of mass positions using coil-magnet actuation [144]. A full description of the cooling system of the cryogenic payload can be found in Section 4.3.5.1.

2.1.2 Suspension Design Considerations

An understanding of suspension systems is necessary for the calculation of resonant modes and to estimate the energy distribution and loss mechanisms of the system. Clever suspension design can allow manipulation of the material losses and energy distribution to minimise the total thermal noise of the system. An example of this would be the shaping of the fused-silica fibres used in current GW detectors. The ear-mirror bond and fibre-ear welded regions have a higher mechanical loss than the central region of the fibre [61]. To decrease the mechanical loss of the system it is important to minimise the energy which is dissipated in these regions. A fibre suspending some mass will have natural modes of oscillation. The energy dissipation of these modes is concentrated at the fibre bending point: the point along the fibre at which the normal to the mass's top surface intersects with the vertical axis as shown in Figure 2.3(b). The fibres are shaped to push this natural bending point further along its length and away from the bond and weld regions, resulting in a much lower energy dissipation in the lossy regions and creating a lower-noise suspension system [60].

The mechanical resonances found in a suspended mirror system are outlined in Section 2.2. An introduction to the mechanical losses of the suspension elements and calculation of the loss contribution to thermal noise via the fluctuation-dissipation theorem is in Section 2.3. Third generation GW detectors are planned to operate at cryogenic temperatures and as such suspension designs must move away from the fused-silica optics currently in use. Silicon and sapphire are both being investigated for use in these detectors due to their desirable material properties at low temperatures. Silicon will be used for the optics in the GCIF cryogenic cavity and is discussed as a suspension material in Section 2.4. The production of silicon is outlined, then the silicon material properties are introduced which will be used throughout this thesis.

2.2 Resonant Modes of a Suspension

All suspended mirror systems will have their own mechanical resonances arising from the geometry of the system. Most relevant in suspension thermal noise calculation are the pendulum modes (longitudinal in x), bounce (or vertical) modes, and violin modes. For a single mass suspended from two wires (which at their top are fixed and lossless) the pendulum mode, bounce mode, and violin mode of this system are as shown in Figure 2.8. There are additional modes corresponding to the other degrees of freedom of the masses: these are the pitch mode, roll mode, yaw mode, and y-longitudinal mode.

In this body of work the coordinate x refers to motion normal to the suspended mirror's face (along the axis of the interferometer arm), y motion is perpendicular to this, and z is motion in the vertical plane. Masses can be suspended from fibres, thin metal wires, or from rectangular cross-section ribbons. Here the case is considered for a mass suspended from two fibres of a circular cross-section which are attached to the mid-plane of a cylindrical

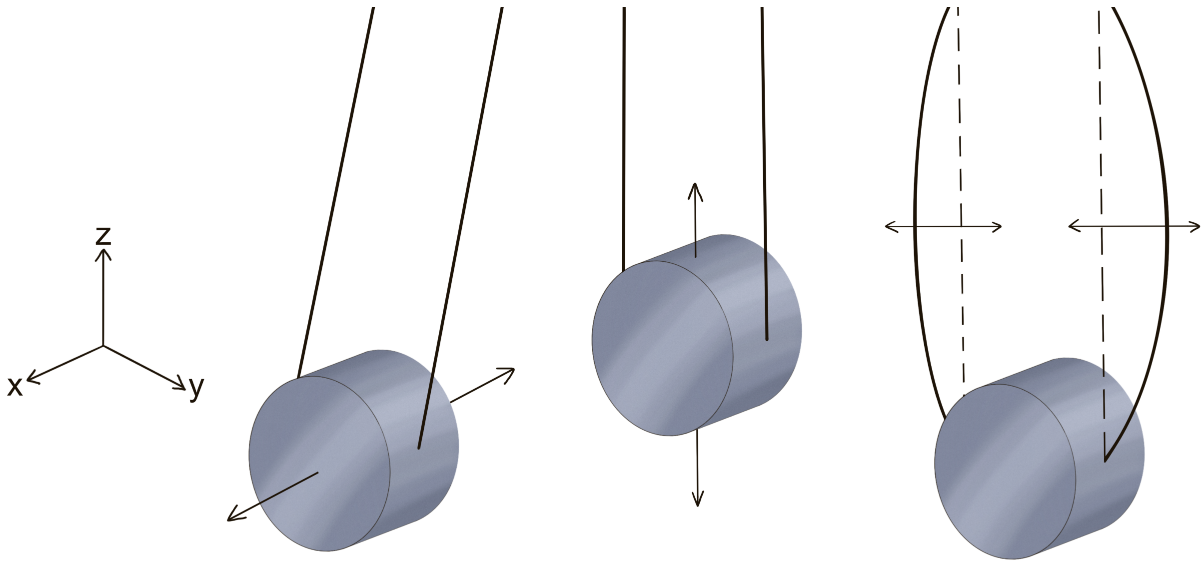

Figure 2.8: From left to right the suspension pendulum mode, vertical mode, and first order violin mode.

mass at the bottom end and a rigid, lossless clamp at the top end. Suspension resonant modes and the cross-coupling between these modes depends on the number of suspension fibres used and the attachment points of these fibres on the mass. An overview of resonant modes for this case is given in this chapter with a full analytical description of pendulum dynamics found in [193].

2.2.1 Pendulum Mode

The pendulum (or longitudinal) frequency of a single pendulum undergoing simple harmonic motion for small displacements has an angular frequency

$$\omega = \sqrt{\frac{k}{m}}. \tag{2.4}$$

In the case of a simple pendulum the restoring force comes from the gravitational restoring force, therefore the restoring spring constant, k, can be calculated by equating Hooke's law and Newton's 2nd Law such that

$$k = \frac{m\,g}{l} \tag{2.5}$$

for pendulum mass, m, gravitational acceleration, g, and pendulum fibre length, l [58]. Substituting Equation 2.5 into Equation 2.4 outputs the pendulum frequency, f_{pend}, as

$$f_{\mathrm{pend}} = \frac{1}{2\,\pi}\sqrt{\frac{k}{m}} = \frac{1}{2\,\pi}\sqrt{\frac{g}{l}}. \tag{2.6}$$

From Equation 2.1 a single-stage pendulum has a magnitude response of mirror motion divided by the ground motion which is flat and unity at frequencies below the suspension resonant peak at f_{pend} and drops off as $1/f^2$ above the resonant frequency. Seismic motion coupling to the mirror is therefore attenuated at frequencies where $f_{\mathrm{ground}} > f_{\mathrm{pend}}$. Multi-stage pendulum designs are typically used in detectors as the magnitude response of an n-stage pendulum decreases with frequency at $1/f^{2n}$ above the pendulum resonant modes, providing greater seismic isolation to the mirror than a single-stage pendulum of the same length [163].

2.2.2 Vertical Mode

The frequency of the vertical mode depends largely on the fibres used in the suspension and requires calculation of the vertical spring constant of the fibre,

$$k_{\mathrm{vert}} = \frac{Y\pi r^2}{l}, \tag{2.7}$$

for Young's modulus of fibre material Y, fibre radius r, and fibre length l [193]. For a freely moving mass suspended from n fibres, the vertical mode can be calculated as

$$f_{\mathrm{vert}} = \frac{1}{2\,\pi}\sqrt{\frac{n\,k_{\mathrm{vert}}}{m}}. \tag{2.8}$$

A small fraction of the energy in the vertical resonances is often observed coupling into longitudinal mirror motion. There are various sources of coupling arising from technical and fundamental sources such as mirror curvature or suspension misalignment. For kilometre-scale ground-based detectors there is additional coupling due to the curvature of the Earth. The cavity mirrors are suspended a sufficient distance apart that the local gravitational field at one mirror is not parallel to the local gravitational field of the second mirror, and the mirrors faces will not naturally hang parallel. Any vertical motion of one of these cavity mirrors therefore couples into longitudinal mirror motion and changes the length of the cavity by some small amount. This combination of effects is estimated to produce a cross-coupling factor of 0.1% in aLIGO and GEO600 [153, 160]. For a small lab-based experiment the local gravitational field will be uniform at both mirrors. A 0.1% cross-coupling could be used as a conservative estimate for the GCIF however this is likely an over-estimate with the coupling factor lying closer to 0.01% [57, 149].

2.2.3 Violin Mode

Violin modes are caused by the resonance of standing waves in the suspension fibres. There will be multiple standing waves per fibre, with the lowest frequency mode being that of a single standing wave. Calculation of suspension violin modes is done as for a simple waves-on-a-string case. Firstly, the mass per unit length in the fibre is found from fibre radius and fibre material density, $\pi\, r^2\, \rho$. This is used to calculate the velocity of a wave travelling along the fibre,

$$v_{\text{wave}} = \sqrt{\frac{F}{\pi\, r^2\, \rho}}, \tag{2.9}$$

with F the tension in the fibre. The frequency of the p^{th} resonant violin mode is

$$f_{\text{violin,p}} = \frac{p\, v_{\text{wave}}}{2\, l} \tag{2.10}$$

if the bending points of the violin modes are assumed to lie at the end of each fibre. In reality this bending length is some small distance along the fibre length which marginally increases $f_{\text{violin,p}}$.

2.2.4 Other Resonant Modes

The pitch and yaw modes are caused by tilt of the mass around the y- and z- axes respectively, shown in Figure 2.9. Careful centring of the laser beam on the mirror can lessen the pitch and yaw coupling to cavity length noise but any slight beam misalignment will produce a longitudinal signal response from the interferometer at these resonances. The percentage coupling of these modes from a beam offset can be calculated through trigonometry. The yaw resonances of the aLIGO suspension occur at low frequencies, typically outside the sensitivity band, and are strongly decoupled from the other mirror resonances, however the pitch modes can be observed [17, 168]. The yaw and pitch modes are damped via coil-magnet actuation from the reaction chain [97, 145].

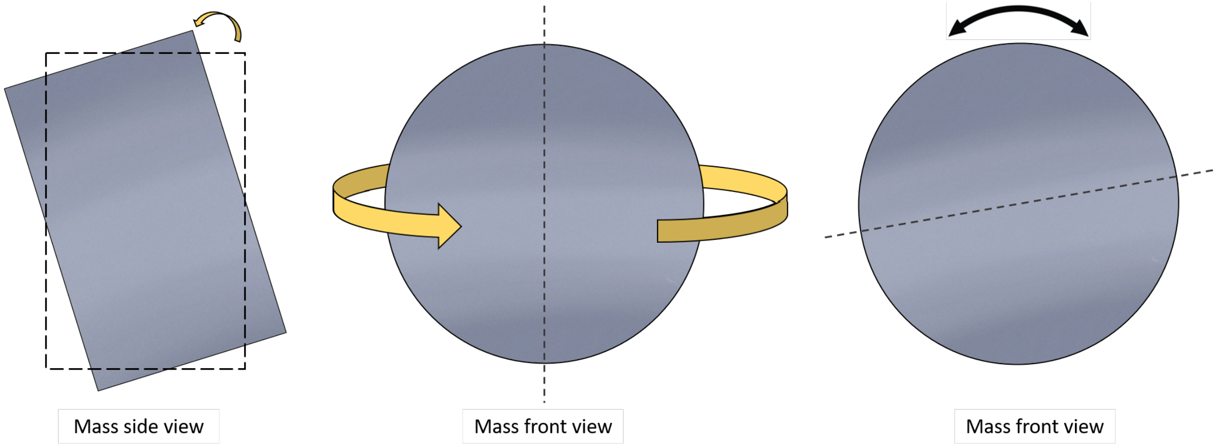

Figure 2.9: From left to right the mass pitch mode, yaw mode, and roll mode.

The roll mode is the name given to tilt of the mass around the x-axis, also shown in Figure 2.9. There is no first order coupling of roll to longitudinal motion however some higher order coupling can be seen which arises from technical sources. Currently the aLIGO detectors have no damping of the test mass roll resonance.

The y-longitudinal mode is motion of the mass linearly along the y-axis. There is no direct coupling of this mode to mirror motion along the cavity axis, however the concave mirror surface will affect the cavity resonance condition if there is y-axis displacement. A mismatch of the laser beam wavefront and the mirror curvature will decrease resonance in the cavity, leading to some loss of power in the interferometer arm. This is controlled via actuation of the mirror position to lock the cavity on-resonance [17].

2.3 Thermal Noise of a Suspension

One of the main aims of the 10 m interferometer prototype at the GCIF is a direct measurement of coating thermal noise between 100 Hz - 1 kHz using materials suitable for use in 3G cryogenic detectors. To achieve this goal the thermal noise of the suspension chain must be kept as low as possible to allow coating thermal noise to dominate a thermal noise measurement. Introduced in Section 1.2.3.3 and detailed in Section 2.3.3, the thermal displacement motion is directly linked to the mechanical losses of the suspension materials through the fluctuation-dissipation theorem. A full consideration of suspension thermal noise must therefore start with the mechanical loss mechanisms in the system. This section outlines the calculation of the displacement thermal noise of a 1 kg mass attached to four crystalline fibres as planned for use in the GCIF.

2.3.1 Mechanical Loss Contributions to Thermal Noise

Mechanical loss in a material describes the degree to which the material dissipates energy through its internal atomic bonds. For a driving force incident on one side of some material, this force must propagate through that material from bond to bond across some finite time. During the propagation there will be some energy lost due to internal frictions inside the material. Mechanical loss is quantified by a loss angle, $\phi(\omega)$: the ratio of the imaginary to real parts of the material Young's modulus. This represents the phase angle at which the strain response lags behind the driving force [58, 78]. There are four mechanisms contributing to the mechanical loss of the suspension: the material bulk loss, surface loss, thermoelastic loss, and weld/bond loss, detailed below. These mechanisms contribute to the mechanical loss of the key suspension modes (pendulum, vertical, and violin) with the level of contribution depending upon the suspension's energy distribution, and the dissipation dilution of each resonant mode, D_{mode}. The dilution factor is the ratio of the energy stored in the gravitational field compared to that in the fibre, and is covered in Section 2.3.2.

Estimations of the energy distribution across a suspension system are typically calculated through finite-element analysis (FEA) [60]. This approach requires detailed knowledge of the suspension geometry including fibre neck shaping, fibre bending points, and ear dimensions. These are not yet known for the GCIF cryogenic suspension therefore the energy ratio was assumed to be similar to that of the aLIGO fused-silica suspensions: the pendulum and first-order violin modes have fibre energy dissipation occurring at a similar bending point position where 90% of the energy is stored in the bending region and 10% in the bond. In the excitation of the vertical mode, energy dissipation is concentrated in the mid-point of the fibre (far from the welded regions at the fibre ends) and there is a negligible contribution from the weld region [60, 62]. The loss contribution for the pendulum mode is thus given by

$$\phi_{\text{pend}} = \frac{0.9 \times (\phi_{\text{bulk}} + \phi_{\text{surface,mode}} + \phi_{\text{thermoelastic}}) + 0.1 \times \phi_{\text{joint}}}{D_{\text{mode}}}, \tag{2.11}$$

for the vertical mode by

$$\phi_{\text{vert}} = \phi_{\text{bulk}} + \phi_{\text{surface,mode}}, \tag{2.12}$$

and for the violin mode by

$$\phi_{\text{violin}} = \frac{0.9 \times (\phi_{\text{bulk}} + \phi_{\text{surface,mode}} + \phi_{\text{thermoelastic}}) + 0.1 \times \phi_{\text{joint}}}{D_{\text{mode}}} \tag{2.13}$$

[59]. All four loss mechanisms contribute to loss in the pendulum mode and violin modes, while the vertical mode contains no contribution from thermoelastic loss or joint loss and no dilution as detailed in Section 2.3.2. A recommendation for future work would be to perform FEA studies with a more mature design of the GCIF triple suspension to find the exact energy distribution across the fibre.

2.3.1.1 Bulk Loss

Bulk loss is the mechanical loss caused by internal frictions from crystal lattice imperfections or impurities in the substrate. For samples with a large volume-to-surface ratio this loss is measured through 'ringdown' measurements which observe the Q-factor of various material resonances. These measurements can be performed at different frequencies and temperatures to build up a picture of the variation across temperature and frequency space [102, 147]. This mechanism is observed to be affected by material doping and varies with both temperature and frequency as seen in [196, 218].

2.3.1.2 Surface Loss

Surface loss is also caused by crystal lattice imperfections and is typically lossier than the material's bulk loss down to some depth h. This is due to damage at the surface interface where the surface material can be subjected to shallow damage, cracks, and impurities from materials in contact with the surface [94, 111]. The frequency-independent surface loss contribution can be calculated by multiplication of the mechanical loss of the material's surface, ϕ_{s}, the depth h of the surface loss mechanisms, and the ratio of surface area to volume, $\frac{S}{V}$, equalling $\frac{2}{r}$ for a circular cross-section fibre. A factor $c_{\mathrm{geometric}}$ takes into consideration the radial distribution of strain based on the fibre geometry and class of resonant mode. For the vertical mode the strain is distributed evenly across the fibre radius and $c_{\mathrm{geometric}} = 1$ whereas any transverse modes must account for the stretching and compression across the fibre surface. From [94], $c_{\mathrm{geometric}} = 2$ for transverse oscillations in a fibre geometry. The surface loss contribution to the pendulum and violin modes is therefore

$$\phi_{\mathrm{surface,transverse}} = \frac{4\,h\,\phi_{\mathrm{s}}}{r}, \tag{2.14}$$

while for the vertical mode

$$\phi_{\mathrm{surface,vertical}} = \frac{2\,h\,\phi_{\mathrm{s}}}{r}. \tag{2.15}$$

The surface loss contribution to thermal noise is often considered as the material surface loss multiplied by the depth to which these occur due to the difficulties of measuring h and ϕ_{s} individually [63]. For materials with a high surface-to-volume ratio the surface loss is expected to dominate ringdown measurements [61, 111].

2.3.1.3 Thermoelastic Loss

Thermoelastic loss of a suspension fibre is the loss contribution occurring when local temperature gradients in the fibre cause expansion or contraction, leading to bending in the fibre and hence longitudinal movement of the suspended mass. Thermoelastic loss of a suspension system is dependent on both frequency and temperature and can be calculated as

$$\phi_{\text{thermoelastic}} = \frac{\omega\,\tau}{1+(\omega\,\tau)^2} \times \frac{Y\,T}{\rho\,c_{\mathrm{p}}} \times \left(\alpha - \frac{\beta\,\sigma_0}{Y}\right)^2. \tag{2.16}$$

τ is the characteristic time for heat to cross the fibre,

$$\tau = \frac{(2\,r)^2\,\rho\,c_{\mathrm{p}}}{4.32\,\pi\,\kappa}, \tag{2.17}$$

where c_{p} is the specific heat capacity at constant pressure, α is the coefficient of thermal expansion, $\beta = \frac{1}{Y}\frac{\mathrm{d}Y}{\mathrm{d}T}$ is the thermal elastic coefficient, σ_0 is fibre tensile stress, and κ is the thermal conductivity. For many materials the thermal elastic coefficient, β, is negative, and thermoelastic loss increases as fibre tensile stress, σ_0, is increased. The fused-silica in use at room-temperature detectors has a positive thermal elastic coefficient, a property which can be exploited to 'null' thermoelastic loss through a careful choice of fibre stress as $\sigma_0 = \frac{\alpha Y}{\beta}$ [58]. Thermoelastic loss in cryogenic suspension materials will be reduced by the low operation temperature rather than through thermoelastic nulling.

2.3.1.4 Weld and Bond Loss

In a suspension system with fibres suspending some mass there must be jointing between the fibre-ear interfaces and the ear-mass interfaces, with this jointing region contributing some loss to the suspension. Welding and bonding are techniques currently utilised in the fused-silica suspensions at aLIGO, with both jointing methods being investigated for their suitability for use in 3G detectors. For 3G detectors, this join region must have low mechanical loss and low energy dissipation to reduce the contribution to thermal noise, and be suitable for use in vacuum and at cryogenic temperatures. Jointing contribution has been observed in room-temperature detectors and is expected to be observed in cryogenic crystalline suspensions [59].

Hydroxide-catalysis (HC) bonding is a technique developed at Stanford University [77]. It has been demonstrated to work effectively in vacuum and at low temperature and is currently in use at the aLIGO and AdV detectors [13, 17]. In HC bonding, two silica surfaces have a silicate-like network created between them using an alkaline bonding solution which acts to liberate silicate ions from the two surfaces. In the bonding process the silicate ions disassociate, combine, and polymerise into siloxane chains and water. The water evaporates and leaves a network of siloxane chains forming a bond between the surfaces. HC bonding of silicon first requires growth of a thermal oxide layer [28]. Cured silicon-silicon HC bonds have a tensile strength of 30 ± 17 MPa, with a bond and oxide layer combined thickness of 200 nm [151].

The welding of crystalline fibres is currently an area of active research. Direct silicon-silicon welding was first attempted at the University of Glasgow via laser heating. In this process silicon's quick transition from crystalline solid to low viscosity liquid made it difficult to create the viscous molten zone necessary for controllable welding to take place [74]. More promising results have been observed by the German Leibniz-Institut für Kristallzüchtung through demonstration of silicon-silicon fibre welding using the float zone method, producing welds for diameters up to 28 mm [32]. The float zone method for the manufacture of high-purity silicon is discussed in Section 2.4.1.1. Further laser welding techniques are being tested for silica-clad silicon fibres using a metal-semiconductor alloy in the weld, with the cladding providing containment of the molten area and a higher control of the welding region [84]. The early status of this research means there is not yet available data on the mechanical loss of silicon-silicon welds and so the mechanical loss of a silicon-silicon HC bond was used as the jointing loss for thermal noise modelling in this work. In the case of silicon welding proving unfeasible, one solution for the GCIF suspensions could be HC bonding.

2.3.2 Dissipation Dilution

For the mass of a pendulum displaced longitudinally by some amount x, there will be some restoring force on this mass both from the gravitational restoring force

$$E_{\mathrm{g}} = \frac{1}{2}\, k_{\mathrm{g}}\, x^2, \tag{2.18}$$

and from the stiffness of the pendulum fibres,

$$E_{\mathrm{f}} = \frac{1}{2}\, k_{\mathrm{f}}\, x^2. \tag{2.19}$$

$k_{\mathrm{g}} = \frac{m\,g}{l}$ is the lossless gravitational spring constant while $k_{\mathrm{f}} = \frac{n\,\sqrt{F\,Y\,I}}{2\,l^2}$ is the elastic spring constant for pendulum fibre with some longitudinally displaced mass [164], and I is the second moment of area of the fibre

$$I = \frac{\pi\, r^4}{4}. \tag{2.20}$$

The pendulum mode dilution factor, D_{long}, is the ratio of the energy stored in the gravitational field compared to that in the fibre. This can be calculated as

$$D_{\mathrm{pend}} = \frac{k_{\mathrm{g}}}{k_{\mathrm{f}}} = \frac{\frac{m\,g}{l}}{\frac{n\,\sqrt{F\,Y\,I}}{2\,l^2}} = \frac{2\,m\,g\,l}{n\,c_{\mathrm{bend}}\,\sqrt{F\,Y\,I}}, \tag{2.21}$$

where we must introduce multiplication by a bending factor, c_{bend}. This factor $c_{\mathrm{bend}} = 2$ for a fibre with two bending points while $c_{\mathrm{bend}} = 1$ for a fibre with a single bending point as shown in Figure 2.10. No single-fibre suspensions are considered in this thesis therefore $c_{\mathrm{bend}} = 2$ for all cases in this work. Excitation of a pure vertical mode should produce no longitudinal movement of the test mass, therefore the dilution factor of this mode is $D_{\mathrm{vert}} = 1$. It is predicted in [134] that the total energy stored in the violin mode resonance of the fibres is half that of the energy stored in the pendulum mode. From this, it follows that the dilution factor for the violin mode is half that of the pendulum mode, $D_{\mathrm{violin}} = \frac{1}{2} \times D_{\mathrm{pend}}$, as used in [85].

2.3.3 Displacement Noise

The fluctuation-dissipation theorem states that a measurable parameter of a dissipative system is subject to random thermal fluctuations. In the case of a fibre suspension, the thermal fluctuations describe the thermal displacement noise while the dissipative part is described by the imaginary part of the mechanical impedance, contributed to via mech-

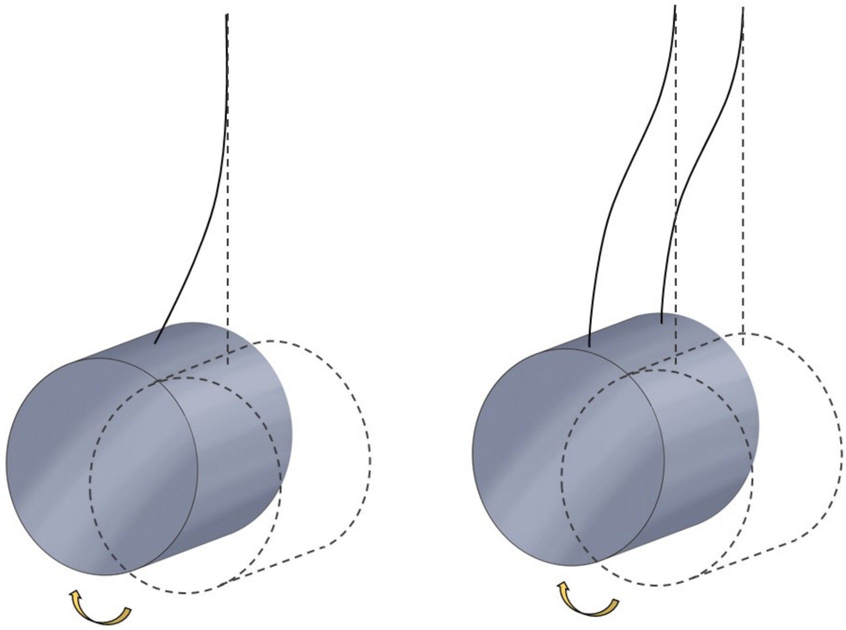

Figure 2.10: Depiction of a fibre with a single bending point (left) and a fibre with two bending points (right).

anical losses. This theorem therefore allows calculation of displacement thermal noise for known mechanical loss of a system. In this system of a single test mass suspended by n fibres the motion of the mass in the longitudinal direction is the displacement noise. The root-mean-square thermal displacement for some frequency, ω, can be given by

$$x_{\mathrm{rms,mode}}\left(\omega\right) = \sqrt{\frac{4\,k_{\mathrm{B}}\,T}{m\,\omega} \times \frac{\phi_{\mathrm{mode}} \times \omega_{\mathrm{mode}}^{2}}{\left(\phi_{\mathrm{mode}} \times \omega_{\mathrm{mode}}^{2}\right)^{2} + \left(\omega_{\mathrm{mode}}^{2} - \omega^{2}\right)^{2}}}, \tag{2.22}$$

for a given resonant frequency, ω_{mode}, being the pendulum, vertical, or violin modes [35, 36]. In the case of the pendulum mode, the frequency of the pendulum mode, ω_{pend}, and the mechanical loss of this mode, ϕ_{pend}, can simply be subbed in for ω_{mode} and ϕ_{mode} respectively.

When considering the displacement noise contribution from the vertical mode, the fraction of vertical loss which is transferred to longitudinal motion must be considered as introduced in Section 2.2.2. The vertical mode resonance frequency, ω_{vert}, and mechanical loss, ϕ_{vert}, replace ω_{mode} and ϕ_{mode} in this equation before multiplication of the vertical-to-horizontal cross-coupling factor, with 0.1% typically being used for aLIGO.

In the excitation of a violin mode the oscillation of the fibres will cause a small longitudinal motion of the test mass. The scale of this motion is determined by the *mass coupling factor*, which must be found for each violin mode, p, as the ratio of the mass and fibres, giving

$$\sqrt{\frac{2\,m_{\text{fibre}}}{\pi^2\,m\,p^2}}, \tag{2.23}$$

where fibre mass can be calculated as

$$m_{\text{fibre}} = \rho\,\pi\,r^2\,l. \tag{2.24}$$

The test mass thermal displacement, $x_{\text{rms,violin,p}}(\omega)$, for each violin mode can be found by inputting ϕ_{violin} and $\omega_{\text{violin,p}}$ to the thermal displacement equation, and multiplying this by the mass coupling factor for each resonant mode, giving

$$x_{\text{rms,violin,p}}(\omega) = \sqrt{\frac{2\,m_{\text{fibre}}}{\pi^2\,m\,p^2} \times \frac{4\,k_{\text{B}}\,T}{m\,\omega} \times \frac{\phi_{\text{violin}} \times \omega^2_{\text{violin,p}}}{\left(\phi_{\text{violin}} \times \omega^2_{\text{violin,p}}\right)^2 + \left(\omega^2_{\text{violin,p}} - \omega^2\right)^2}}. \tag{2.25}$$

The total test mass thermal displacement noise across the frequency space from all resonant modes can be summed in quadrature as

$$x_{\text{rms,total}}(\omega) = \sqrt{x_{\text{rms,pend}}(\omega)^2 + x_{\text{rms,vertical}}(\omega)^2 + 4 \times x_{\text{rms,violin,p}}(\omega)^2}. \tag{2.26}$$

The violin modes are independent and should each be added in quadrature. Circular fibres will oscillate in various directions, however a $4\times$ multiplication factor is assumed as a worst-case contribution of violin mode noise whereby the four fibres are excited only in the direction of the arm cavity, and directly injecting motion to the test mass along this axis [57].

2.4 Crystalline Suspensions

2.4.1 Silicon

GW detectors are moving towards cryogenic operation to reduce the thermal noise currently limiting the mid-detection band. Crystalline silicon and sapphire are being investigated to replace the fused-silica which is used in 2nd generation room temperature suspensions. Silicon is an attractive choice as a suspension material due to its high thermal conductivity and low mechanical loss at low temperatures, and in particular due to the nulling of its coefficient of thermal expansion at 18 K and 123 K which renders the thermoelastic loss contribution to thermal noise negligible at these temperatures.

2.4.1.1 Silicon Production

Silicon dioxide (SiO_2) can commonly be found in nature as quartz and can be reduced with carbon in an electric arc furnace for the commercial production of crystalline silicon. From elemental silicon a single-crystal ingot can be manufactured via the Czochralski method or the float zone method. In the Czochralski method, elemental silicon is melted in a crucible before the introduction of a seed crystal which contacts the molten silicon and is drawn upwards to create a single-crystal ingot. The contacting of the molten silicon and the crucible in the Czochralski method can be a source of impurities in the final silicon ingot, making the crucible-less float zone method preferred for higher purity ingot production [203]. In the float zone method, a zone of molten silicon is created on the bottom of an ingot by inductive heating. A seed crystal is introduced to the molten zone and the ingot is passed through the induction coil, creating a single-crystal ingot [214]. In this method the molten silicon zone must have a high enough surface tension to retain its shape which places a limit on ingot diameter of near 200 mm [146].

The ingots can be sliced into wafers, which are used commonly in the semiconductor production industry. From these wafers silicon ribbons can be easily manufactured by laser cutting [203].

Silicon fibre production for optical suspensions is a newly emerging field, but advancements are being driven by the GW community alongside the semiconductor industry. Silicon fibres of tens of centimetres have been drawn by crystal growth techniques in the past, however these are too short for use in GW suspensions [118, 169]. Recent advancements have demonstrated production of fibres >1 m long using the float zone technique [202]. These fibres have a diameter of 3 mm as required for a silicon ET suspension, with a maximum surface deviation of <0.1 mm. Alongside this, promising developments of silica-clad silicon fibres have shown that molten-core drawing can produce fibres on kilometre-length scales with width down to 55 µm [170, 205]. Laser drilling is being investigated as a method of removing the silica cladding without damage to the silicon fibre surface by using a series of laser-drilled pores to weaken the glass layer and an interface layer to encourage a clean separation of materials [41]. Chemical-mechanical processes have also been attempted [41]. Unclad silicon fibres of appropriate dimension and surface quality for 3G detectors have yet to be demonstrated at the time of writing. Design assumptions concerning silicon fibre production have had to be made throughout the course of this work which are based upon the current available technologies in an evolving field.

2.4.1.2 Silicon Material Properties

The pure silicon considered for 3G detectors has a bulk loss which decreases with decreasing temperature. A loss peak from introduced oxygen atoms is often seen in silicon manufactured via the Czochralski method [147]. The bulk loss values used in this body of work were taken from measurements of a high-Q silicon crystal manufactured via the Czochralski method, with data from 1.3 – 300 K [142]. The surface loss of silicon is taken as $h\,\phi_{\mathrm{s}} = 5 \times 10^{-13}$ from loss measurements of high-purity thin cantilevers between 5 – 300 K [63, 111]. Silicon-silicon hydroxide catalysis mechanical bond loss increases with increasing bond thickness and increasing temperature [151]. The mechanical loss value for silicon HC bonding was taken from [200], with bond loss increasing from 5×10^{-4} to 3×10^{-2} between 8 – 300 K for a bond thickness lying between 390 – 1120 nm. The contributions of these intrinsic and extrinsic silicon properties depend on the geometry and design of the suspension they will be used in. Figure 2.11 shows a graphical representation of these losses for the suspension design detailed in Chapter 4.

The temperature independent material properties of silicon are shown in Table 2.1. The Young's modulus, Y, of a material is a measure of stiffness under an applied load, with silicon Y varying with crystal axis orientation. The density and Young's modulus values were taken from [148] and [106] respectively, with Y for the (110) crystal orientation. These

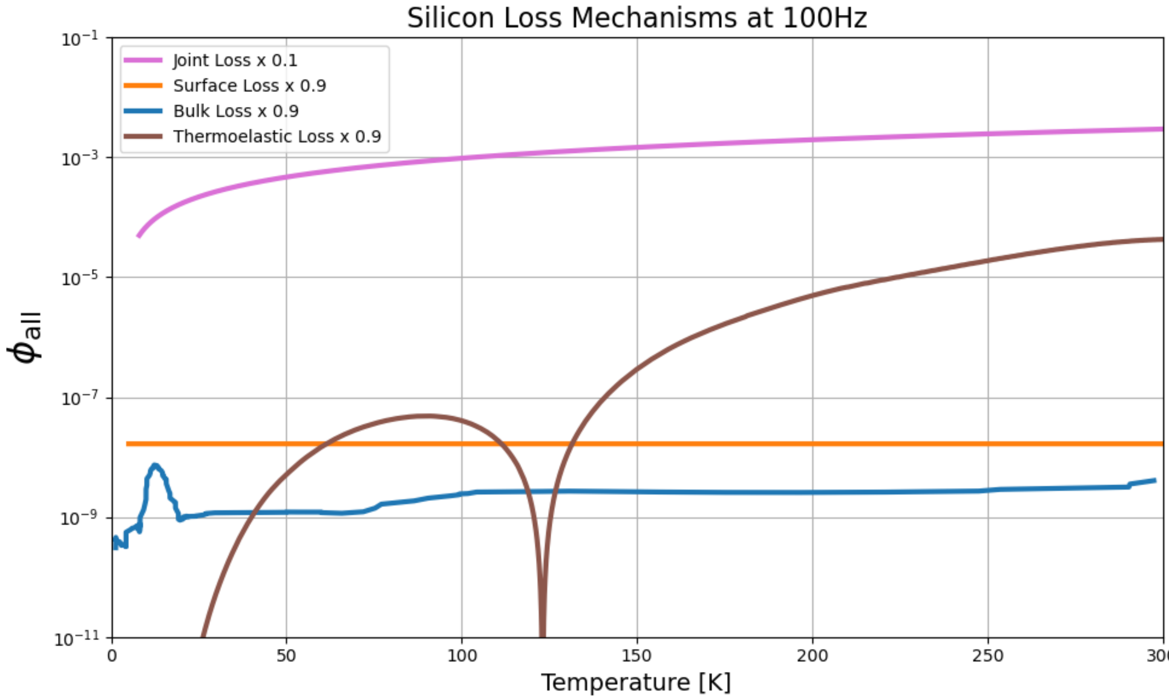

Figure 2.11: Silicon loss mechanism values used throughout this thesis, showing the joint loss, surface loss, bulk loss, and thermoelastic loss. Joint loss (pink) was taken from [200] for a silicon-silicon HC bond. Surface loss data (orange) was taken from [111]. The bulk loss values (blue) were taken from [142]. Thermoelastic loss (brown) was calculated using Eq. 2.16 for a frequency of 100 Hz.

Symbol	**Property**	**Value**
ρ	Density	2300 kgm^{-3}
ν	Poisson's Ratio	0.3
Y	Young's Modulus	169 GPa
ϵ	Surface Emissivity	0.5
σ_0	Breaking Stress	200 - 300 MPa

Table 2.1: Silicon temperature independent material properties.

properties have a variance of $\leq$1% between $0 - 300$ K so were treated as independent of temperature in this body of work. Poisson's ratio was taken from the COMSOL Material Library [55] and surface emissivity from [112]. Breaking stress was measured in tension at room temperature by [178].

Table 2.2 shows the material properties of silicon which vary with temperature. The specific heat capacity and thermal expansion coefficient data was taken from [106]. Silicon's thermal conductivity is strongly dependent upon the geometry of sample measured (see Section 4.3.3) and the doping of the crystal. A suspension material with a high thermal conductivity is essential for 3G detectors. The masses and fibres must be able to quickly dissipate the heat flux from the laser cavity to remain at cryogenic temperatures, and also to reduce thermal gradients across the face of the test mass which would contribute to thermal displacement noise in the cavity. Pure (undoped) silicon exhibits the highest thermal conductivity and lower thermal noise, and so a high purity silicon was considered in this body of work. The thermal conductivity values for high purity silicon were taken from [194], a recommended thermal conductivity curve based on measurements of 80 silicon samples of differing experimental methods, dimensions, and purities.

2.4.2 Sapphire

Sapphire is currently in use as the monolithic stage material at the KAGRA site [15]. The use of sapphire is being considered in two planned 3rd generation detectors (the Einstein Telescope and LIGO Voyager) however it is not currently the baseline material choice for either detector, which have both chosen to favour silicon suspensions [8, 119]. As a crystalline material sapphire benefits from a high thermal conductivity and low mechanical loss at low temperatures which are comparable to those of silicon. Heat extraction through sapphire fibres has been previously explored by Japanese groups for KAGRA which provides an advantage over the research and development which would be necessary for silicon suspension fibres [117]. Sapphire welding technology is currently under development in a slightly more advanced state than that of silicon - welds with no visible defects have been demonstrated between sapphire fibres of diameters between 0.425 mm – 2 mm, although characterisation of the welded regions is yet to be performed [172]. Despite these advantages sapphire does not share the nulling of silicon's coefficient of thermal expansion, and displays potentially weaker HC bonding than silicon [65]. One further potential limitation of sapphire's use in 3G detectors is the lack of availability of high-purity sapphire boules in large enough diameters for manufacturing of the optic. Here, silicon has the advantage as the main material used in the ever-growing semiconductor industry which drives the development of large, high-purity silicon crystals, and is the reason why silicon is the baseline material choice for ET and LIGO Voyager [183]. This thesis focuses on the use of silicon in the monolithic stage of cryogenic suspensions, as the material chosen for the GCIF cavity optics and as the baseline choice for 3rd generation GW detector optics, but considers the inclusion of sapphire fibres as part of a multi-material suspension design for the GCIF.

2.5 Conclusion

This chapter provides an introduction to suspension dynamics and to the displacement thermal noise contribution to the test mass from the suspension system. This thermal noise is affected by the mechanical loss of the suspension materials, contributed to via bulk loss, surface loss, thermoelastic loss, and the loss of some jointing region. The mechanical loss of the system can be minimised through careful suspension design to reduce energy dissipation in high-loss regions, the creation of a high-Q system to contain thermal noise in the suspension resonances, and through suspension material choice. Silicon and sapphire are discussed as the potential material choices for 3rd generation GW detectors, with silicon the material focused on within this body of work.

Temperature T [K]	**Specific Heat Capacity** C_{p} [$\mathrm{Jkg^{-1}K^{-1}}$]	**Coeff. of Thermal Expansion** α [$\mathrm{K^{-1}}$]	**Thermal Conductivity** κ [$\mathrm{Wm^{-1}K^{-1}}$]
4.2	0.034	1.11×10^{-10}	260
20	3.41	-2.5×10^{-9}	4940
40	44.1	-1.6×10^{-7}	3530
60	115	-3.65×10^{-7}	2110
80	188	-4.65×10^{-7}	1340
100	259	-4.23×10^{-7}	884
120	328	-2.83×10^{-8}	615
140	395	4.81×10^{-7}	460
160	456	8.93×10^{-7}	368
180	511	1.22×10^{-6}	307
200	557	1.52×10^{-6}	264
220	597	1.83×10^{-6}	230
240	632	2.07×10^{-6}	202
260	665	2.26×10^{-6}	180
280	691	2.49×10^{-6}	162
300	713	2.65×10^{-6}	148

Table 2.2: Silicon temperature dependent material properties.

Chapter 3

ROOM TEMPERATURE DOUBLE SUSPENSIONS

3.1 Suspension Design for the GCIF Room Temperature Optics

The Glasgow prototype interferometer will comprise of a 10 cm cryogenic cavity for testing of optical coatings and suspended silicon mirror controllability, a 10 m long room-temperature arm cavity for laser stabilisation, and various table-top systems for monitoring and laser input optics, with the layout shown in Figure 1.11. The design of this cryogenic cavity requires suspended mirrors suitable for use at extremely low temperatures and compatible with the near-infrared laser wavelengths chosen for next generation GW detectors. The design of the cryogenic cavity suspension is discussed in detail in Chapter 4. Alongside the cryogenic suspensions, the room-temperature suspensions play an equally important role, and the design of two room-temperature suspensions are outlined in this chapter. The first suspension design considered is for the 10 m arm cavity mirror suspensions. Secondly, the design is discussed for a small beam-steering suspension which will sit between the 10 m reference cavity and the cryogenic cavity to steer the beam into the cryostat. These suspensions must not couple too much additional noise to the beam during this journey.

Section 3.2 of this chapter describes the detailed modelling performed on two pre-existing double suspensions, including the verification of the model through comparison of resonant mode outputs, and calculation of the length noise added to the system by these suspensions. This modelling framework was used to design modifications to the two existing suspensions in Section 3.3, with the aim of meeting the length noise requirement of the laboratory. The 10 m cavity suspension was then characterised through measurements of displacement transfer functions with the methodology covered in Section 3.4. The results of the modelling and characterisation work are discussed in Section 3.5, providing a comparison of the model to the physical suspensions, and between the unmodified and modified suspension builds.

3.1.1 Design Sensitivity Requirements

The optical coating used on the cryogenic cavity optics will contribute thermal noise to the system which is sensed as length noise in the cavity. To make a direct measurement of coating thermal noise, this noise must be the dominant noise source in the cryogenic cavity and thus is set as the laboratory sensitivity requirement across the desired frequency range of 100 Hz – 1 kHz. The lower frequency limit is set by the seismic noise spectrum. The upper 1 kHz upper bound is set by the cavity control loop bandwidth, limited by the sample rate of the control and data acquisition system of the GCIF. Displacement thermal noise calculations were considered for a silica-tantala coating (such as might be used in the GCIF direct thermal noise experiment) following the method in [99] by the author with the help of Dr. I. Martin. For the two beam spot sizes incident on the cryogenic cavity mirrors this calculates a coating length noise of $1 \times 10^{-17}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ on the initial test mass, or ITM (with a 83 µm beam radius) and $1.4 \times 10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ on the end test mass, or ETM (with a 588 µm beam radius), calculated at 100 Hz and 123 K operation. At 18 K operation for a frequency of 100 Hz the coating length noise is $4 \times 10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ on the ITM and $5.7 \times 10^{-19}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ on the ETM. These values have a $^1/_{\sqrt{f}}$ dependence on frequency above 100 Hz.

The cryogenic cavity is considered to be formed by two identical and independent cavity optics, each with an identical and independent optical coating however with differing beam spot sizes on the ITM and ETM. The contribution of coating noise to the cryogenic cavity length noise is therefore the quadratic sum of the length noise of the two coatings. From this, the coating cryogenic cavity length noise is $1.0 \times 10^{-17}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz at 123 K operation, and $4.0 \times 10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz at 18 K operation. These values are used as the cryogenic cavity length noise requirement throughout this body of work.

The room temperature suspensions contribute noise to the interferometric system as seismic noise couples to optic motion. The room temperature suspensions do not form the cryogenic cavity therefore this is not directly sensed as cryogenic cavity length noise, however motion of the optic introduces frequency noise to the laser beam which is later sensed as length noise in the cryogenic cavity. To provide a design sensitivity requirement for the room temperature suspensions, the cryogenic cavity length noise requirement must be appropriately converted to optic length noise for each suspension. The room temperature optics must be suitable for use during both the 123 K and 18 K phases of the GCIF experiment, therefore the lower noise requirement of 18 K is used as the design requirement in the remainder of this chapter.

3.1.1.1 300g Double Suspension

For light of frequency f_{laser} reflecting in a cavity of length L_{cavity}, any noise imposed by the changing frequency of the laser beam, δf_{laser}, is indistinguishable from a physical change in length of the cavity, δL_{cavity} [171]. This can be expressed as

$$\frac{\delta f_{\text{laser}}}{f_{\text{laser}}} = \frac{\delta L_{\text{cavity}}}{L_{\text{cavity}}}. \tag{3.1}$$

The cryogenic cavity length noise $\delta L_{\text{cryo. cavity}}$, in a noise amplitude spectral density of $\text{m}/\sqrt{\text{Hz}}$, can be converted to stabilised frequency noise δf_{laser} in units of $\text{Hz}/\sqrt{\text{Hz}}$ by multiplication of the frequency of the 1550 nm laser (1.94×10^{14} Hz) and division of the length of the 10 cm cryogenic cavity. Thus in order to measure the length noise of the cryogenic cavity the laser frequency noise must be more stable than the cryogenic cavity length fluctuations. The cryogenic cavity length noise, $\delta L_{\text{cryo cavity}}$, can be used to determine the stabilised frequency noise requirement δf_{laser} using Equation 3.1. Calculating the 10 m reference cavity noise requirement, $\delta L_{\text{ref. cavity}}$, is then a simple matter of applying Equation 3.1 to the 10 m cavity with this frequency noise and rearranging, giving

$$\delta L_{\text{ref. cavity}} = \delta L_{\text{cryo. cavity}} \times \frac{L_{\text{ref. cavity}}}{L_{\text{cryo. cavity}}}. \tag{3.2}$$

Reference cavity length noise is then converted to single optic length noise by division of $\sqrt{2}$ for the inverse of the quadrature summation of independent, identical elements. For the length noise of a single 300 g double suspension this is

$$\delta L_{\text{300g optic}} = \frac{100}{\sqrt{2}} \times \delta L_{\text{cryo. cavity}}. \tag{3.3}$$

The multiplication of 100 arises from simplification of the cavity lengths term where $L_{\mathrm{ref.\,cavity}} = 10\,\mathrm{m}$ and $L_{\mathrm{cryo.\,cavity}} = 0.1\,\mathrm{m}$. Using Equation 3.3 to convert the cryogenic cavity length noise requirement to optic length noise for the 300 g double suspension produces a design sensitivity of $2.8 \times 10^{-16}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz.

3.1.1.2 Auxiliary Suspension

The auxiliary suspension will be used as a beam steering optic between the 10 m reference cavity and the cryostat housing the 10 cm cryogenic cavity, shown in Figure 1.11. A pair of steering mirrors will be used for beam alignment to the cavity, however only the second of these mirrors will be housed in the small vacuum tank situated beside the cryostat. This suspension therefore has a more stringent design requirement than the first steering suspension, needing low optic length noise while maintaining a compact size. The noise introduced by the auxiliary suspension optic motion cannot therefore be a simple conversion between cavity length noise and laser frequency noise. As in Section 3.1.1.1, the cryogenic cavity length noise can be converted to equivalent laser frequency noise requirement as

$$\delta f_{\mathrm{laser}} = \delta L_{\mathrm{cryo.\,cavity}} \times \frac{f_{\mathrm{laser}}}{L_{\mathrm{cryo.\,cavity}}}. \tag{3.4}$$

The instantaneous frequency of an oscillating signal is defined as the rate of change of phase with time, t, such that the laser frequency noise can be expressed as the time-derivative of the phase noise,

$$\delta f_{\mathrm{laser}} = \frac{1}{2\pi}\frac{\mathrm{d}(\delta\phi)}{\mathrm{d}t}, \tag{3.5}$$

Fourier decomposition of a signal allows the phase noise in some frequency bin at a given frequency, f, to be described as an oscillation at that frequency with an amplitude A, such that

$$\delta\phi = A\,\sin\left(2\pi f t\right). \tag{3.6}$$

Substituting this phase oscillation into Equation 3.5 produces the frequency noise in each frequency bin as

$$\delta f_{\mathrm{laser}} = f \times A\,\cos\left(2\pi f t\right). \tag{3.7}$$

The noise spectrum in each frequency bin expresses only the signal amplitude, thus the frequency noise in each bin is $A\,f$, and the phase noise spectrum can be found as

$$\delta\phi = \frac{1}{f} \times \delta f_{\mathrm{laser}}. \tag{3.8}$$

For the conversion of this phase noise into optic length noise, consider a laser beam propagating across a length of vacuum, l, which accumulates a phase of

$$\phi = -\frac{2\,\pi\, f_{\mathrm{laser}}\, l}{c}, \tag{3.9}$$

for c the speed of light in a vacuum [171]. It can be seen that any motion of the beam steering optic along the axis of propagation, $\delta L_{\mathrm{aux.\,optic}}$, modifies the beam path length and introduces phase noise to the beam since $\delta\phi \propto \delta L_{\mathrm{aux.\,optic}}$. This phase noise can therefore be written as

$$\delta\phi = \frac{w \times \delta L_{\mathrm{aux.\,optic}}}{c} = \frac{2\,\pi}{\lambda_{\mathrm{laser}}} \times \delta L_{\mathrm{aux.\,optic}}, \tag{3.10}$$

with laser wavelength λ_{laser}. In reality, the auxiliary suspension sits at a 45° angle to the beam propagation direction, introducing an additional factor of $\sqrt{2}$ to the change in path length from optic motion.

Combining Equations 3.4, 3.8, and 3.10 the conversion from cryogenic cavity length noise to auxiliary suspension optic motion is

$$\delta L_{\mathrm{aux.\,optic}} = \delta L_{\mathrm{cryo.\,cavity}} \times \frac{f_{\mathrm{laser}}}{L_{\mathrm{cryo.\,cavity}}} \times \frac{1}{f} \times \frac{1}{\sqrt{2}}\,\frac{\lambda_{\mathrm{laser}}}{2\,\pi}. \tag{3.11}$$

Converting the cryogenic cavity length noise requirement into the limit on auxiliary suspension optic motion produces $1.3 \times 10^{-10}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz.

3.2 Modelling of Existing Suspensions

Several double-stage suspensions were designed for use as cavity optics and beam-steering optics in past projects in the Glasgow 10 m laboratory. Re-purposing these suspensions for use in the facility's upgrade is cheaper, faster, and more efficient than custom-building new designs. The existing suspensions were characterised and modifications were made to ensure the existing suspensions met the noise requirements for the planned GCIF thermal noise experiment.

The suspensions chosen for re-use as the 10 m cavity optics were designed by N. Gordon and R. Jones [88]. They have a total pendulum length of 60 cm with an upper aluminium mass of 290 g and suspended optic of 300 g. The suspension sits on an aluminium base plate to which four aluminium struts are bolted. The struts support the suspension top

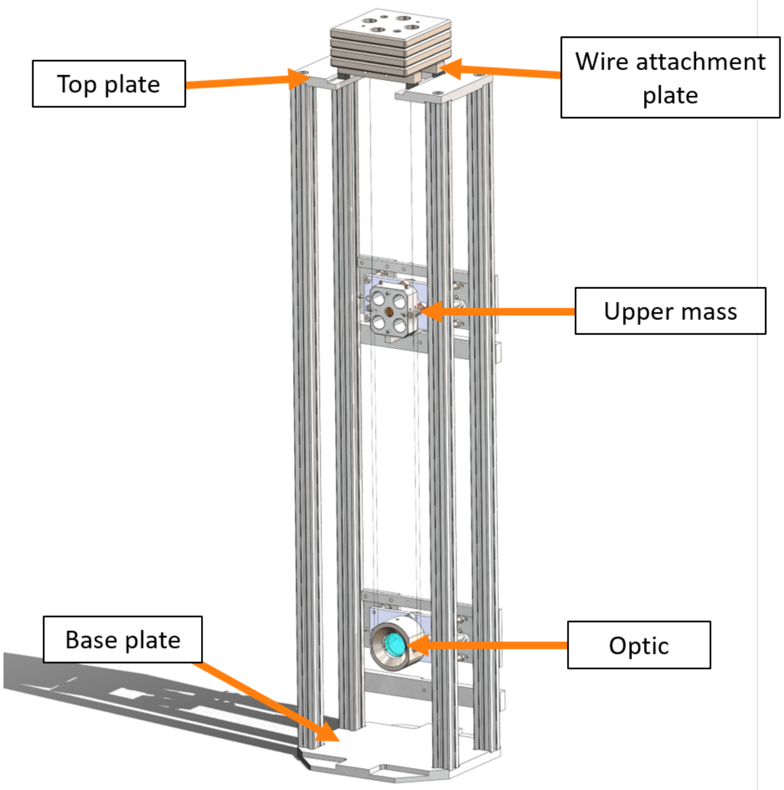

Figure 3.1: Computer-aided design image of the 300 g suspension isometric view, showing the base, struts, and top plate of the suspension support structure. The wire attachment plate sits on four rubber blocks, providing an attachment point for the two wires attached to the upper mass. The optic is suspended from the upper mass by four wires.

plate and provide an attachment point for the coil-magnet actuators used for alignment and damping of the upper mass. The wire attachment plate is separated from the top plate by four rubber blocks to provide additional seismic attenuation to the masses. This suspension is referred to as the 300g suspension from hereon, and is shown in Figure 3.1.

A small double-stage beam steering suspension was designed for use in a previous GCIF experiment by J. Hennig and R. Jones [103]. This is referred to as the auxiliary suspension in [103] and this body of work. The auxiliary suspension test mass comprises of an aluminium ring optic holder and 30 mm mirror optic for a total test mass of 75.9 g. A 75.4 g aluminium upper mass holds four coil-magnet actuators for alignment and damping of the longitudinal, yaw, and pitch modes. 50 µm diameter steel wires suspend the masses with two 100 mm long wires to the upper mass and four 150 mm long wires to the optic. The back-plate's base features a small footprint of $< 50 \times 80$ mm which allows multiple steering mirrors to be used within each vacuum tank. An image of the auxiliary suspension can be seen in Figure 3.2 with a full description of the design found in [103].

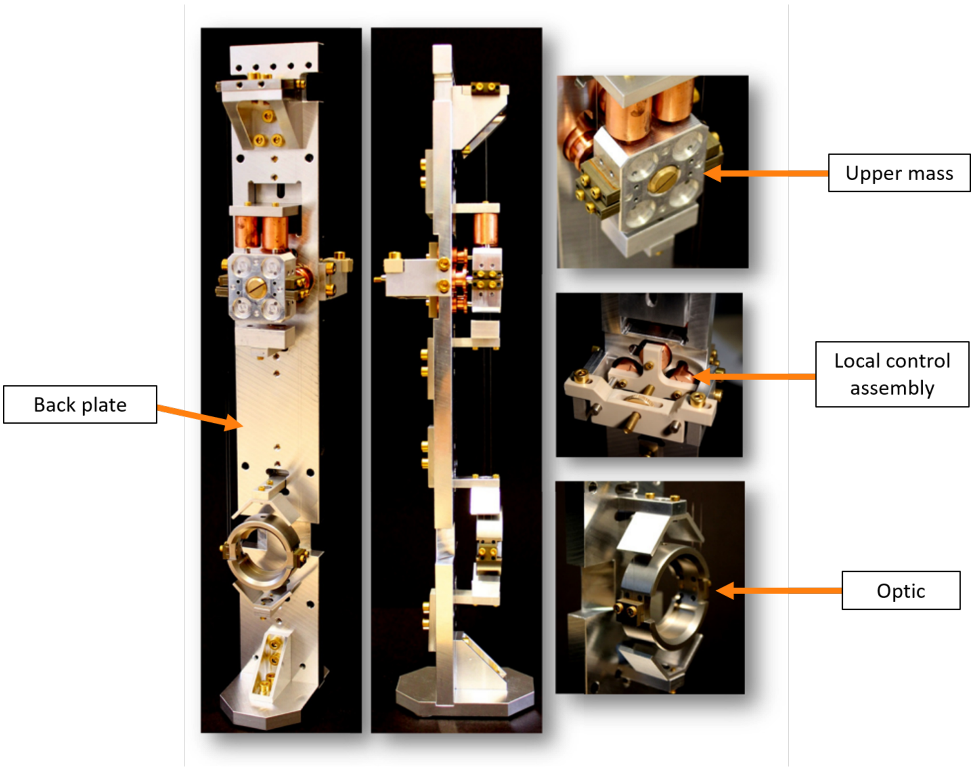

Figure 3.2: L-R: Photographs of the auxiliary suspension isometric and side view. The upper mass, local control assembly and optic are shown from top to bottom on the right of the image. Image from [103].

3.2.1 Creating the Mathematica Models

The first computer model of triple-stage suspended optics in the gravitational wave community was developed in MATLAB by C. Torrie using the analytical approach outlined in [193]. This was a detailed model of suspension dynamics which calculated resonant modes and transfer functions of the system, although the by-hand approach necessitated some simplifying assumptions. Torrie modelled the masses as rigid bodies connected by massless wires with longitudinal elasticity. For simplicity the system is assumed to be symmetric and therefore any cross couplings between the degrees of freedom caused by asymmetries are ignored: only the independent vertical, longitudinal-pitch, transverse-roll, and yaw motions are considered. Additionally, this model accounts for blade springs by adding elasticity to the wires rather than through an independently modelled element [25]. This model was later implemented into Mathematica and updated by M. Barton. Mathematica is a technical computing software which uses a symbolic algebra system powerful enough to calculate the cross-couplings of an asymmetric model - a real possibility for a physical suspension build if, for example, there are slight differences between wire lengths [212]. The updated model also includes the bending elasticity and stretching of the wire and uses separate rigid body elements for the blade spring tips. A full comparison between the MATLAB and Mathematica models can be found in [25].

These Mathematica suspension models are currently in use by the LIGO scientific collaboration to provide resonant frequencies, transfer functions, and thermal noise estimates for most suspensions used at the aLIGO interferometer [24]. An existing Mathematica model for a double pendulum suspension, [23], was used as the basis of the suspension models in this body of work and was adapted for use with the GCIF suspensions by the author with the help of M. Barton. Separate models were made for the auxiliary beam-steering double suspension and for the 10 m cavity 300 g double suspension. The input parameters of [23] were updated to reflect the mass, shape, moments of inertia, and material of the upper mass and optic, the length, radius, and material of the wires, and the position of the wire attachment points on the masses for both suspension models. The Mathematica modelling considers the suspended stages from the upper wire mounting point to the optic, however does not currently model an element for the suspension support structure. This must be taken into account when comparing a model output to measurements of a physical suspension.

3.2.2 Verification of the Model by Comparison of Resonant Modes

Suspension Mode Type	**Direct Analytical Calculation** [Hz]	**Mathematica Model Output** [Hz]	**Simulated Output from [103]** [Hz]
x-longitudinal	1.05, 2.45	1.07, 2.48	1.05, 2.68
y-longitudinal	1.07, 2.68	1.08, 2.68	1.08, 2.68
yaw	1.70, 4.08	1.70, 4.29	1.70, 4.28
pitch	4.27, 15.30	3.25, 15.22	4.35, 15.8
vertical	31.00, 90.30	31.65, 90.03	32.0, 90.3
roll	38.21, 123.12	37.82, 123.09	38.2, 123.3

Table 3.1: Resonant mode frequencies for the auxiliary suspension, calculated through a direct analytical approach; through a Mathematica model of the masses and wires; and simulated suspension outputs taken from [103].

The first resonant modes of each degree of freedom were output from the auxiliary suspension Mathematica model and compared to those calculated via a direct analytical approach, and to a computer model created by Hennig during the original build of the auxiliary suspension [103]. These show the resonances of relative motion between the two suspended masses. Hennig later compared his mode outputs to measurements of the resonant modes as will be done in Section 3.4. This comparison verified the accuracy of the modelling before any suspension input parameters were adjusted. The direct analytical approach used the methods outlined in Chapter 2 and [193], with the results in column two of Table 3.1. Similarly, the output modes for the 300 g double suspension were compared to measurements taken by Gordon during the initial suspension build. Gordon's experiment involved the excitation of suspensions resonances using an electromagnetic actuator and measuring the motion of the mass using a laser vibrometer, with the results shown in column two of Table 3.2 and [88].

The auxiliary suspension resonant modes agree well between the models. The Mathematica model modes match Hennig's simulation modes to within 0.5 Hz for all resonances aside from the two pitch modes, and match the direct analytical calculations to within 0.5 Hz for all aside from the lower-frequency pitch and vertical modes. The lower pitch mode of the Mathematica model differs from the other two calculation methods by ~1.1 Hz, likely caused by a slight difference in the vertical position of the wire attachment points between the models. The other output mode discrepancies are less than 1 Hz and could be attributed to slight differences in the sourced material properties. Overall this comparison demonstrates good agreement between the three methods and validates the accuracy of the Mathematica double suspension model. This allows the suspension parameters of the model to be adjusted while producing believable outputs, allowing different suspension designs to be explored for use in the GCIF.

A similar check was performed for the 300 g double suspension as seen in Table 3.2. The y-longitudinal motion was not measured for the physical suspension, however it is expected that these modes would approximately match the x-longitudinal modes. The missing yaw measurement is due to a vibrometer misalignment [88]. The Mathematica modelled resonances for the 300 g suspension show greater differences to the published modes than seen for the auxiliary suspension. This is expected as the auxiliary suspension published resonances are for an modelled ideal suspension, whereas the 300 g published results were measured on a physical suspension. Accounting for machining tolerances and build quality, the physical suspension resonances are expected to deviate slightly compared to the model. Comparison of results show that most modelled resonances match the experimental measurements to within 1 Hz, displaying a good agreement between the

Suspension Mode Type	**Mathematica Model Output** [Hz]	**Measured Modes from [88]** [Hz]
x-longitudinal	0.70, 1.73	0.69, 1.71
y-longitudinal	0.70, 1.74	
yaw	1.15, 3.12	1.07
pitch	1.99, 12.55	2.93, 12.20
vertical	19.86, 55.50	19.00, 54.70
roll	30.63, 80.46	29.8, 89.20

Table 3.2: Resonant mode frequencies for the 300 g double suspension output from a Mathematica model of the masses and wires, and measured on a physical suspension [88].

methods. The only exception to this is the second order roll mode which shows a difference of 8.74 Hz. Small variations to the moment of inertia of the masses, wire radius, wire length, and wire attachment points were made to investigate possible differences between the physical suspension and the model, however no configuration was found which resulted in a closer match for the second order roll mode while maintaining a match between the other modes. As discussed in Chapter 2 the roll mode is not expected to couple strongly to the longitudinal optic motion therefore this discrepancy was not a large concern in this study.

3.2.3 Optic Length Noise

Transfer functions were calculated in the Mathematica model showing the effect of motion at the top wire mounting point to motion of the optic. The transfer functions of interest when considering length noise are longitudinal suspension base motion to longitudinal optic motion, vertical suspension base motion to vertical optic motion, and longitudinal suspension base motion to optic pitch motion. These show the strongest coupling to longitudinal optic motion, and therefore have the strongest effect on cryogenic cavity length noise. The transfer functions of upper wire mounting point motion to optic motion were calculated using the Mathematica model for both suspensions with the example of the 300 g double suspension shown in Figure 3.3. For the auxiliary suspension this is assumed to be the same as suspension base motion to optic motion. The 300 g suspension contains a layer of rubber between the suspension base and the upper wire mounting point which is considered in the following section.

The suspension optic motion expected during experimental operation can be predicted by multiplication of the modelled transfer functions by the expected suspension base seismic noise. Figure 3.5 shows seismic data provided by the GCIF cryostat manufacturer, Leiden Cryogenics, for typical noise levels found inside their cryostat inner shields, represented in brown and orange. This data includes cryocooler noise and so provides a worst-case estimate of laboratory ground motion in both horizontal and vertical directions between 1 – 1500 Hz.

The room temperature suspensions will be placed upon optical breadboards inside large vacuum chambers which are shown in the lab layout diagram, Figure 1.11. These optical breadboards will sit upon a passive isolation stack to reduce the seismic motion coupling into the optical components. To estimate the seismic motion at the base of the suspensions, the attenuation of the ground motion through these passive isolation stacks must first be considered.

3.2.3.1 Additional Passive Isolation

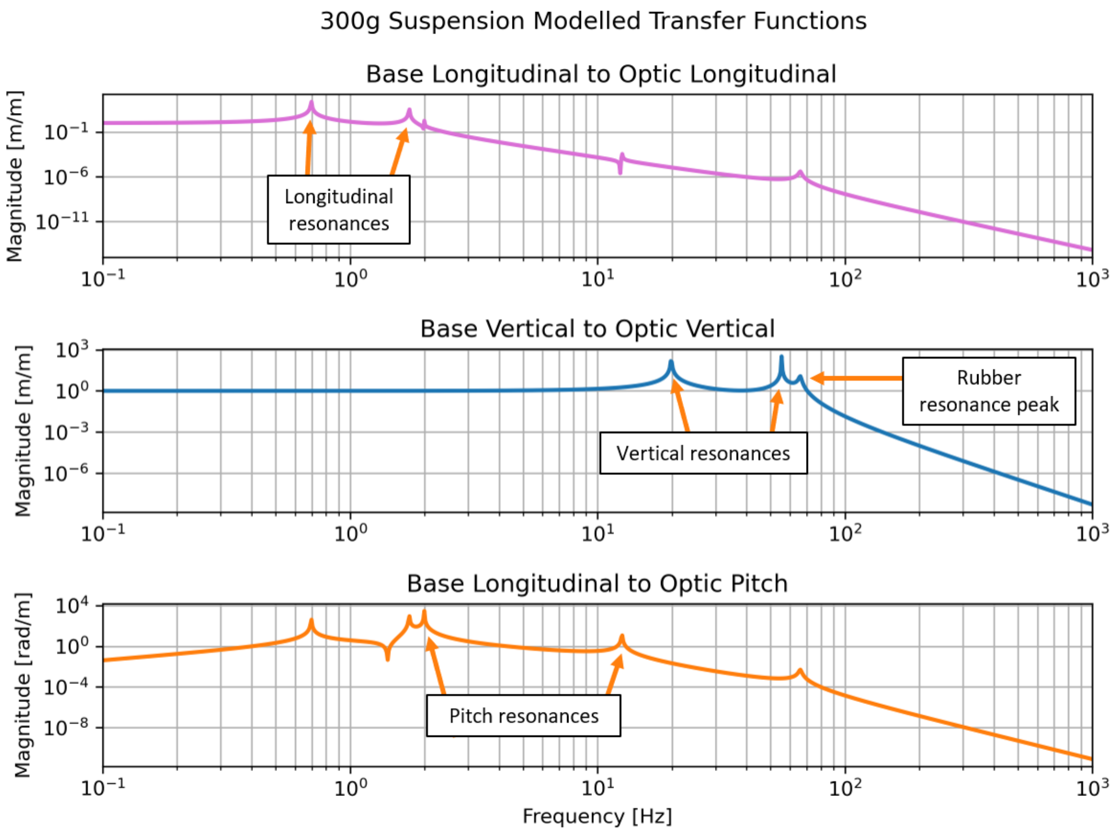

Figure 3.3: The magnitude output of modelled transfer functions for the 300 g double suspension. The transfer functions are between motion at the suspension base to optic motion and therefore does not include the effects of the passive isolation stacks. The top plot is longitudinal base motion to longitudinal optic motion, shown in pink. The middle plot is vertical base motion to vertical optic motion, shown in blue. The lower plot shows longitudinal base motion to optic pitch motion, shown in orange. The suspension resonances are labelled, with some coupling observed between longitudinal and pitch motions.

Passive isolation stacks exist for the 1 m diameter vacuum chambers from previous experiments conducted in the laboratory. These isolation stacks consist of two alternating layers of rings of aluminium extrusion and rubber blocks. Measurements in [88] provide data on the resonant frequency of these rubber blocks varying with load, then approximates the trend between resonant frequency and mass load to predict the resonant mode of the isolation stacks with the weight of the experimental payload for a previous GCIF experiment. The isolation stacks will be re-used in the upgrade of the 10 m laboratory and so the resulting resonances from [88] are used to predict the seismic attenuation of the stacks in this work, where 20.5 Hz and 24 Hz resonances were chosen to match the stack configuration and mass load expected for the GCIF. During a 2016 experiment an additional layer of rubber blocks was installed at the feet of the vacuum chambers with an estimated resonance frequency of 65 Hz, which was included in the isolation estimation.

Equation 2.1 was used to model the transfer function of the isolation stack against frequency, with unity before the resonance, a peak at the resonance frequency, and a $1/f^2$ drop-off above the resonance for each layer of the stack. The rubber at the 300 g suspension mounting point was also accounted for this way. Measurements of the suspension transfer function (discussed in Section 3.4) showed the resonance of this rubber to be 66 Hz. To check the accuracy of this prediction it is recommended that a measurement of seismic displacement is taken on the optical platforms and compared to the seismic displacement measured on the laboratory floor. Additionally, the predicted resonant frequencies of the stack could be re-calculated once the true loading of the breadboards is known for the GCIF experiment. The suspension base to optic transfer functions of interest are shown as an example for the 300 g suspension in Figure 3.3 which includes the passive isolation from the mounting point rubber blocks.

The smaller vacuum tank for the beam steering suspension is a new addition to the laboratory and does not yet contain a passive isolation stack. The seismic isolation performance of a stack suitable for use in this small tank was estimated through measurements of a loaded breadboard sitting on rubber blocks. The experimental set-up is shown in Figure 3.4 (a). Three rubber cylinders of a 30 mm diameter and 25 mm height were placed in an equilateral triangle configuration underneath a Thorlabs 0.3 m $\times$ 0.3 m optical breadboard, such as would fit within the small tank. Two Brüel & Kjaer Type 4379 accelerometers were used to measure vertical acceleration. One accelerometer was placed on the laboratory bench and one on the isolation stack. Each accelerometer was connected to a pre-amplifier then to one of the two input channels of a Stanford Instruments SR785 spectrum analyser. The spectrum analyser measured the two inputs simultaneously and output the ratio between the two channels. The laboratory bench was gently excited by knocking of the beam tube under the bench such that the excitation was visible above the seismic background noise. Measurements of the ratio and coherence between the two accelerometers were taken using the signal analyser for a frequency range of 0 – 400 Hz. This technique was repeated as the load was increased from 11.8 kg, to 16.5 kg, to 19.5 kg

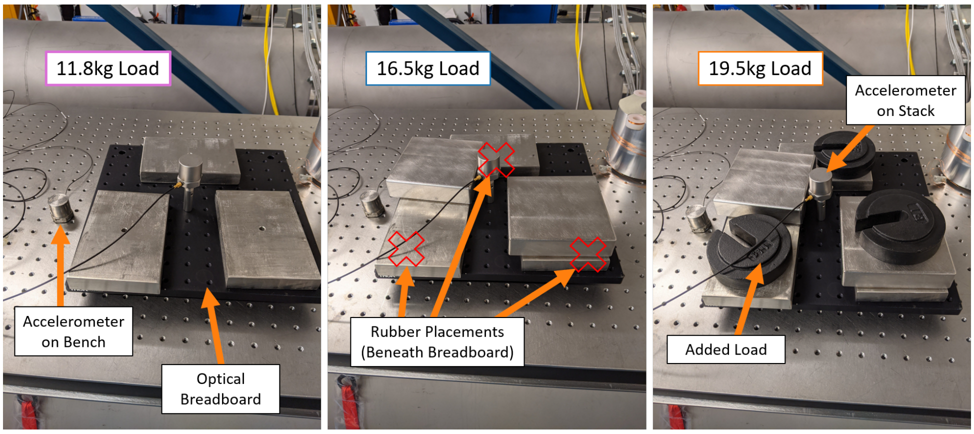

(a) Photograph of experimental set-up.

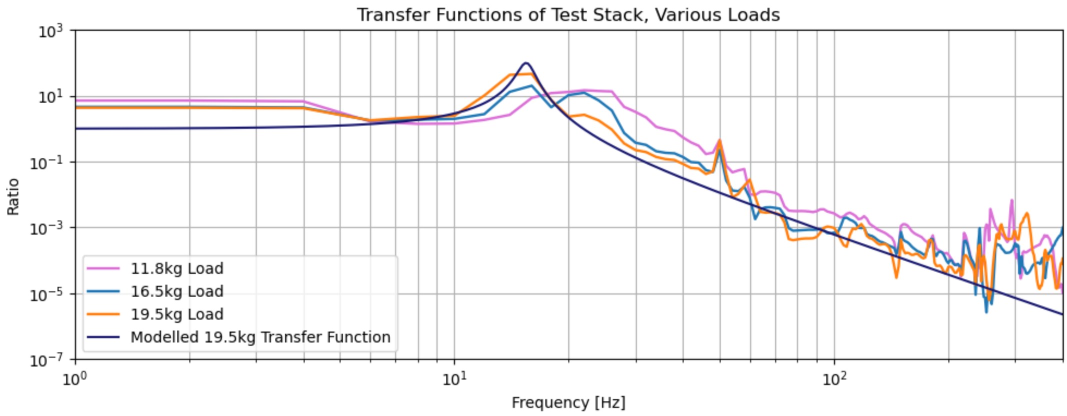

(b) Transfer function outputs as measured between the two accelerometers.

Figure 3.4: (a) Photograph of experimental set-up for the measurement of isolation stack transfer functions. The three load scenarios are shown from increasing load L-R. The two accelerometer placements are visible. The rubber is sitting underneath the breadboard with placements represented by red crosses. (b) Transfer function outputs as measured between the two accelerometers shown for the 11.8 kg load in pink, the 16.5 kg load in blue, and the 19.5 kg load in orange. The modelled transfer function based on the 19.5 kg load is shown in navy.

(with load including breadboard mass). The observed resonant frequency of the stack decreased as the load was increased which is consistent with expectations from the vertical resonant frequency equation (Equation 2.8). The measured transfer functions of the prototype stack are shown in Figure 3.4 (b).

These measurements were used to provide an estimation of the resonant frequencies of loaded isolation stacks and were not intended to be taken as the final transfer functions for small tank isolation. In place of using the measured transfer function, which runs into instrumental noise at about 200 Hz, Equation 2.1 was used to model the isolation up to 1500 Hz to match the frequency range of the seismic data. The modelled transfer function of the 19.5 kg loaded stack is shown in Figure 3.4 (b) and has a 15.5 Hz resonant frequency and $1/f^2$ drop-off above the resonance.

3.2.3.2 Calculation of Optic Length Noise

The seismic motion is multiplied by the isolation stack transfer functions to estimate the noise reaching the base of the suspensions. The longitudinal isolation of the passive stacks was assumed to be similar to the vertical performance [83, 103]. This is shown in Figure 3.5 in blue and cyan for the auxiliary suspension and light pink and purple for the 300 g double suspension, for both vertical and horizontal directions.

The horizontal seismic motion at the suspension base is multiplied by the longitudinal to longitudinal suspension transfer function. This outputs the longitudinal optic motion caused by horizontal seismic noise.

The vertical seismic motion at the suspension base is multiplied by the vertical to vertical suspension transfer function, outputting the vertical optic motion caused by vertical seismic noise. A 0.1% coupling is assumed from optic vertical motion to optic longitudinal motion, as detailed in Section 2.2.2.

Similarly, the horizontal seismic motion at the suspension base is multiplied by the longitudinal to pitch suspension transfer function to output the optic pitch caused by horizontal seismic noise. As mentioned in 2.2.4 the pitch-to-length coupling factor can be calculated through trigonometry and is taken as a maximum of 4% for a 0.5 mm offset on the 25 mm diameter optics of the double suspensions.

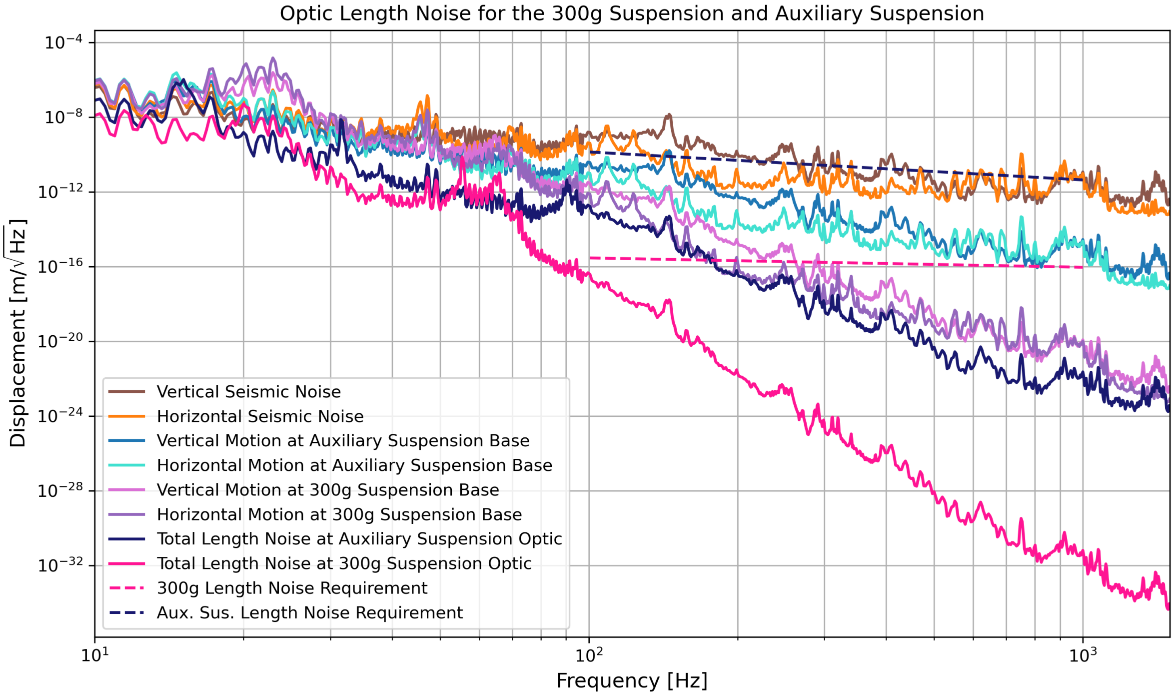

Figure 3.5: Modelled optic length noise of the unmodified auxiliary and 300 g suspensions, shown in comparison to seismic noise and length noise requirements. The horizontal and vertical seismic data is shown in orange and brown. The seismic noise attenuated by passive isolation before the suspension base is shown in cyan and blue for the auxiliary suspension, and purple and pink for the 300 g suspension. The total optic length noise is plotted in dark blue for the auxiliary suspension and in dark pink for the 300 g. The plotted frequency range is 10 – 1500 Hz to increase visibility of the experimental sensitivity region.

The total length noise of the optic is calculated as the combination of the optic-longitudinal, optic-vertical, and optic-pitch sources of longitudinal optic motion in quadrature. This is shown in navy in Figure 3.5 for the auxiliary suspension, and in dark pink for the 300 g double suspension. Also represented by dashed lines of the same colour scheme are the optic length noise requirements which were calculated in Section 3.1.1.

The optic length noise of both suspensions sits below the length noise requirement for most of the 100 Hz – 1 kHz experimental sensitivity range, with the 300 g suspension only just below the requirement at 100 Hz. The total length noise within this sensitivity range has an isolation from seismic noise which varies with $1/f^6$ for the auxiliary suspension and $1/f^{12}$ for the 300 g. The length noise of both suspensions therefore decrease much faster with frequency compared to the noise requirement, which has a $1/\sqrt{f}$ dependence for the 300 g suspension and $1/f^{1.5}$ for the auxiliary suspension. At 1 kHz the 300 g suspension length noise is over 15 orders of magnitude below the requirement, with the auxiliary suspension length noise over 10 orders of magnitude below its requirement. This gives an extremely comfortable margin at this higher frequency, but there is a much smaller margin at 100 Hz: the 300 g noise lies just 7.6× lower than the noise requirement. The auxiliary suspension is a more comfortable 825.8× lower than the requirement at 100 Hz. Reduction of this length noise to well below the sensitivity requirement is critical for allowing a direct measurement of coating thermal noise. A secondary motivation for length noise reduction is suspension controllability at low frequencies. If the suspension motion is higher than the control capabilities of the sensors or actuators then the control bandwidth of the 10 m cavity will be limited, impacting the laser stability and coupling additional noise into the experiment. The frequency-dependent drop off is assumed by the model to continue up to 1500 Hz, however it can be seen in Figure 3.5 that this drop-off is only required up to a few hundred hertz as above this, the optic length noise lies many times below the sensitivity requirement. Modifications to the room temperature suspension designs are outlined in Section 3.3 which aim to give a larger margin to the isolation of both suspensions around the 100 Hz region, with an emphasis on the isolation of the 300 g suspension.

3.2.4 Force-Displacement Transfer Functions

The seismic base motion to optic motion transfer functions were used to estimate the optic length noise from input seismic noise, however the $1/f^4$ drop from base motion to optic motion means these transfer functions are difficult to measure directly in an experimental set-up. Force-displacement transfer functions were instead used to allow comparison of the physical 300 g double suspension to the Mathematica model. These describe the optic longitudinal motion resulting from an input 1 Nm torque on the upper mass, where driving of the upper mass results in a lower $1/f^2$ translation to optic motion. The transfer functions of primary interest are those which can be directly measured on the suspension: the longitudinal-longitudinal, the pitch-longitudinal, and the yaw-longitudinal. In this work these are referred to by convention as force-displacement transfer functions, but have units of [m/m] for the longitudinal and [m/Nm] for the angular functions.

The four coil-magnet actuators placed at the 300 g suspension's upper mass allow easy drive of the upper mass longitudinal, pitch, and yaw motions. These can be seen in Figure 3.17. The coil-magnet spacing was measured to allow the driving force to be simulated in the model as an input force vector applied at each actuator position. The input force vector can be applied in the positive x-direction to simulate longitudinal motion of the upper mass. Pitch or yaw motion can be simulated by applying force in opposing x-directions to the top and bottom or left and right respective coil pairs.

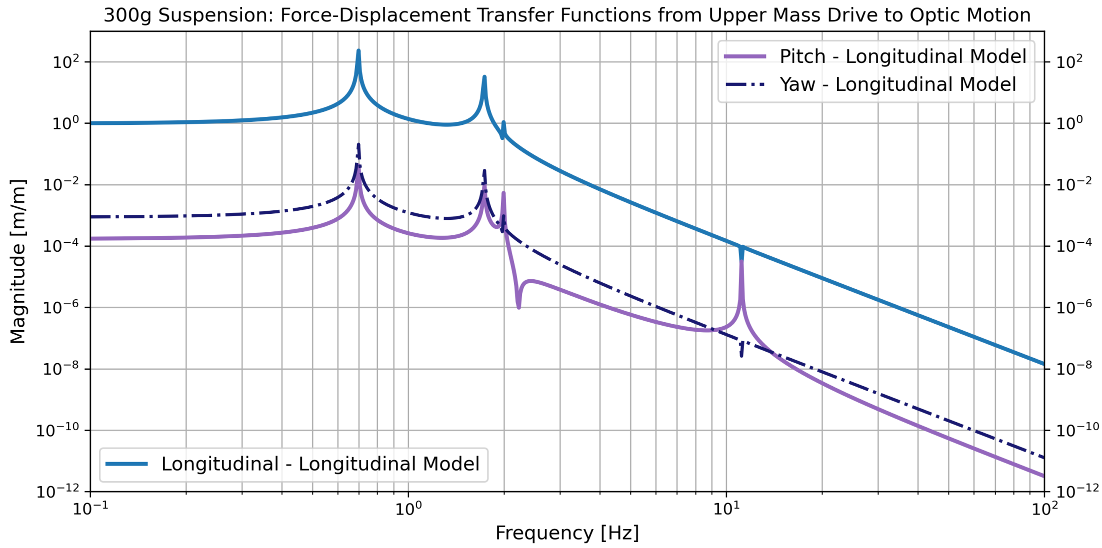

Figure 3.6: Modelled force-displacement transfer functions for upper mass drive to optic motion of the 300 g double suspension. Upper mass longitudinal drive to optic longitudinal motion is shown in blue with the units of the left axis. The upper mass pitch and yaw motions to optic longitudinal motion are shown in purple and dash-dotted dark blue respectively, and correspond to the right axis units.

The outputs of the simulation for the 300 g suspension are shown in Figure 3.6. The peaks of the resonant modes are the same as those shown in column two of Table 3.2, noting that this Figure displays longitudinal optic motion. These modelled outputs are later compared to measured data as part of the suspension characterisation in Section 3.5.1.2.

3.3 Suspension Modifications

The predicted optic length noise is compared to length noise requirements for the 300 g double suspension and the auxiliary suspension in Figure 3.5. The optic length noises lie well below the noise requirement for the majority of the experimental sensitivity band, however in the 100 − 150 Hz region the optic length noise could be decreased to give an extra margin to the isolation of both suspensions. This is essential for the 300 g suspension which has a length noise 7.6 × below the requirement at 100 Hz: an additional 10× isolation at 100 Hz would be ideal. Simulations are presented for potential modifications for the auxiliary suspensions should they be required in future GCIF projects, producing a suggested design which is able to reduce optic length noise while maintaining the compact size and existing suspension parts. Due to the greater need for improved 300 g suspension performance the focus of this work was centred around the mechanical design and improvements for the reference cavity suspensions. For the 300 g suspension the detailed design and modification process is described. This includes the characterisation of existing blade springs, details of additional parts manufactured for the upgrade, and the re-hanging process of the modified suspension.

The modifications to the 300 g double suspension are outlined in Section 3.3.1, and the auxiliary suspension modifications in Section 3.3.2. For both suspensions the conceptual design is described, and the suspension base motion to optic motion transfer functions are modelled. These are later used to predict the improvement to optic length noise for both modified suspensions.

3.3.1 Modification of the 300g Double Suspension

A closer examination of the longitudinal, vertical, and pitch motion contributions to the 300 g optic total length noise are shown in Figure 3.7. This plot shows the length noise contributions which are described in Section 3.2.3.2, created by the multiplication of seismic noise, suspension transfer functions, and any coupling factor to longitudinal motion. These contributions are combined in quadrature for the 300 g total optic length noise represented in dotted dark pink. Comparison of the contributions plotted in Figure 3.7 show that the vertical contribution is dominant in the entire experimental sensitivity range. Based on this, it was decided to decrease the suspension optic length noise through the addition of vertical isolation to the suspension.

3.3.1.1 Conceptual Design

In the re-design of the 300 g suspension two steel blade springs were added to the top wire mounting point. Blade springs have a lower spring constant than the upper stage suspension wires alone, reducing the resonance frequency of the vertical mode. This therefore improves the vertical isolation of the suspension at 100 Hz by shifting the beginning of the $^1/_{f^4}$ slope to a lower frequency.

Blade springs used in a suspension chain must have an uncoupled resonant frequency, f_s, well below the experimental frequency range of interest, have a breaking stress above that induced by the weight of the suspension chain, and be of a suitable size to fit within the design of the suspension chain. The uncoupled frequency of a cantilever blade is calculated

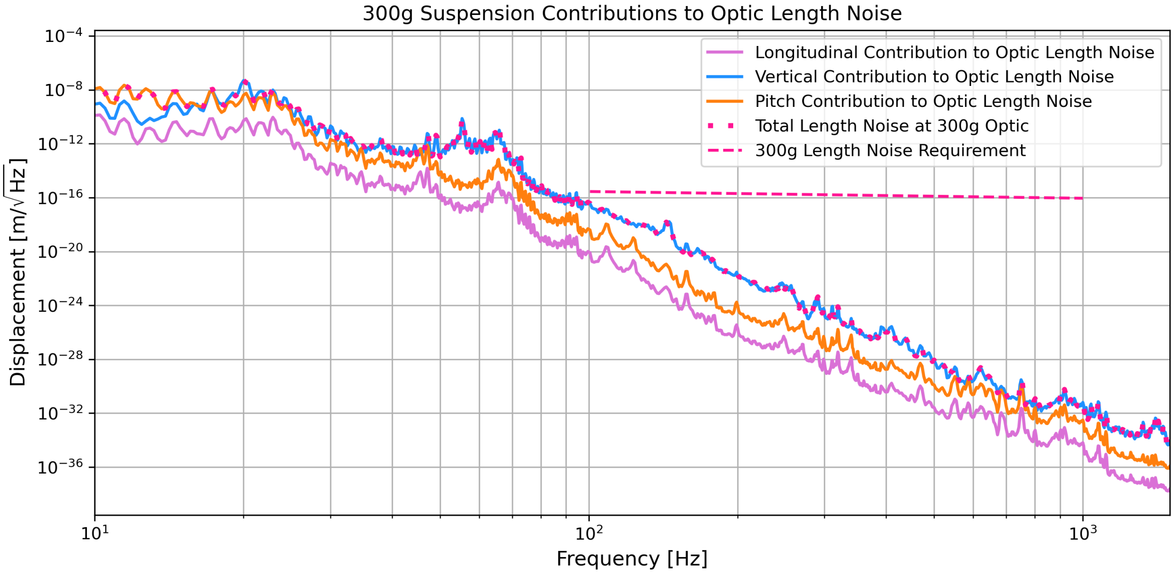

Figure 3.7: 300g suspension contributions to optic length noise. The longitudinal contribution is shown in pink, the vertical contribution in blue, and the pitch contribution in orange. The optic total length noise is represented by a dotted dark pink line to match the previous colour convention while maintaining visibility of the vertical contribution, and the optic length noise requirement is represented by the dashed dark pink.

by modelling the blade as a spring with spring constant, k_{s}. This spring constant is equal to the applied load over the blade vertical deflection, with this deflection dependent on the chosen geometry and the material's elasticity.

$$k_{\mathrm{s}} = \frac{Y\,w\,t^3}{6\,l^3} \tag{3.12}$$

for material Young's modulus Y, blade width at clamping end w, blade thickness t, and blade length l as defined in Figure 3.8 [154]. From this the blade resonant frequency in hertz is calculated from simple harmonic motion as

$$f_{\mathrm{s}} = \frac{1}{2\,\pi}\sqrt{\frac{k_{\mathrm{s}}}{m_{\mathrm{load}}}}, \tag{3.13}$$

where m_{load} is the mass of the applied load in kg [217].

Blade springs existing from a previous 10 m laboratory experiment were examined for suitability of use at the 300 g mounting point. These stainless steel blade springs measured 125 mm in length with a width of 18 mm tapering to 5 mm and a 1 mm thickness, and were therefore of a suitable size for installation on the 300 g suspension. As an approach, the re-use of existing suitable materials is efficient, cheaper, and more environmentally friendly than the manufacture of new materials. Inputting the blade spring geometry into Equations 3.12 and 3.13 produced an expected resonant mode of $f_{\mathrm{s}} = 5.68$ Hz for a 300 g

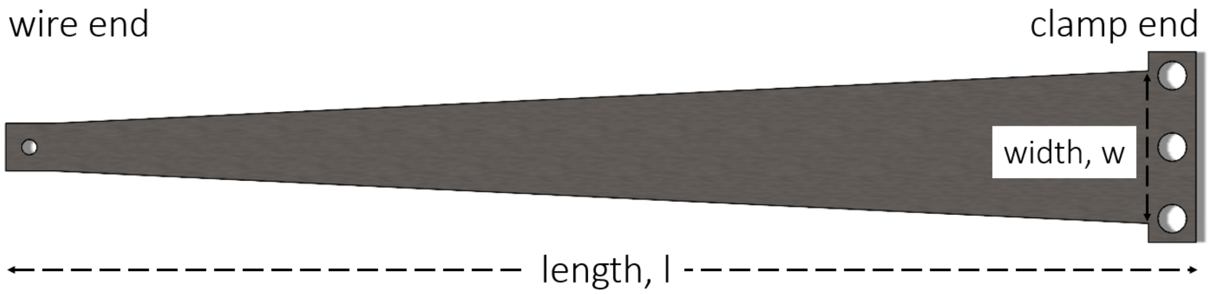

Figure 3.8: Diagram of a blade spring showing the wire and clamp ends, and the width and length dimensions.

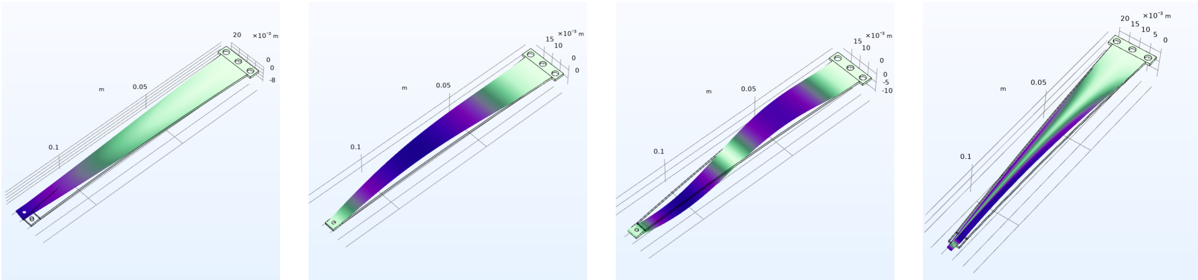

(a) 62.7 Hz transverse mode.

(b) 239.7 Hz first-order xylophone mode.

(c) 843.5 Hz second-order xylophone mode.

(d) 1182.5 Hz torsional mode.

Figure 3.9: FEA predicted blade spring resonances of the higher order modes <1.5 kHz showing (a) the transverse mode, (b) and (c) the first and second order xylophone modes, and (d) the torsional mode.

load per blade spring. A quick estimation with Equation 2.3 shows that reduction of the upper stage vertical mode from 19.76 Hz to 5.68 Hz provides roughly a 12 $\times$ improvement in vertical isolation at 100 Hz, meeting the order of magnitude requirement. These existing blade springs were therefore used in the modelling stage for the suspension modification.

3.3.1.2 Modelling the Suspension Modifications

The modified suspension was modelled to ensure that the addition of blade springs meets the desired reduction of optic length noise. Firstly, finite element analysis (FEA) was performed on the blade springs to estimate the blade spring resonant mode and deflection under load. The resonance mode output is then compared to the result from Equation 3.13 for verification of the FEA model, while an accurate value for the loaded blade deflection is necessary for the mechanical redesign of the suspension to maintain the correct optic height during the re-build. The blade springs were added to the Mathematica dynamic model to output transfer functions for the modified suspension, which were used to estimate the optic length noise of the redesign. Further computer-aided design (CAD) modelling of the mechanical design was undertaken to show the dimensions of the new suspension design and to facilitate the manufacture of necessary parts, as described in Section 3.3.1.4.

The Solidworks program was used to create a CAD model of the steel blades with the appropriate dimensions. This model was imported into COMSOL for finite element analysis within the Solid Mechanics interface, with a Fixed Constraint allowing simulation of the base-side clamping, and a mass per unit length added at the wire end to represent the 300 g load. The body material was set as AISI 4030 steel to match the existing blade springs. The convergence of outputs with varying element sizes was checked to ensure an accurate resonance mode estimation, after which the frequency of the fundamental blade spring resonance was output as 5.79 Hz. This differs by $\leq 2\%$ to the result of Equation 3.13 and can be taken as a verification of the COMSOL model. A small variation in results is expected due to small differences in the two approaches, such as the FEA inclusion of the 20×5 mm clamp end, and the inclusion of bending point variations due to clamping which the equation cannot account for. The FEA displayed the weighted blade to have a deflection of 7.8 mm, therefore needing a launch angle of $7°$ to have a flat tip once weighted.

The FEA model was also used to explore the higher frequency resonances of the blade spring which may appear in the sensitivity range ≤ 1 kHz. A transverse mode (lateral movement of the blade tip in the x-direction) is present at 62.7 Hz, a first-order xylophone mode at 239.7 Hz, and a second-order xylophone mode at 843.5 Hz. There is a blade spring torsional mode (rotation of the blade tip around the y-axis) present at 1182.5 Hz, with the other higher order resonant modes having frequencies above 1.5 kHz and not expected to impact sensitivity. These mode shapes are shown in Figure 3.9, and compared to blade spring measurements in Section 3.5.1.3.

The Mathematica suspensions framework has been rigorously developed to be capable of accounting for various suspension concepts such as the elasticity of wires of differing geometries alongside any mass-wire attachment points, the coupling of motion between numerous suspended stages, and the inclusion of vertical blade springs. The existing 300 g suspension model was updated to include the effect of blade springs at the attachment points of the two top wires, with the 5.47 Hz blade spring resonance used to calculate the net vertical elasticity of the upper stage blades in the model. The upper wire length was increased to account for the new attachment point of the upper wires which takes into consideration the deflection of the blade. The modified suspension resonant modes and transfer functions were then calculated as in Section 3.2. The results of this modelling work are shown in detail in Section 3.5, and predict the necessary improvement to optic length noise. This conceptual design was therefore carried forward to detailed design and rebuild phases.

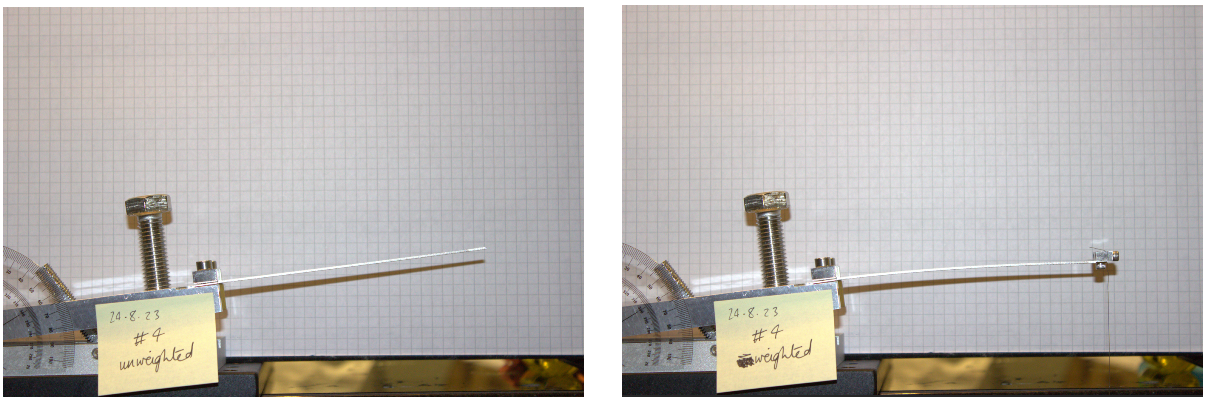

(a) Unweighted blade with a 7° launch angle.

(b) Blade spring with a 300 g load at the tip.

Figure 3.10: Photographs of the unweighted and weighted blade springs which were used for curvature and tip gradient analysis.

3.3.1.3 Characterisation of the Blade Springs

Twelve laser cut steel blade springs of the above dimensions were available for reuse in the GCIF. These blade springs had previously been used to test the validity of using laser-cut blade springs by measuring deflection under high loading. The blade springs were loaded with masses between 0.5 kg – 5.3 kg, with some blades loaded past the blade yield point where there will be plastic deformation. As such, all twelve blade springs were characterised to ensure the blades chosen were of suitable quality for use in the suspension upgrade. The blade springs were firstly visually inspected for surface damage and deformation, with two of the unweighted blades showing a high curved deformation from their previous weighting and a further six showing slight signs of curved deformation. Four of the blades appeared flat and it was predicted that these would most closely match the desired loaded curvature. No surface damage was visible on any blade.

Each blade spring was then clamped to a 7° launch platform and photographed unweighted, then photographed weighted with a 300 g load with the set-up shown in Figure 3.10. The high-resolution images were used to calculate the curvature profile of the loaded blades by marking points along the blade surface and fitting a second-degree polynomial along the curve. This curve fitting was also performed on the FEA model of the weighted blade to provide a comparison between the predicted curvature and the weighted blades. Additionally, points on the images were used to estimate the gradient of the tip of each blade spring, which must lie horizontal to avoid unequal tensioning on the upper mass wire.

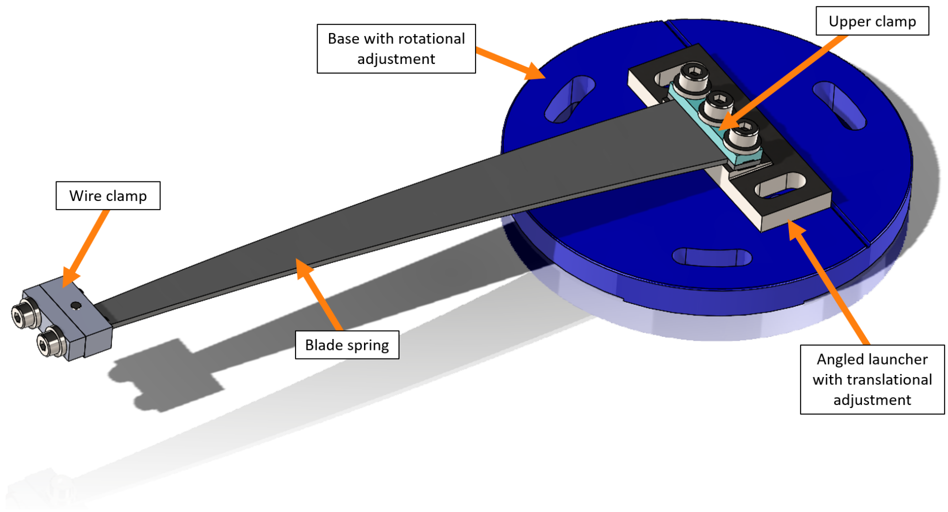

Figure 3.11: Computer-aided design of the blade spring clamp assembly showing the aluminium base piece in blue, the steel 7° angled launcher and the steel upper clamp. The blade spring is shown in the context of the base assembly with the wire clamp parts attached to the blade tip.

Two blade springs will be installed on each modified suspension, and two modified suspensions are needed to form the 10 m cavity. Four blade springs were therefore chosen for use which had profiles best matching the FEA predicted curvature, and a tip gradient closest to horizontal. As expected, the four blade springs which showed no initial deformations had the closest match to the FEA curvature and also near-horizontal tip ends with the 300 g load and 7° launch angle.

3.3.1.4 Detailed Design and the Modification Process

With the conceptual design meeting optic length noise requirements and the characterisation of the blade springs completed, the detailed design of the suspension was modelled through computer-aided design to show the dimensions of the new suspension design and facilitate the manufacture of necessary parts.

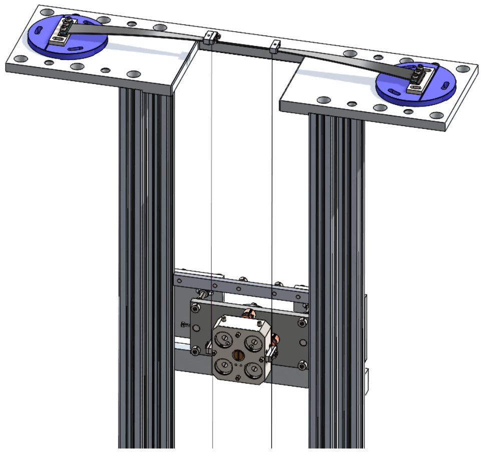

(a) CAD image of blade springs.

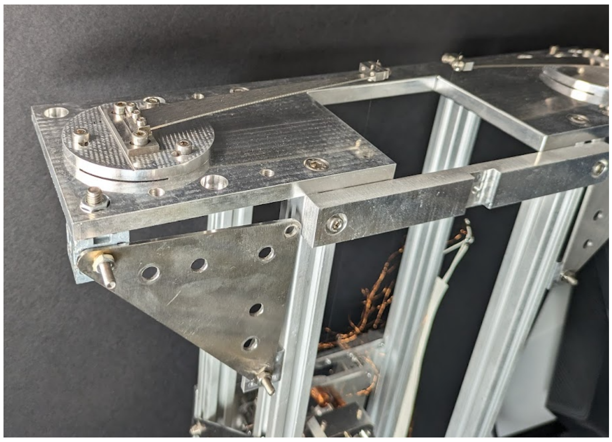

(b) Photograph of blade springs.

Figure 3.12: (a) Computer-aided design and (b) photograph of two blade springs installed on the 300 g suspension at the top wire mounting point. The blade spring mounting assembly is shown on the top surface of the modified top plate.

A portable, adjustable blade clamp was designed which would allow the rotational and translational adjustment of the blade springs during installation. This clamp was separated into several parts both for ease of manufacture and to allow new pieces to be easily swapped in. A close-up of the blade clamp design is shown in Figure 3.11. The aluminium base piece attaches to the suspension's top plate, with slots allowing rotational adjustment of the blade. An angled launcher can be attached to the base piece, with slots for translational adjustment of the blade position. If necessary, launch pieces with differing angles could be manufactured and swapped in for adjustment of the launch angle. The blade is placed on the angled launcher and sandwiched between the launcher and upper clamp. The two pieces in direct contact with the blade are manufactured from stainless steel to provide a stiff contact point, and contain a groove up their centre to ensure a line contact with the blade along the launch edge.

A crossed-blade design was initially considered which could fit the addition of the blade spring base assembly on the existing suspension's top plate. The blade springs could be fitted on this top plate and reach the required wire mounting point, however this design produced insufficient clearance around the blade tip for attachment of the wire to the wire clamps. Instead, the top plate was lengthened as shown in Figure 3.12 to provide easier access to the blade ends. Screw holes were added along to the length of the new top plate to provide attachment points for any future additions such as motion limiters to hold the blades at the correct height when working on the suspension, or gusset plates which would reduce resonance motion of the lengthened top plate.

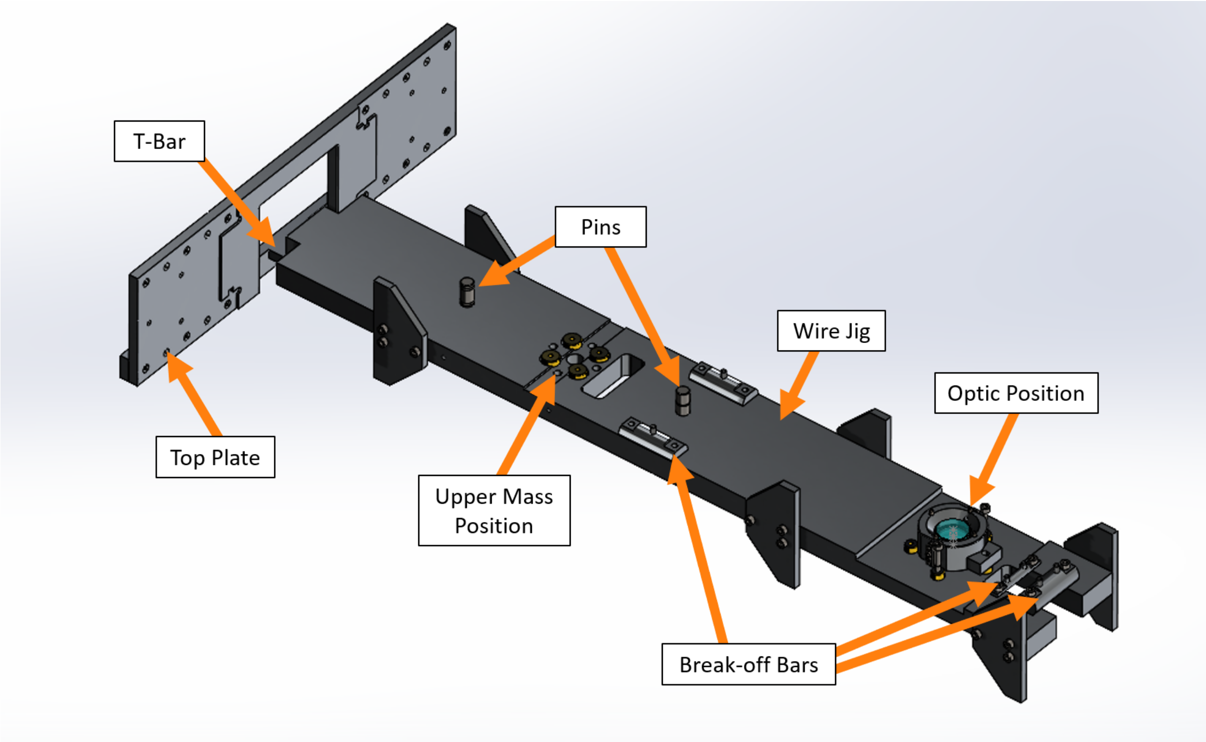

Figure 3.13: The wire jig used for the re-hanging of the modified 300 g suspension, showing the T-bar and new top plate on the left. The upper mass and optic positions are marked, as are the pins and wire break-off bars necessary for the tensioning and positioning of the wires.

The 'wire jig' is an aluminium back plate used in the hanging of the suspension to give the correct lengths, attachment points, and pre-tensioning of the wires. The suspension wire jig needed the addition of a T-bar to accommodate the new dimensions of the top plate which can be seen alongside the wire jig in Figure 3.13. In the hanging of the suspension, the top plate is removed from the suspension structure. The blade assembly is mounted on the top plate, with the top plate then attached to the wire jig and laid horizontally. The upper mass and optic are placed on the wire jig to correctly set the lengths of the wires. Wire loops are added at the pins which are laid over the break-off bars and weighted to set the appropriate wire tension, with the upper stage wires attached to the blade spring wire clamps. The wires are clamped at the sides of the upper mass and optic before any excess wire is cut and the weights are removed. The jig is lifted into a vertical position, allowing the top plate to be removed and re-attached (with the blades, wires, and masses) to the suspension structure.

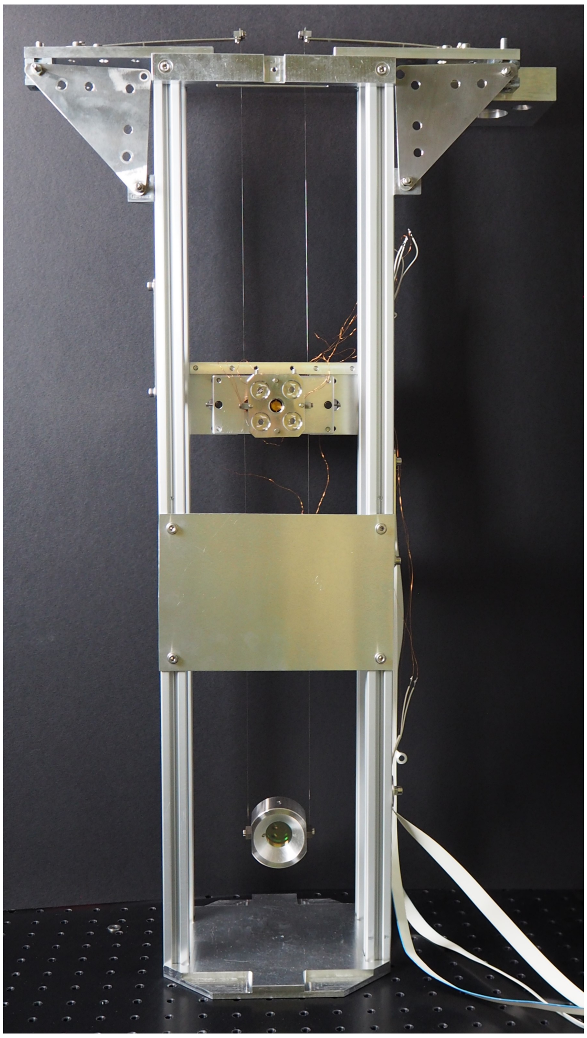

(a) Front view.

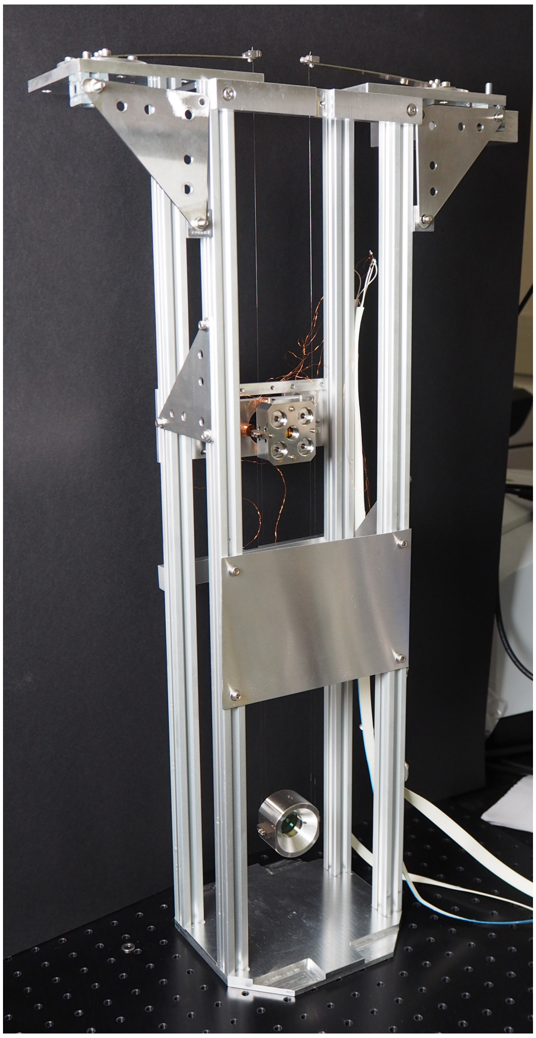

(b) Isometric view.

Figure 3.14: Photographs of modified 300 g double suspension with blade springs and elongated top plate installed.

3.3.2 Modification of the Auxiliary Suspension

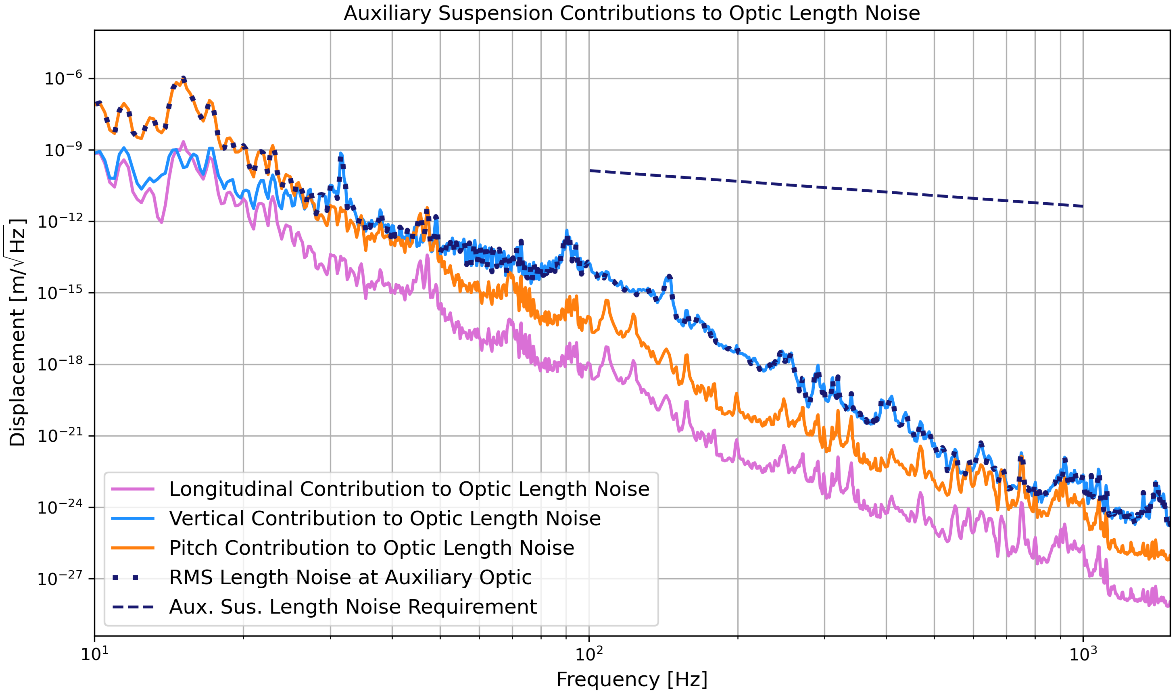

Figure 3.15: Auxiliary suspension contributions to optic length noise. The longitudinal contribution is shown in pink, the vertical contribution in blue, and the pitch contribution in orange. The optic total length noise is represented by a dotted dark blue line to match the previous colour convention while maintaining visibility of the vertical contribution.

The longitudinal, vertical, and pitch motion contributions to the auxiliary suspension optic total length noise are shown in Figure 3.15, calculated with the method discussed in Section 3.2.3.2 as was done for the 300 g suspension. Similarly to the 300 g suspension, the vertical contribution dominates the auxiliary suspension length noise in the entirety of the 100 Hz – 1 kHz sensitivity range. Additional vertical isolation was therefore added to the suspension to reduce the overall length noise.

3.3.2.1 Conceptual Design

The auxiliary suspension length noise, already sitting 825.8× lower than the noise requirement at 100 Hz, does not need such an involved modification as the 300 g suspension and its compact size creates a strict limit on space for any structural modifications. The addition of a rubber block beneath the wire mounting plate provided a simple way of increasing the vertical isolation and providing an additional margin around the length noise isolation requirements. This passive isolation technique is already utilised in the

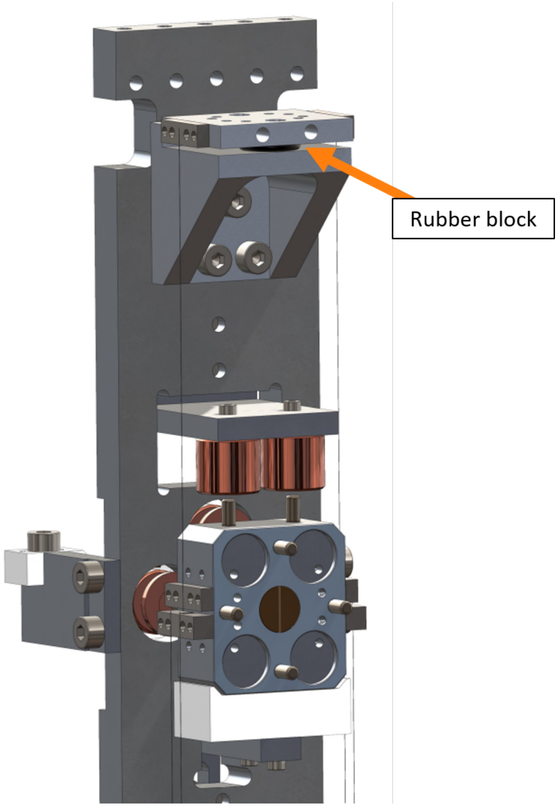

Figure 3.16: Computer-aided design image of the modified auxiliary suspension.

room-temperature vacuum chambers, with good isolation performance demonstrated and no current evidence of long-term continual drift in resonance frequency. This can easily be integrated into the existing suspension design, needing only a small modification to the existing top plate while maintaining the compact size of the original design.

A small rubber block would be placed beneath the clamping block for the wire mounting points, sitting on the angled bracket which attaches to the suspension's back plate as shown in Figure 3.16. Additional mass could be attached to the clamping block to appropriately weight the rubber for a chosen resonance frequency.

3.3.2.2 Modelling the Suspension Modifications

The resonance frequency of weighted rubber blocks were measured with the experimental method described in Section 3.2.3.1. Three smaller rubber blocks of $10 \times 10 \times 5$ mm formed from a soft rubber were used, and weighted with the 3 kg Thorlabs breadboard. The resonance of this weighted rubber was measured to be 32 Hz. As in Section 3.2.3.1 the transfer function of this stack was modelled to show the magnitude drop-off to higher frequencies than the measurement allowed.

The Mathematica transfer functions of the unmodified suspension were multiplied by the transfer function of the additional rubber stack with the same method as for the passive isolation stacks. This provides an estimation of the seismic noise coupling to the auxiliary suspension optic with the addition of rubber at the mounting point. A detailed discussion of the modelling results is given in Section 3.5.

3.4 Suspension Characterisation Methodology

A characterisation methodology was developed for the 300 g suspension to allow transfer functions of the physical suspension to be measured. These transfer functions could then be directly measured for the unmodified and modified suspension and compared to the modelled outputs. This comparison further verifies the modelling, and also demonstrates the development of the suspension through the modification process. For full characterisation of the suspension two processes were developed.

The upper mass longitudinal, pitch, and yaw motion coupling to optic long motion is considered in Section 3.4.1. Due to the driving force on the upper mass these three transfer functions are referred to in this work as 'force-displacement transfer functions'. The longitudinal optic motion resulting from various upper mass resonances is valuable when considering the optic length noise in an interferometer configuration, as optic length noise couples to cryogenic cavity length noise and can affect cavity measurements.

The domination of vertical-longitudinal coupling in the total length noise of the 300 g optic required a suspension modification process targeting the reduction of optic vertical motion. It was therefore necessary to directly measure the vertical-vertical transfer function of the suspension between the base plate and upper mass to measure the effectiveness of this modification. The measurement process is outlined in Section 3.4.2.

3.4.1 Force-Displacement Transfer Functions

The force-displacement transfer functions were measured between 0.1 – 100 Hz. This frequency range was chosen based on the model in Figure 3.6 to provide visibility of the pre-resonance motion, the resonance peaks, and the drop-off of the slope at frequencies above the resonances.

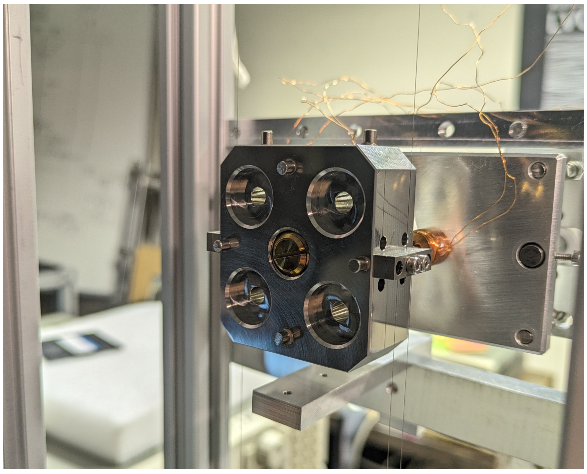

(a) Upper mass isometric view.

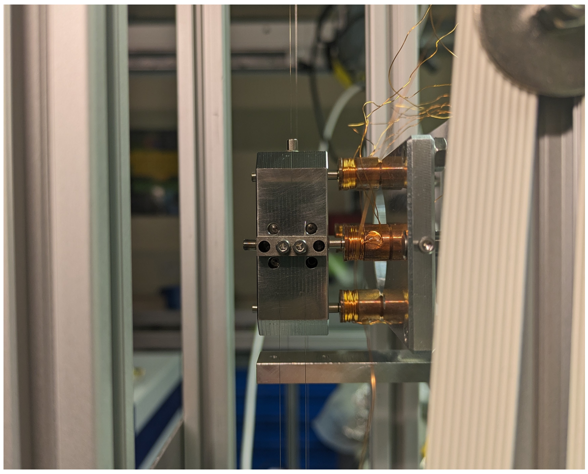

(b) Upper mass side view.

Figure 3.17: Photographs of upper mass isometric and side views, with (b) image showing the coil-magnet actuators used to drive upper mass motion.

In the experimental set-up the suspension base was placed on four rubber blocks to reduce the coupling of laboratory seismic noise. The upper mass was ensured to be free-hanging, with the coil actuators carefully positioned around the upper mass magnets to input a strong drive signal with no coil-magnet contact. A close-up of the upper mass, magnets, and coils is shown in Figure 3.17. The coils were attached via a BNC cable to the output port of a Stanford Research Systems SR785 dynamic signal analyser [179], with a signal used to drive the upper mass motion in longitudinal, pitch, and yaw via the electromagnetic actuators. The longitudinal motion of the optic was sensed using a Polytec OFV-505 vibrometer and recorded by the signal analyser, with the velocity output of the vibrometer being later converted to position data via integration. The vibrometer beam spot was centred on the optic to reduce pitch or yaw motions coupling to the longitudinal measurement.

A swept sine measurement was first performed over a wide frequency range to test the frequency response of the system. This technique measures the response of a system at a discrete frequency to a sine wave generated at that frequency, which is repeated across the desired measurement range. Measuring at each discrete frequency means the entire drive signal is concentrated at that frequency, giving this measurement technique a very high signal to noise ratio. Care must be taken around system resonances however, as the high drive signal can excite the resonance and if there is not sufficient settling time between measurements the excess energy can affect the measurement in the surrounding bins. For this reason, the SR785's FFT (Fast Fourier Transform) frequency response measurement mode was used to measure the frequency response around the resonance peaks. This mode employs band-limited white noise as the drive signal, effectively applying a small signal at every bin in the frequency range. The smaller signal per bin results in a reduced amount of drive and the technique has limited utility away from the resonances due to reduced signal to noise, however by exciting every bin simultaneously the distortion effect around the resonances from the sine wave approach is avoided. Combining the two approaches can give a clean signal across the whole band of interest. During analysis of the output data the vibrometer measurement is divided by the input drive signal to produce the suspension transfer function between upper mass motion and longitudinal optic motion. The results of the force-displacement transfer function measurements are presented in Section 3.5.1.2, and are comprised of a combination of swept sine measurements from 0.1 – 10 Hz and 10 – 100 Hz, and FFT measurements performed around each resonance peak.

3.4.2 Vertical-Vertical Transfer Function

Measurement of the vertical-vertical transfer function was performed for a frequency span of 10 Hz – 1 kHz, best fitting the region of interest modelled in Figure 3.3. The planned suspension modifications (discussed in Section 3.3) centre around softening the vertical mode's spring constant at the upper mass wire mounting point, leading to a preference of measuring the entire suspension transfer function from base to optic. At frequencies above the vertical resonances the transfer function therefore has a $^1/_{f^4}$ relationship with frequency. The higher frequency span coupled with this steeper signal drop-off results in a more challenging measurement than the force-displacement transfer functions as at higher frequencies the signal to noise quickly becomes a limiting factor. Clean measurements of an acceptable sensitivity across the wider measurement span took several iterations to achieve. The final methodology is described first, followed by a discussion on the development of this method.

The suspension base was placed on four rubber blocks, with a linear voice coil actuator of $8\,\Omega$ resistance clamped to the base plate for vertical drive of the suspension base. A SR785 signal analyser provided an input signal via a 5 V biased unipolar driver, and a ± 5 V range was input to the coil. This input drive signal was also connected to channel one of the signal analyser. A Brüel & Kjaer 4368 accelerometer was placed on the suspension base plate, connected to a Stanford Research Systems SR560 low-noise pre-amplifier which was then input to channel two of the signal analyser. Measurements were taken of the difference between the accelerometer signal and drive signal from 10 Hz – 1 kHz for the unmodified suspension, and 1 Hz – 1 kHz for the modified suspension. The accelerometer was then placed on the upper mass and the measurement was repeated, with both accelerometer placements shown in Figure 3.18. As for the force-displacement characterisation, a swept-sine measurement was first used across a large frequency span to observe the frequency

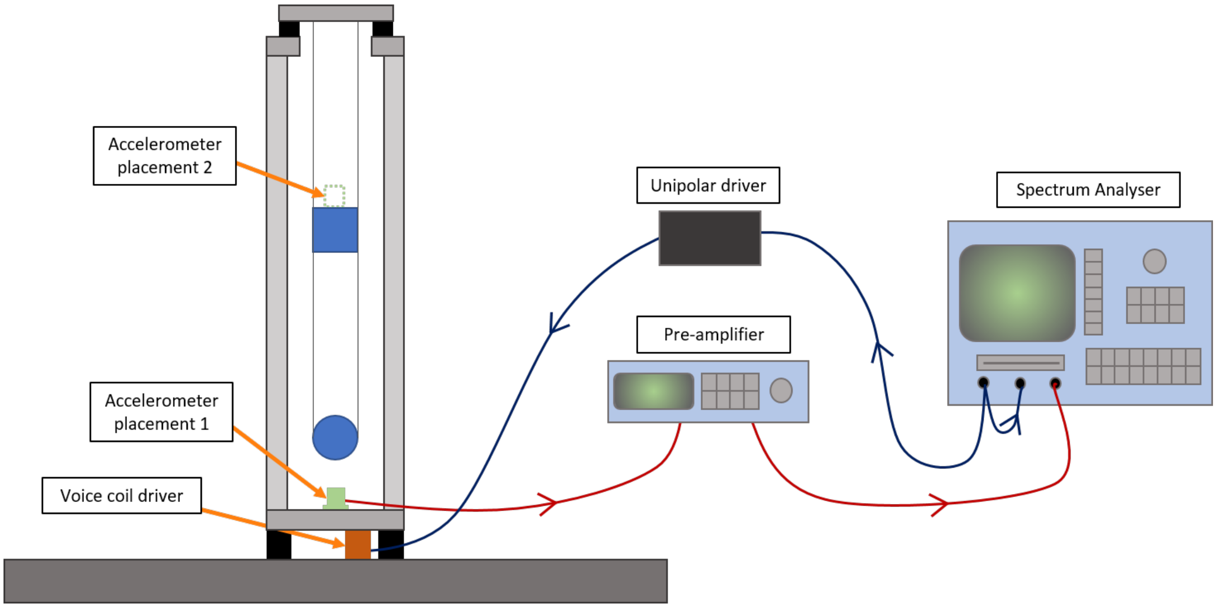

Figure 3.18: Experimental set-up for vertical-vertical transfer function measurements.

response. This was followed by FFT measurements of the response around the resonance peaks. During analysis of the output data the upper mass measurement is divided by the base plate measurement to produce the suspension transfer function between the vertical base plate motion and vertical optic motion.

Development of this methodology began with the build of a vertical shaker platform. Initially three piezoelectric actuators (PZTs) were mounted to a 0.3 m × 0.3 m breadboard and the vertical drive was tested with an accelerometer. The PZTs were not able to drive the mass of the board fast enough at frequencies above 100 Hz, with the exact frequency of the drop-off dependent on mass load. A voice coil actuator was then mounted beneath the breadboard for the second iteration of the vertical shaker platform which showed improvements but was still limited at 100 Hz when the weight of the suspension was added to the breadboard. This design eventually moved to direct mounting of the actuator on the suspension base to reduce the mass load: for a given drive force from the coil the reduction of the mass load increases the acceleration of the suspension base through Newton's second law of motion, further increasing the drive signal sensed by the accelerometers. The voice coil, made up of a magnet placed inside a solenoid, has a driving force proportional to current where the magnitude of the base motion delivered by the coil was initially limited by the 100 mA current output of the spectrum analyser. A custom driver box capable of sourcing 5 A was installed between the spectrum analyser and the coil driver, providing sufficient drive to the suspension base.

A Polytec OFV-505 laser vibrometer was also used to measure the displacement of the platform base and the suspension optic. This technique was favoured in early iterations for providing a no-contact measurement of the suspended masses, however the instrumental noise floor limited the measurement visibility after the magnitude drop-off at frequencies above 110 Hz. A Brüel & Kjaer Type 4368 accelerometer was then balanced on the top surface of the optic to sense optic vertical motion. This smaller accelerometer was favoured over the larger Type 4379 accelerometers as seen in Section 3.2.3.1 to decrease the mass added to the suspension during the measurement. With only one of these accelerometers available, this led to separate measurements of the suspension base motion and optic motion which were combined after the measurement to create the base-optic transfer function. Adding the accelerometer to the measured system affected the transfer functions slightly through the introduction of cable drag and the change of the moment of inertia (MOI) of the measured mass. The 30 g mass increase from the accelerometer was included in the Mathematica modelling, however cable drag and MOI changes were not modelled.

The $1/f^6$ slope of the unmodified suspension transfer function was visible above the noise floor up to 150 Hz with this measurement method, however this was not a sufficient distance from the resonant peaks to provide good visibility of this slope. Moving the accelerometer to the upper mass and instead measuring the base-upper mass transfer function produces a $1/f^4$ magnitude drop off for the unmodified suspension and provided slope visibility up to a more comfortable 300 Hz. The results of the vertical-vertical transfer function measurements are presented in Section 3.5.1.3.

3.5 Results of the Characterisation and Modelling

3.5.1 300g Double Suspension

This section shows the results of the modelling and characterisation work performed on the 300 g suspension using the methods outlined throughout this chapter. The results of the modelling work performed on the unmodified and modified suspensions are shown in Section 3.5.1.1. A discussion of the predicted comparison between the two models is shown, as are the transfer functions which are later used in Section 3.5.3 to predict the 300 g optic length noise for the two cases. The characterisation work is discussed in Section 3.5.1.2 with a comparison of the modelled force-displacement transfer functions to those measured on the physical suspension. This comparison is shown for both the unmodified and modified suspension. Finally, the results of the vertical-vertical transfer function measurements are presented in Section 3.5.1.3. The data is shown against the theoretical predictions, and improvements to the suspension design are discussed.

3.5.1.1 Modelled Transfer Functions

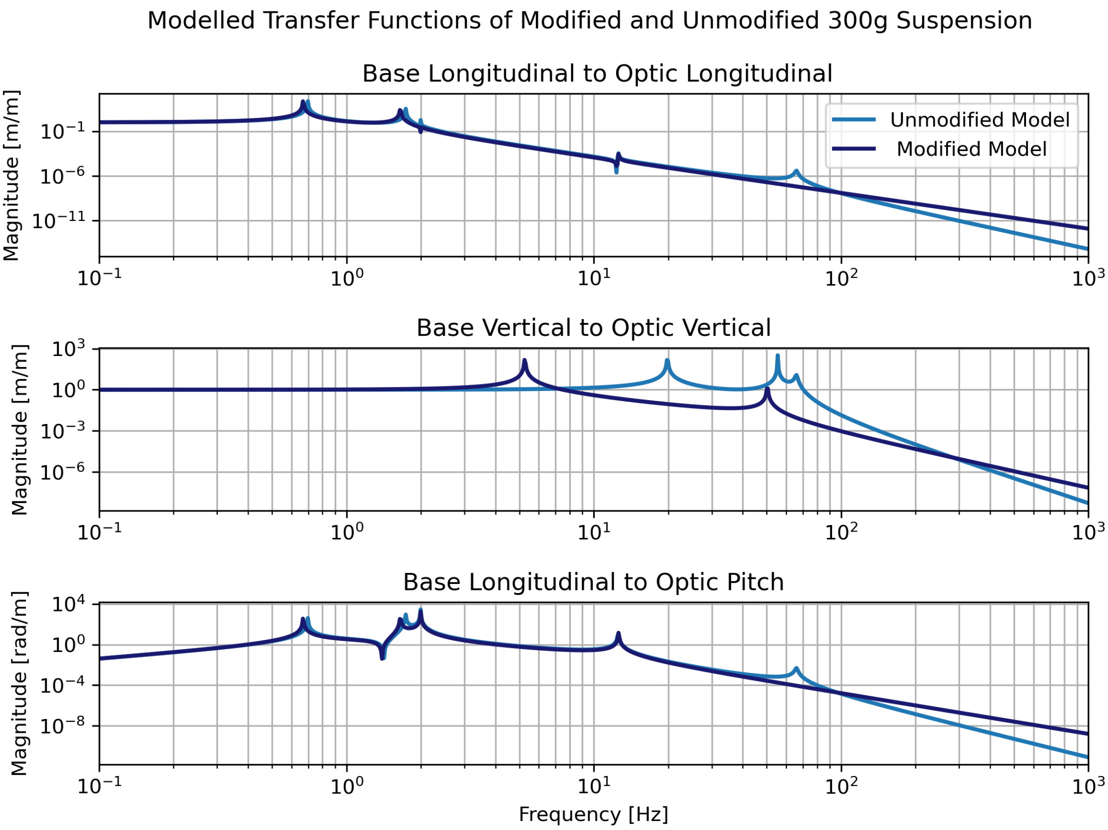

Figure 3.19: Comparison of the unmodified and modified 300 g suspension modelled transfer functions. The transfer functions are between motion at the suspension base to optic motion. The top plot is longitudinal base motion to longitudinal optic motion, the middle plot is vertical base motion to vertical optic motion, the lower plot shows longitudinal base motion to optic pitch motion. The outputs of the unmodified suspension are shown in light blue, and the modified suspension outputs in dark blue.

The modelled base-motion to optic-motion transfer functions were calculated with the method discussed in Section 3.2.3 for both the unmodified suspension and the modified suspension. The longitudinal-longitudinal, vertical-vertical, and longitudinal-pitch transfer functions can be seen in Figure 3.19 The unmodified suspension transfer functions are the same as those shown in Figure 3.3, and are shown again here for comparison between the two cases.

As expected, the addition of blade springs to the 300 g suspension does not largely affect the longitudinal or pitch resonant modes, but does have a significant affect on the vertical-vertical transfer function. The 19.86 Hz vertical mode of the unmodified suspension is reduced to 5.27 Hz on the modified suspension due to the low spring constant of the blade springs lowering the resonant frequency in comparison to the wire-only case. The addition of blade springs also slightly affects the second vertical resonance which is reduced from 55.50 Hz to 50.36 Hz. The unmodified suspension contains the 66 Hz resonance and extra $1/f^2$ drop-off from the mounting point rubber blocks. The blade springs in the modified suspension replace this rubber and so display an overall $1/f^4$ slope of magnitude against frequency. The modified suspension is predicted to couple less seismic motion to the optic in the 100 – 300 Hz region, which was critical for reducing the 300 g optic length noise. The increased vertical-vertical coupling of the modified suspension between 300 Hz – 1 kHz is not a concern, as the optic length noise at 300 Hz already lies over $10^8 \times$ lower than the noise requirement, with optic length noise still decreasing with increasing frequency. If desired, the rubber blocks could be re-integrated into the modified design below the top plate and act in addition to the blade springs to increase the high frequency seismic isolation, however were omitted in this case to reduce the number of resonant peaks near the experimental sensitivity band.

The longitudinal-longitudinal and longitudinal-pitch transfer functions were not directly measured and so are not shown in comparison to experimental data, however these transfer functions are later used to calculate the expected optic length noise in Section 3.5.3. The vertical-vertical model predictions are shown in comparison to the measured data in Section 3.5.1.3.

3.5.1.2 Modelled and Measured Force-Displacement Transfer Functions

The longitudinal, pitch, and yaw couplings of the upper mass to longitudinal optic motion are shown in Figure 3.20 for the unmodified suspension, and in Figure 3.21 for the modified suspension. They represent a torque of 1 Nm incident on the upper mass and the resulting longitudinal optic motion. The measured transfer functions, in pink, are the results of the methodology described in Section 3.4.1, normalised to a torque of 1 Nm to match the model. The dashed blue line represents the initial modelling output and the dashed dark blue line is a modification to the modelling output which is detailed at the end of this section.

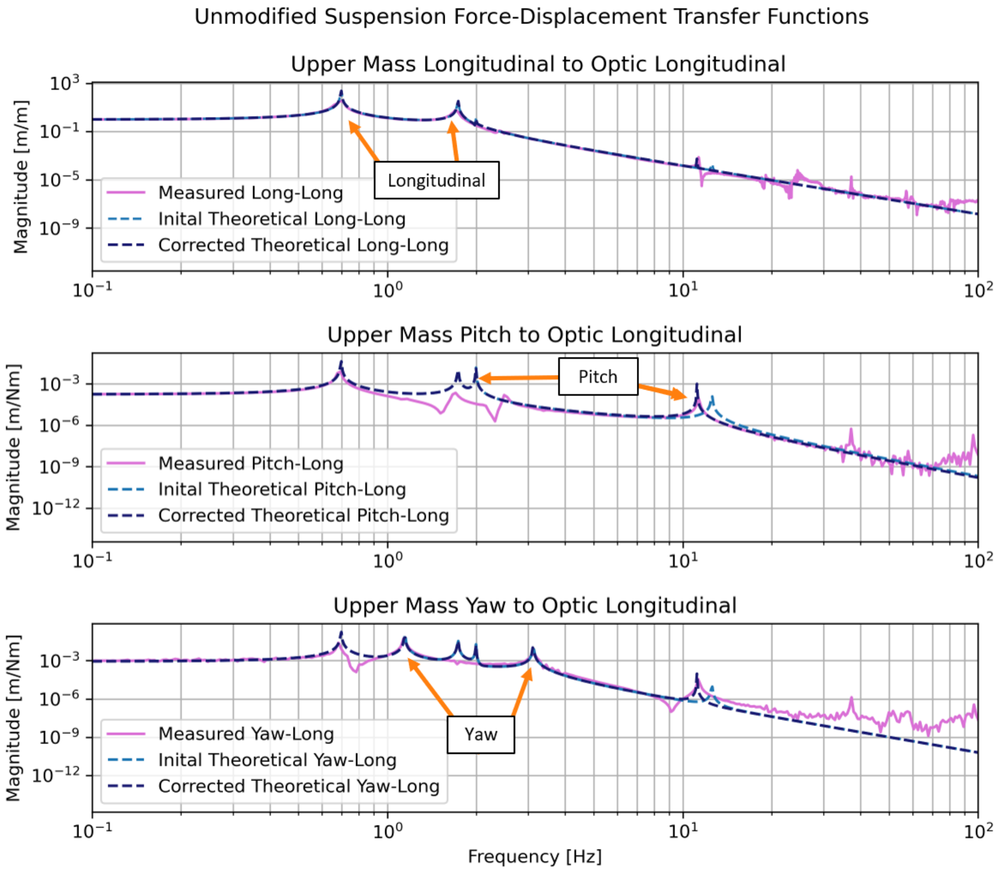

Figure 3.20: Comparison of modelled and measured force-displacement transfer functions for upper mass drive to optic motion of the unmodified suspension. Top: upper mass longitudinal drive to optic longitudinal motion. Centre: upper mass pitch to optic longitudinal. Bottom: upper mass yaw to optic longitudinal. The measured data is represented by a solid pink line. The modelling outputs are dashed, with the initial outputs in blue and a small correction to the pitch modes in dark blue.

For both the unmodified and modified suspensions the predicted longitudinal resonances can easily be identified in the data, and match the model predictions well. The nature of driving an input force on suspended mass means there will be a small pitch motion expected at the optic, as perfect longitudinal drive is difficult to achieve. These pitch resonances are seen in the measurement at 2.4 Hz and 11.4 Hz for the unmodified suspension, and 3.4 Hz and 11.2 Hz for the modified suspension. The modelled longitudinal-longitudinal transfer function does predict peaks for the pitch resonances but these are very small in magnitude. The higher magnitude of the pitch resonances was expected in the characterisation measurement due to the difficulty of perfectly aligning the vibrometer beam in the centre of the optic as the slight misalignment causes optic pitch motion to be sensed as longitudinal optic motion by the vibrometer.

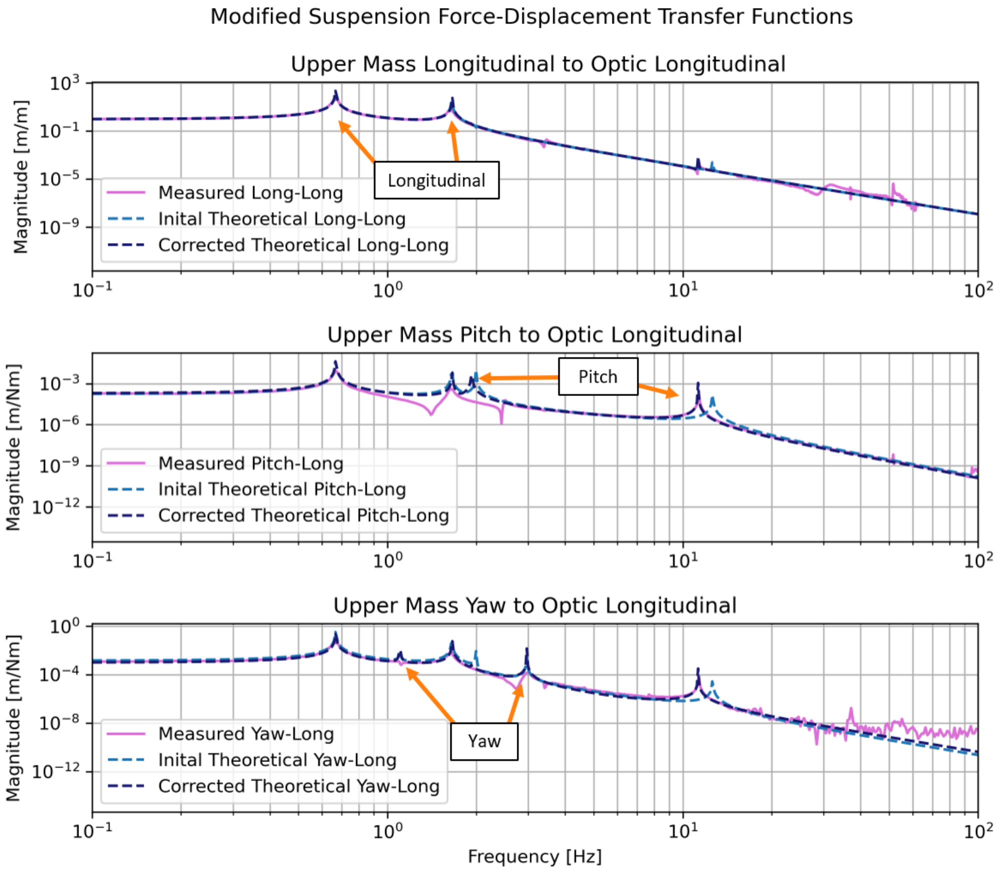

Figure 3.21: Comparison of modelled and measured force-displacement transfer functions for upper mass drive to optic motion of the modified suspension. Top: upper mass longitudinal drive to optic longitudinal motion. Centre: upper mass pitch to optic longitudinal. Bottom: upper mass yaw to optic longitudinal. The measured data is represented by a solid pink line. The modelling outputs are dashed, with the initial outputs in blue and a small correction to the pitch modes in dark blue.

The measured pitch-longitudinal and yaw-longitudinal functions are more complex and show the coupling of multiple suspension resonances. In these cases the predicted angle to longitudinal transfer functions alone did not explain the full measurement signal, so were combined with other modelled transfer functions to give a more complete understanding of the resulting optic motion. These are described separately for the pitch and yaw motions below.

The pitch-longitudinal force-displacement transfer functions for the unmodified and modified suspensions (centre plots of Figures 3.20 and 3.21) display similar features and so are described together below. In both unmodified and modified suspension cases, peaks are observed at the longitudinal resonances around 0.7 Hz and 1.7 Hz which match the frequencies predicted by the model. There is a back reaction peak visible between these resonances which comes from the suspension support structure. This is not included in the theoretical model as the Mathematica modelling does not account for the suspension support structure. The first predicted pitch resonance occurs in the unmodified suspension model at 2 Hz, and at 1.9 Hz for the modified suspension model. In both cases, the predicted resonance is lower in frequency than the 2.5 Hz pitch resonance observed on the suspension. Similarly at higher frequencies the 11.3 Hz measured resonance is predicted at near 12.5 Hz for both models. This shows a discrepancy between the model predictions of pitch modes and those observed on the physical suspension. The predicted pitch-longitudinal transfer function accounts for the measured optic motion between 0.1 – 2 Hz. Above 2 Hz the pitch-longitudinal model predicts the discussed resonant frequencies, but the slope drop-off with frequency is steeper than is observed in the measured data. This means there is an additional contribution to the optic longitudinal motion which is limiting the magnitude of the pitch-longitudinal drop-off.

This difference can be explained if the upper mass pitch motion coupling to optic pitch motion is considered. Some fraction of the optic pitch motion will be sensed by the vibrometer as longitudinal motion if there is even a small misalignment of the beam, which is seen to be the case from the longitudinal-longitudinal transfer function discussed above. A comparison of the results to the pitch-pitch transfer function shows that a 0.1% pitch-pitch coupling to longitudinal motion would perfectly match the slope measured experimentally. This 0.1% coupling would occur for a beam misalignment of $<$1 mm, which seems likely in this set-up as the beam centring was performed by eye. The resulting combination of pitch-longitudinal and 0.1% pitch-pitch is displayed by the light blue dashed line in Figures 3.20 and 3.21.

The modelled yaw-longitudinal transfer functions (lower plots of Figures 3.20 and 3.21) primarily predict visibility of the two longitudinal optic resonances. The predicted 0.7 Hz peak matches the observed longitudinal resonance in both suspension cases, with a back-reaction peak visible in the unmodified suspension data. A second longitudinal peak is visible in the models and data at around 1.7 Hz. This is especially visible in the modified suspension, with an amplitude which agrees well with the measurement. The unmodified suspension has only a small peak measured at this longitudinal resonance and has more energy in the yaw and pitch modes than the yaw-longitudinal model predicted.

The additional yaw and pitch resonances visible in the yaw-longitudinal measurement suggest there is some misalignment causing coupling in the system, such as a vibrometer beam misalignment, or some misalignment in the suspension build which causes upper mass yaw motion to drive optic pitch and yaw motion more strongly than predicted (or a combination of both effects). This coupling is not observed as strongly for the modified suspension. One possible cause is the upper mass magnets not being perfectly centred inside the copper coil actuators, leading to uneven drive on the upper mass which could couple into other suspension modes. Another possibility is that this misalignment is from the suspension's wire attachment points and was reduced by the re-hanging of the suspension during the modification.

As with the pitch-longitudinal transfer functions, the predicted yaw-longitudinal transfer functions do not account for the entire longitudinal motion of the optic: yaw and pitch resonance are observed in the measured data. Upper mass yaw to optic yaw coupling to longitudinal of 0.1% accounts for the observed 1.14 Hz and 3.1 Hz peaks in the unmodified suspension, and for the 1.1 Hz and 3.0 Hz peaks of the modified suspension. The modelled yaw resonances match those measured on the suspension. A pitch resonance is also visible in the measured data at 11.4 Hz. This resonance, and the slope of the transfer function drop-off with frequency, matches a dominance of the upper mass pitch to optic longitudinal transfer function at frequencies above 10 Hz. This contribution is likely due to a combination of vertical vibrometer beam misalignment, as well as a slightly misaligned or unbalanced upper mass which has some energy in a pitching motion.

With the comparison of the Mathematica model and measured resonances it is easy to see there is good agreement between the theoretical and measured resonances. However, there a discrepancy of the 11.4 Hz pitch mode which is predicted by the model at 12.5 Hz. The modelled wire attachment points were slightly adjusted to see if the model could closer match the resonances of the physical suspension while adjusting parameters to within build and machining tolerances. It was found that a change to the optic's x-direction wire attachment of <0.5 mm (0.45 mm) produced a model which matches the 11.4 Hz resonance of the physical suspension. The force-displacement transfer functions of this 'corrected' model are shown in Figures 3.20 and 3.21 by the dashed purple line. This demonstrates the intricacies of both the modelling and the build of such a sensitive physical system, and the care needed to ensure the input of realistic and accurate model parameters. The measured transfer functions are dominated by noise close to 100 Hz for the UM longitudinal and pitch motion plots, and nearer 40 Hz for UM yaw motion, a property also observed in [103] which is attributed to pendulum characteristics and the characterisation of such steep drops in response with increasing frequency. These difficulties show the necessity of accurate experimental characterisation, with the outputs of this work being extremely valuable for the commissioning and running of experiments utilising these suspensions.

3.5.1.3 Measurements of the Vertical-Vertical Transfer Function

The results of the vertical-vertical transfer functions measurements are shown in Figure 3.22. The unmodified suspension measurement is represented in yellow and the modified suspension measurement in blue. Theoretical predictions of the vertical-vertical transfer function were output using the Mathematica model which are represented in orange for the unmodified suspension and green for the modified suspension. These predicted curves display the transfer function between suspension base motion to upper mass motion, and so differ from those in Figure 3.19 which show suspension base motion to optic motion. Additionally, the model outputs in Figure 3.22 include the 30 g weight of the accelerometer at the upper mass.

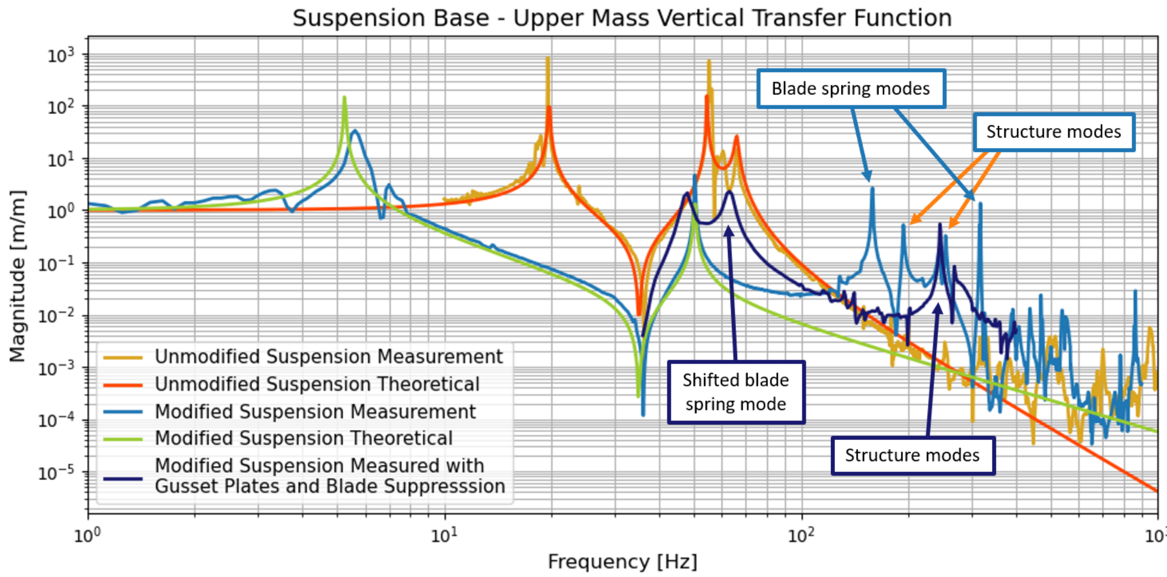

Figure 3.22: Comparison of the modelled and measured vertical-vertical transfer functions between the suspension base and upper mass. The modelling outputs are shown in orange for the unmodified suspension and green for the modified suspension. The measured data is represented by a yellow line for the unmodified suspension and blue for the modified suspension. Measurements of the modified suspension including gusset plates and blade mode suppression is shown in dark blue between 35 – 400 Hz.

The Mathematica model's transfer function of the unmodified suspension predicts resonance peaks at 19.5 Hz and 54.4 Hz, occurring from motion between the upper wire mounting points and the optic. As noted earlier, this model does not include the rubber blocks at the top of the suspension structure. With the method used in Section 3.2.3.1 an additional resonance peak was added to account for the rubber resonant frequency where a peak of 66 Hz best fitted the peak observed during the measurement. This differs from the 29.5 Hz resonance frequency predicted for this rubber in [88], emphasising the need for transfer function measurements of the suspensions and passive isolation stacks during the commissioning of the GCIF.

The unmodified suspension measurements were taken between 10 Hz – 1 kHz with the methodology of Section 3.4.2. This data displays large resonance peaks at 19.5 Hz and 55 Hz which match well with the 19.5 Hz and 54.4 Hz predictions of the model, and are caused by resonant motion between the upper mass and the optic. There is an additional resonance in the data between 57 Hz – 67 Hz which is not well defined by the measurement. Modelling the fitting of the rubber resonance peak with the magnitude drop-off in frequency from 80 – 300 Hz predicts this peak to lie close to 66 Hz. Above this 66 Hz resonance a clean drop-off in magnitude of $1/f^6$ is observed up, expected from the three stages of isolation provided by the rubber and double-stage pendulum. This slope is clearly visible until 300 Hz above which point the measurement runs into the experimental noise floor which is limited by the pre-amplifier.

The modified suspension transfer function was measured from 1 Hz – 1 kHz, with this range chosen to provide visibility of the the low frequency resonance. This resonance is observed between 5 – 6 Hz and appears to contain a peak at 5.6 Hz, however noise in the data below 8 Hz means this is not well enough defined to give an accurate measurement of the resonance frequency. The modelled resonance is predicted slightly lower at 5.2 Hz, which lies within the uncertainty of the peak position of the measurement. The noise around the peak is thought to come from the balancing of the accelerometer on the upper mass: the change in upper mass MOI and the addition of accelerometer cable drag were not included in the modelling but are assumed to affect the energy in this resonance peak, lowering the Q of the suspension and splitting the mode. This peak shape was consistent over multiple data runs and tests with high levels of averaging. The outputs around the second vertical resonance show good agreement between model and predictions, both displaying a peak at 50 Hz.

At frequencies above the 50 Hz peak there begins to be disagreement between the model and measurements. The model predicts there will be a clean $1/f^4$ drop in magnitude which would be expected to be observed until the limit of the noise floor at 300 Hz, however the measured data displays a large resonance peak at 158 Hz, and a range of lesser peaks from 180 – 350 Hz. The origins of these peaks are explored in the following section.

Higher Frequency Peak Investigation:
The peaks observed between 158 – 350 Hz in the modified suspension measurements were not visible in the unmodified suspension measurements, and were not at resonances predicted in the modelling. It was therefore assumed that these peaks were the result of the changes made to the suspension during the modification process. Due to extensive modelling of the suspensions themselves the explanation must lie with elements of the modification that were not included in the model.

It was suspected that suspension frame resonances could arise from the new top plate, which was an elongated and heavier version of the original design. This can affect suspension resonances by shifting them into the measured frequency span or by driving motion of the new top plate. Gusset plates were added to the midpoint of the suspension legs on all sides and to the underside of the top plate, and mass was bolted to the ends of the top plate. This identified the majority of the new peaks as suspension structure resonances. The gusset plates added stiffness to the legs, reducing the midpoint motion and pushing these resonant frequencies higher in frequency space, and a number of peaks were observed to decrease in magnitude.

The 158 Hz resonance was not affected by any changes to the top plate or suspension support structure. It was found that the 158 Hz peak was composed of two independent, overlapping 158 Hz resonances which each come from one of the steel blade springs. It was possible to independently adjust these resonances by the balancing of rubber blocks on the top surface of the blade spring tips. The additional mass on the blade pushes the resonance down in frequency, where a higher additional mass results in a lower frequency resonance. The 158 Hz mode was easily able to be reduced to 63 Hz in this manner. The clearing of modes from adding the gusset plates also showed a second blade spring resonance at 316 Hz, exactly twice the frequency of the 158 Hz peak. The 316 Hz resonance was also affected by the addition of mass to the blade spring. Higher frequency blade modes were unable to be found as the measured signal runs into the noise floor at higher frequencies.

It was thought that the blade spring resonance could arise from a torsional mode or from a xylophone mode as can be seen in Figure 3.9, although these were not appearing at the resonance frequencies predicted through the FEA modelling. To identify the type of resonance, measurements were taken of the blade motion on either side of its central y-axis using the laser vibrometer. The base plate motion was driven using the coil-magnet actu-

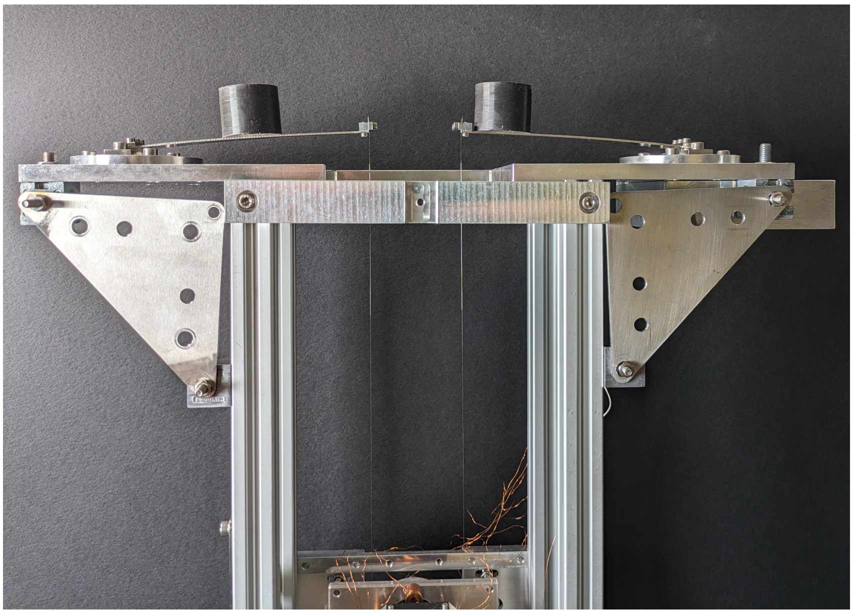

Figure 3.23: An example of blade spring weighting showing suppression of the first order xylophone mode on the left of the image, and weighting of the blade tip on the right.

ator as in Section 3.4.2 to excite the blade spring resonances while the vibrometer beam was directed to the underside of the blade spring using a 45° mirror. These measurements showed there was no phase flip between the blade edges therefore the origin could not be a torsional mode.

A xylophone mode origin was investigated through the positioning of rubber blocks along the length of each blade spring as shown in Figure 3.23. It was observed that either the first or the second order blade modes could be independently damped by careful positioning of the blocks along the blade length. This supports the xylophone resonance theory where mode damping could occur when rubber blocks are placed at the anti-node of each resonance. Further support for a xylophone type origin comes from the mode observed at 316 Hz, which is thought to be a second order harmonic of the 158 Hz mode. It is expected to observe higher order resonances for xylophone-type modes at near-integer multiples of the first-order resonance as is the case for the violin modes (see Section 2.2.3). Finally, the balancing of additional mass of up to 30 g on the optic did not affect the observed amplitude or frequency of either resonance. This change to the loading of the blade is predicted by FEA to shift the frequency of a torsional resonance, but not predicted to affect xylophone mode frequencies.

The experimental investigation supported a xylophone mode origin, however the 158 Hz and 316 Hz observed resonant frequencies of these modes differ from the FEA blade resonance predictions of Section 3.3.1.2 which output xylophone modes at 239.7 Hz and 843.5 Hz. The FEA models only a single blade spring, with the optic weight modelled as an 'Added Mass' along the blade tip edge, and the clamping as a 'Fixed Constraint' over the clamping-end boundaries. Various parameters of the FEA model were adjusted to explore the 82 Hz disagreement between the predicted and observed xylophone resonances to within machining or experimental tolerances. The blade spring geometry, stiffness, and mass loading location were all checked against the real blade and adjusted to within measurement tolerances. Wire end clamps were imported into the FEA model. The location and softness of the Fixed Constraint clamping at the blade base was also investigated. No reasonable change in any of these parameters produced a first-order xylophone mode below 200 Hz. Future investigations could develop the finite-element modelling through the addition of the top plate clamping elements to make the simulated blade clamping more realistic, or through the addition of the entire suspension structure which would better model the cross-coupling between the blade spring and structure resonances. This work would be extremely intensive in both time and computational power. Since the experimental investigation had identified the origin of the higher order modes it was decided that this further FEA analysis was beyond the scope of this work.

The dark blue line in Figure 3.22 shows the vertical-vertical transfer function measurement from 35 – 400 Hz after the addition of gusset plates and with rubber blocks on the blade springs. The 55 Hz resonance is slightly lowered to 47.5 Hz when the blade springs are weighted, but the largest difference is to the blade spring resonances. The 158 Hz resonance has been reduced to 63 Hz, while the 316 Hz mode can not be observed when the blades are weighted to a 63 Hz frequency and is assumed to have been damped by the rubber blocks. The use of gusset plates and blade mode suppression produces a much cleaner transfer function between 100 – 200 Hz for the modified suspension.

Overall the suspension modifications have had the desired effect on the two vertical resonances, lowering both in frequency and reducing the optic's vertical motion at 100 Hz. However, the modification has resulted in some unwanted additional resonances which affect the vertical transfer function in the frequency range of interest. It has been shown that these additional modes arise from the blade spring xylophone modes and from the support structure. They can be adjusted through stiffening of the support with gusset plates, weighting of the top plate, and by suppression of blade spring motion.

Reduction of the additional modes could be furthered with future development of the support structure, for example investigating widening the support legs to the corner of the top plate to reduce 'flapping' top plate modes, finding the optimum weighting of the top plate to push resonances out of the sensitivity range, or improving on the gusset plate placements. The rubber blocks on the blade spring surface has been shown to easily damp the blade xylophone modes, reducing them in frequency to below the experimental sensitivity range. Coil-magnet actuators could be installed on the blades to replace the rubber blocks, allowing specific suppression of unwanted blade spring modes and providing greater control to the wire mounting point than was possible pre-modification.

3.5.2 Auxiliary Suspension

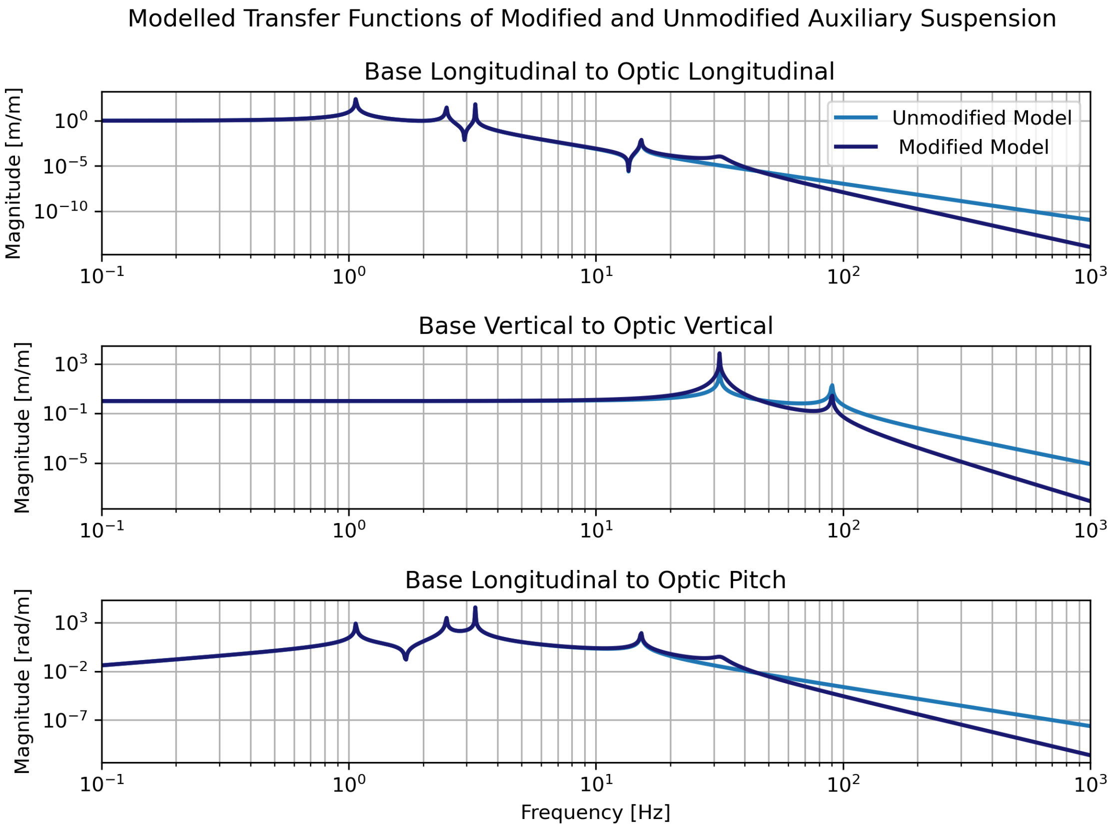

Figure 3.24: Comparison of the unmodified and modified auxiliary suspension modelled transfer functions. The transfer functions are between motion at the suspension base to optic motion. The top plot is longitudinal base motion to longitudinal optic motion, the middle plot is vertical base motion to vertical optic motion, the lower plot shows longitudinal base motion to optic pitch motion. The outputs of the unmodified suspension are shown in light blue, and the modified suspension outputs in dark blue.

The modelled transfer functions are shown in Figure 3.24 for the auxiliary suspension. The Mathematica output of the unmodified suspension is shown in light blue, with the dark blue showing the model containing rubber at the top of the support tower. This rubber has a measured resonance of 32 Hz, with this peak showing in the longitudinal-longitudinal, vertical-vertical, and longitudinal-pitch transfer functions between base motion and optic motion. The addition of the rubber reduces the coupling of base motion to optic motion in all cases between 100 Hz – 1 kHz compared to the unmodified model.

These three transfer functions are used to estimate the optic length noise of the modified suspension, using the method outlined in Section 3.2.3.2. The vertical-vertical transfer function was shown in Figure 3.15 to dominate the unmodified suspension optic length noise. The addition of the rubber block is estimated to reduce this coupling by 9.1× at 100 Hz and provide an additional $1/f^2$ reduction in magnitude with frequency above 32 Hz, meeting the design requirements.

3.5.3 Optic Length Noise

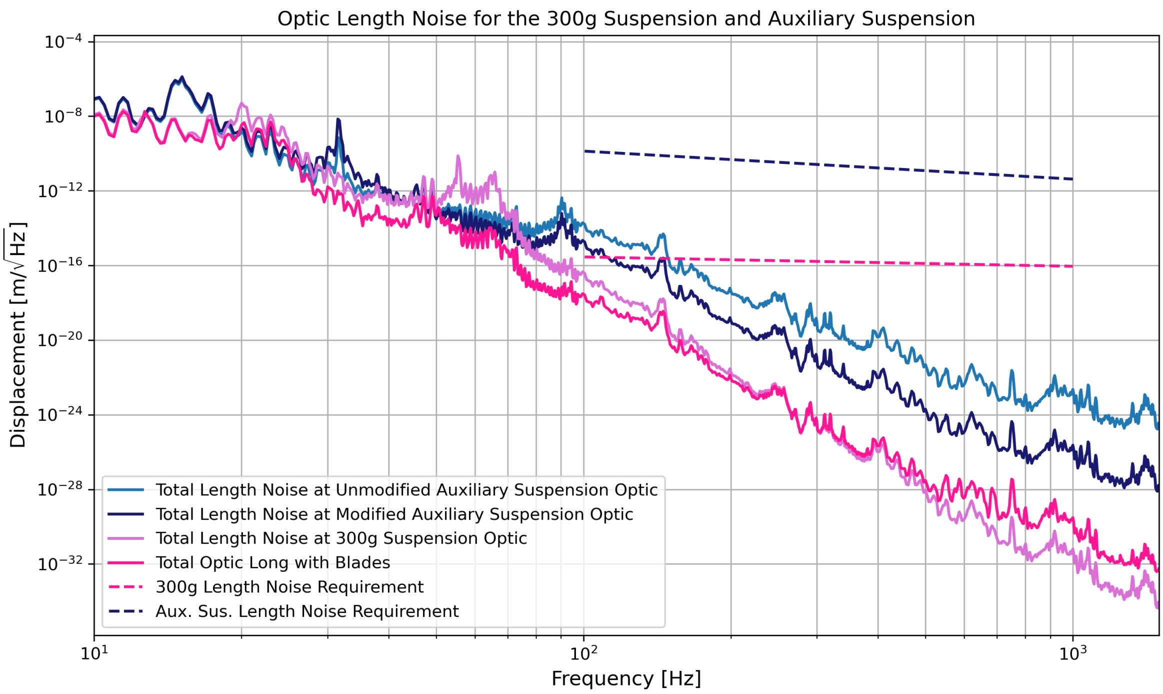

Figure 3.25: Comparison of the optic length noise of the unmodified and modified suspensions, calculated using the modelled suspension transfer functions. The unmodified auxiliary suspension length noise is shown in blue, and the modified version in dark blue. The unmodified 300 g suspension length noise is shown in pink and the modified version in dark pink. The optic length noise requirement are represented by a dashed line for both suspensions, with dark blue for the auxiliary and dark pink for the 300 g.

The unmodified suspension optic length noise and length noise requirements were first shown in Figure 3.5, and are reproduced in Figure 3.25. The unmodified auxiliary suspension optic length noise is shown in blue, with the dashed dark blue representing the requirement. The unmodified 300 g is shown in pink, with the dashed dark pink representing this suspension's length noise requirement. Optic length noise was then calculated for the modified suspensions using the method from Section 3.2.3.2 and the modelled suspension transfer functions, which are shown in Figures 3.19 and 3.24. The modified suspension length noise outputs are represented in solid dark blue for the auxiliary suspension and solid dark pink for the 300 g suspension.

The modelling demonstrates the desired improvement in the length noise of both suspensions across the 100 – 200 Hz frequency span, where the improvement was targeted. At higher frequencies, suspension length noise steeply decreases with frequency so the increase in length noise of the modified 300 g suspension above 200 Hz is not a concern. The predicted decrease in optic length noise at 100 Hz through modification is 8.9× for the auxiliary suspension, and 14.4× for the 300 g suspension when compared to the unmodified designs. This places the modified suspensions 73375× below the 100 Hz length noise requirement for the auxiliary suspension, and 108.5× for the 300 g. This meets the aims of the modification process as outlined in Section 3.3, where a 10× reduction in 300 g optic length noise was desired.

The experimental measurement of the vertical-vertical transfer function was taken between the suspension base and the upper mass. This increased the signal-to-noise ratio of the data by avoiding the extra $1/f^2$ magnitude decrease arising from measuring down an additional pendulum stage at the optic. However, this means the measured data cannot be used to make a prediction of the true length noise of the 300 g suspension. The design work performed has allowed the 300 g suspension to have a potential decrease in length noise of 108.5× at 100 Hz, however some further work is necessary to have the suspension meet this potential. The addition of coil-magnet actuators on the blade springs is recommended for damping of the blade spring xylophone resonances, and has the benefit of providing vertical control at the wire mounting point. Further optimisation of the suspension support structures is recommended to move any top plate and structure resonances away from the frequency span of interest. This can be done either by reducing these to below 100 Hz, or by increasing the resonances to a few hundred hertz where they are unlikely to be visible at the optic due to the $1/f^4$ magnitude drop with frequency. If the first order blade xylophone resonance can be reduced to 50 Hz and the support resonances can be shifted, the suspension's vertical-vertical transfer function is expected to match the modelled transfer function and produce the modelled optic length noise. The slightly crude weighting of the blade spring with rubber blocks is able to reduce the first order xylophone mode from 158 Hz to 63 Hz, so this is expected to be a reasonable aim.

3.6 Conclusion

This chapter covered the modelling, characterisation, and modification work performed on two double-stage suspensions for use in room-temperature areas of the GCIF. The GCIF design sensitivity requirements on cryogenic cavity length noise were outlined, and converted to optic length noise for the 300 g and auxiliary suspension separately which allows for easier comparison to the requirements.

A full dynamical model was made for the two double-stage suspensions. This model output the suspension resonant frequencies which were compared to analytical mode calculations, and to previously published mode frequencies. This verified the dynamical model which later allowed parameter adjustments to result in believable resonant mode predictions.

The dynamical model was then used to calculate transfer functions between different elements of the suspension. The longitudinal, vertical, and pitch motions of the optic from both horizontal and vertical suspension base motion were output by the model. These transfer functions were used to calculate the resulting optic length motion from ground seismic motion for both suspensions, including attenuation from passive isolation stacks between the suspension base and laboratory floor. This modelling was compared to the optic length noise requirements, showing that the auxiliary suspension length noise lies beneath the requirement in the entire 100 Hz – 1 kHz sensitivity range whereas the 300 g suspension length noise lies near the requirement at 100 Hz. This motivated the need for a redesign of the 300 g suspension to improve optic length noise performance. A modification to the auxiliary suspension was also considered to improve the margin to the requirement around the 100 Hz region. The force-displacement transfer functions were modelled to allow a comparison between the model and physical suspension, due to the difficulty of measuring suspension transfer functions directly. This set of outputs considered the optic longitudinal motion resulting from longitudinal, pitch, and yaw driving of the upper mass.

The suspension transfer functions were examined to find the dominant source of optic length noise, found to be vertical-longitudinal coupling of optic motion for both suspensions. This informed the conceptual design of the modifications of both suspensions where a reduction in the magnitude of the vertical base motion to vertical optic motion coupling at 100 Hz was deemed necessary. This was implemented in the 300 g suspension through the addition of steel blade springs at the upper wire mounting points, reducing the frequency of the system's vertical resonances. The auxiliary suspension's much more compact size made the addition of blades difficult – instead, the modified design involved the addition of a rubber block near the top wire mounting point. New suspension models were created for the modified suspensions to predict the transfer functions and system resonances of the updated designs.

A characterisation methodology was developed for the 300 g suspension to provide physical measurements of the unmodified and modified suspension transfer functions. Coil-magnet actuators installed at the upper mass provided a simple way of driving the upper mass in longitudinal, pitch, and yaw motions, while the resulting longitudinal optic motion was sensed with a laser vibrometer. The suspension characterisation demonstrated good agreement with the modelled transfer functions for most cases. The measured force-displacement transfer functions of the unmodified suspension displayed more longitudinal motion from optic pitch and yaw resonances than predicted, thought to be largely due to the difficulty of aligning the vibrometer beam on the optic centre with possible contributions from uneven upper mass drive or small offsets of the wire attachment points.

A methodology was also developed for measurement of the vertical-vertical transfer function between the suspension base and upper mass. A large voice coil actuator was used to drive the suspension base plate in the vertical direction while an accelerometer balanced on the upper mass measured vertical motion. The suspension vertical resonant modes agreed well with the modelling in both the unmodified and modified cases, however measurements of the modified suspension also displayed additional resonance peaks. Investigation of these resonant modes found them to originate from the suspension support structure and from the xylophone resonances of the blade springs. The addition of gusset plates was found to stiffen the suspension structure and shift most structure resonances away from the frequency band of interest. This could be further optimised through adjustments to the support legs and gusset plate placements, and top plate weighting. The xylophone blade spring resonances could be shifted from 158 Hz to 63 Hz by suppression of the blade spring motion with balanced rubber blocks. Development of this work could use coil-magnet actuation on the blade springs to suppress this mode. Shifting the resonance to 50 Hz would give the originally-modelled drop in optic length noise of the modified suspension.

Finally, the optic length noise was calculated using the modelled transfer functions of the unmodified and modified suspensions. This demonstrated the expected decrease in optic length noise resulting from the modification process, and demonstrated that this is expected to meet the experimental noise requirements between 100 Hz – 1 kHz. This work highlights the intricacies of accurate suspension modelling, with the experimental characterisation hugely valueable in the understanding of these suspension systems.

Chapter 4

CRYOGENIC TRIPLE-STAGE SUSPENSION DESIGN

4.1 Suspension Design for the GCIF Cryogenic Cavity

The Glasgow Cryogenic Interferometer Facility aims to take direct measurements of coating thermal noise by observing displacement thermal noise inside a 10 cm cryogenic laser cavity formed by two 1 kg silicon mirrors, and to demonstrate the hanging of a fully crystalline cryogenic suspension. To allow coating thermal noise to dominate this measurement pendulum style suspensions are used to reduce the residual motion of the test mass and isolate from seismic noise. These suspensions must be carefully designed to meet seismic isolation requirements, have low thermal noise, and be suitable for operation at the desired temperature ranges of the experiment: the design of cryogenic, low-noise suspended cavity optics is integral to the success of this experiment. This chapter details the design of such a suspension suitable for use in the GCIF.

Gravitational wave detectors utilise multi-stage pendulums to provide greater seismic isolation to the cavity mirrors than a single-stage pendulum of the same length, as described in Chapter 2. To contribute the least seismic noise into the experimental system it is advantageous to design a pendulum with low-frequency resonant modes such that the coupling of seismic motion to the mirror mass decreases quickly with increasing frequency of motion. This is accomplished through the use of long suspension wires and multiple stages of suspended masses. Despite their length noise advantages multi-stage pendulums add complexity to a design and may introduce additional noise sources such as high mechanical loss jointing regions. The design of a multi-stage pendulum is not just motivated by the experimental noise requirements, but must also factor for the cost, availability, and production limits of materials, as well as laboratory space limits and ease of building. The GCIF cryogenic cavity suspension must have the minimum number of suspended stages necessary to meet the seismic isolation requirements but no more than this to keep the system's cost and complexity down. Section 4.2 outlines the investigation of a multi-stage suspension designed to meet the seismic isolation requirements for the GCIF.

The GCIF test cavity is planned for operation at cryogenic temperatures to reduce the experimental system's thermal noise for a direct measurement of coating thermal noise, and also to allow the research and development of the materials planned for use in 3rd generation gravitational wave detectors. The cavity optics must be closely held at a chosen experimental temperature despite heating from the incident laser beam and ambient heat sources. The GCIF will operate in two phases, firstly running at 123 K and in a later phase at 18 K. Both temperature regimes benefit from the nulling of silicon's thermal expansion coefficient, and therefore are both of interest as operational temperatures. The design of the suspension system must meet the thermal noise and seismic isolation requirements whilst having the capability of maintaining the test mass at these temperatures through a combination of radiative and conductive cooling. Section 4.3 uses finite element analysis to model the heat extraction capabilities of the design from Section 4.2, and builds upon this to outline a design for a suspension suitable for use at cryogenic temperatures. Section 4.4 outlines a full suspension thermal noise calculation using the method outlined in Chapter 2 to ensure that the design meets the experimental sensitivity requirements.

4.1.1 Design Sensitivity Requirements

The comparison between length noise of a single test mass and a single coating is a more intuitive comparison for this chapter than comparing the cryogenic cavity length noise contributions of both elements. The sensitivity requirement is therefore the length noise from the ITM which dominates cavity length noise: $1 \times 10^{-17}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz at 123 K operation, and $4 \times 10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz at 18 K operation with a $1/\sqrt{f}$ dependence on frequency above 100 Hz, as introduced in Chapter 3.

4.2 Seismic Modelling

This section outlines the design work undertaken for the GCIF cryogenic suspension in creating a suspension which meets the seismic noise requirements of the experiment. This work starts with a comparison between the seismic attenuation of a double- or triple-stage pendulum of a length suitable for the cryostat inner shield dimensions in Section 4.2.1. Next, the length noise contribution from a suspension using either fibres or ribbons in the monolithic stage is discussed in Section 4.2.2. Moving forward with the design from Sections 4.2.1 and 4.2.2, the addition of blade springs is examined in Section 4.2.3 with four different designs considered of varying levels of attenuation and complexity. To maintain a monolithic design which includes blade springs in the lower stage the use of silicon blade springs must be considered – Section 4.2.4 presents the finite element analysis of a silicon blade spring. Finally, Section 4.2.5 outlines a conceptual design for a compact seismic isolation platform which could be installed inside the cryostat and add an additional suspended stage for the masses.

4.2.1 Number of Pendulum Stages

The design for the suspension must balance technical simplicity with cryostat height constraints and experimental seismic noise requirements. In the GCIF, the inner shield dimensions of the Leiden Cryogenics cryostat impose a limit on the usable space to a cylinder of 1 m height and 1 m diameter. Based on the availability of crystalline fibre lengths and the 1 m height restriction the logical choices for the suspension were between a double-stage suspension with 45 cm fibre lengths and a triple-stage suspension with 30 cm fibre lengths. Both options were investigated to see if they would meet the noise requirements for the laboratory. For the scope of this study the build of two Mathematica models would have been overly time-consuming. Instead, a direct analytical approach allowed a fast comparison of the seismic attenuation of the two designs. Future sections in this chapter show a detailed seismic study of the cryogenic suspension using a Mathematica model.

Stringent bulk thermal noise and optical absorption requirements for the suspended mirrors necessitate the use of float-zone silicon, which can be manufactured in higher-purity single crystals than silicon ingots produced via the Czochralski method (discussed in Section 2.4.1.1), however current float-zone silicon crystals are limited to a 200 mm diameter [146]. The mass aspect ratio was modelled after the test mass dimensions used at aLIGO,

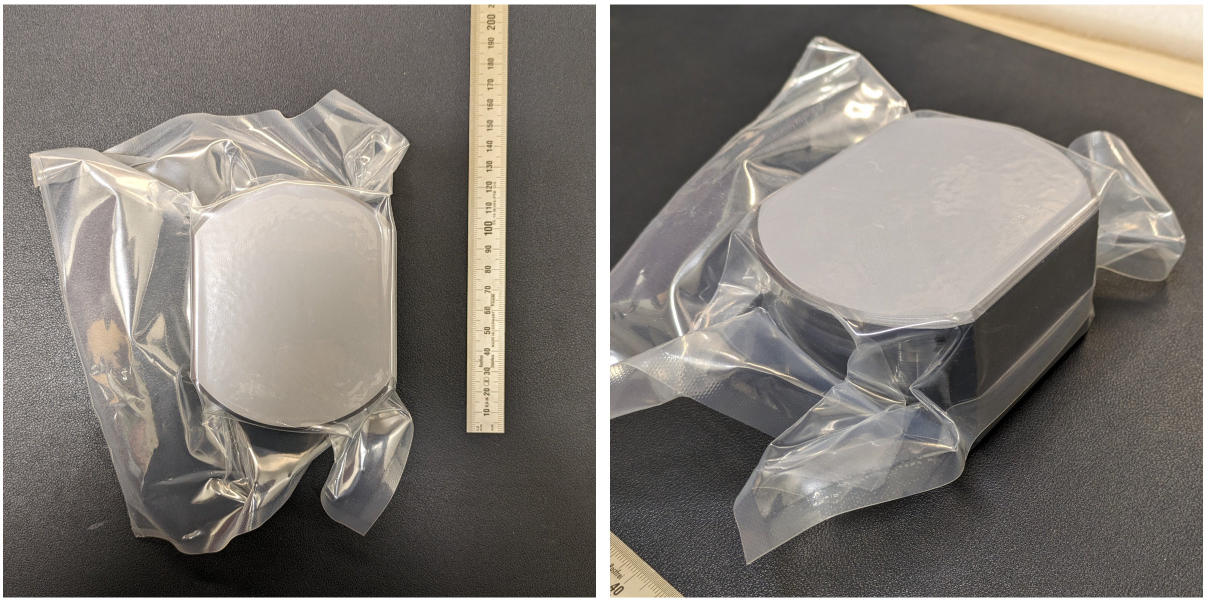

(a) Top view. (b) Isometric view.

Figure 4.1: Float-zone silicon masses for the GCIF cavity suspension.

proven to have good test mass mode visibility [80]. To maintain this aspect ratio while meeting cost and production limitations, a cylinder of 105 mm diameter, 55 mm depth was chosen for the test mass giving a mass of 1108 g before the cutting of bonding flats. The 105 mm diameter face is large enough for the predicted cavity beam spot diameters, and the 1108 g mass allows space for two test chains, two reaction chains, and auxiliary optics to fit within the cryostat to facilitate cavity locking. Polished flats of 55 mm $\times$ 60 mm were cut on each side of the mass to allow future bonding of ears, utilising a design similar to that developed for the fused-silica aLIGO mirror masses. The large flats area will not limit necessary bonding area and will allow the fibre bending point to lie close to the centre of mass of the silicon mass, necessary to ensure mechanical stability of the suspension design. Once material has been removed to produce the flats the final mass is 983 g, shown in Figure 4.1

Data from cryostat manufacturer, Leiden Cryogenics, provides a measurement of seismic noise levels inside a cryostat inner shield as introduced in 3.2.3. This data includes cryocooler noise which is expected to dominate over environmental seismic noise, and so can be used as an estimate of the noise reaching the base of the triple suspension once hanging inside the active cryostat. This data can be seen in Figure 3.5 in Chapter 3, and the horizontal seismic noise data is shown in pink in Figure 4.2. The suspensions developed in this work were thus designed using Leiden Cryogenics's measured seismic data, however it is recommended that the seismic noise levels in the GCIF should be measured after the installation of the cryostat, and the displacement of the mass can be recalculated using these observed levels.

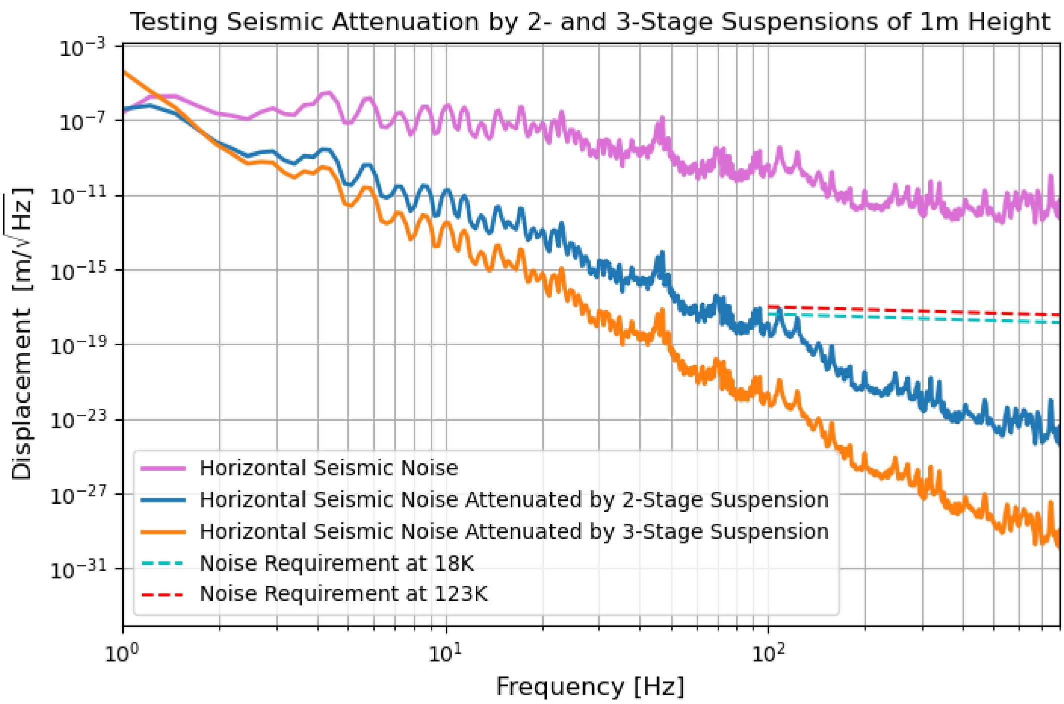

Figure 4.2: Displacement spectrum plotted for horizontal seismic noise measured inside a GCIF-type cryostat (pink), then attenuated by a 2-stage suspension with 50 cm wire lengths (blue) and by a 3-stage suspension with 30 cm wire lengths (orange). Coating length noise requirements shown in dashed red for 123 K and dashed cyan for 18 K.

4.2.1.1 Results from Pendulum Stage Comparison

Chapter 2 covers the derivation of the resonant modes and frequency response function of a simple pendulum. Using Equation 2.1 and Equation 2.6 the longitudinal-to-longitudinal transfer function profile was modelled for a frequency span of 1 Hz – 1 kHz for both a double-stage pendulum with 45 cm fibres, and a triple-stage pendulum with 30 cm wires. Figure 4.2 plots the seismic motion measured in the x-direction (along the axis of the interferometer arm) with the seismic motion in the x-direction attenuated by the double-stage (blue) and triple-stage (orange) suspensions. The longer wire length in the double suspension gave a slightly lower longitudinal resonant mode of $f_{l,double} = 0.74$ Hz compared to the triple suspensions $f_{l,triple} = 0.91$ Hz, however at <1 Hz these values are both outside the frequency span of interest. The double-stage only just meets the desired noise requirement at 100 Hz while the additional $^1/_{f^2}$ attenuation of the triple-stage gave this design an extra three orders of magnitude seismic isolation at this frequency: the lower seismic noise coupling of this design naturally led to the choice of a triple-stage pendulum with 30 cm wire lengths and $\sim$ 1 kg masses moving forward.

4.2.2 Comparing Ribbons and Fibres for the Monolithic Stage

The first monolithic silicon suspension was successfully demonstrated at the University of Glasgow by G. Eddolls [74]. Eddolls experiment showed a 1 kg mass could be hydroxide catalysis bonded to a 250 mm × 2.2 mm × 775 µm silicon ribbon, suspended in a cryostat, and survive a hanging period of over 8 months including a full cryogenic cycle of cooling to 20 K and warming back up to room temperature. As a proven technology for a cryogenic silicon suspension, ribbons were therefore considered in the seismic study for use in the monolithic stage of the GCIF suspension.

Circular cross-section fibres are currently in use at all 2nd generation gravitational wave detectors, with fused-silica fibres at aLIGO and AdV, and sapphire at KAGRA. Fused-silica fibres can be welded to create a monolithic join between fibre and ear, and the fibres can also be shaped in such a way that the bending point of the fibre is far away from this high-loss welded region, reducing the overall thermal noise of the suspension. At the time of writing this kind of shaping has not been investigated experimentally for silicon and the current jointing method is with a hydroxide-catalysis bond. Despite the need for research and development of silicon fibre production and welding, crystalline fibres are promising for creating a low-noise monolithic suspension which led to the consideration of a suspension with four silicon fibres between the intermediate mass and test mass.

For an accurate comparison between the fibre or ribbon monolithic stage it is no longer sufficient to take a direct analytical approach. The cross-sectional area, length, and material are consistent between the ribbon and fibre geometries so only small deviations between the two length noise outputs are expected. The creation of a Mathematica model (as discussed in Chapter 3) for the cryogenic suspension became necessary to allow the effects of bending length, fibre/ribbon stretch, changing of material properties with temperature, and differing material dimensions in the x- and y- axes to be considered in the calculation of the suspension resonant modes. The Mathematica suspension model had been verified in Chapter 3 by comparison of modelled resonant modes to those measured on the suspension and this pre-verified model was used as the basis of the cryogenic suspension model. An extra suspension stage was added and the basic input parameters of masses, wire lengths, and materials were updated to reflect the cryogenic suspension design. Material properties of silicon were added for both room temperature and cryogenic temperatures – these input parameters are discussed in Chapter 2.

In this initial design, two stainless steel blade springs suspend stainless steel wires which, at their lower end, are attached to a 1 kg aluminium upper mass (UM). Four stainless steel wires are attached between the UM and a 1 kg silicon intermediate mass (IM). The experimental design necessitates a cryogenic suspension with a monolithic final stage with 1 kg silicon masses: the intermediate mass and lower 1 kg silicon test mass (TM) make up the monolithic stage. In this model, rectangular cross-section silicon ribbons and circular cross-section silicon fibres are considered for TM suspension in the monolithic stage.

The ribbon width in the y-dimension, w, and thickness in the x-dimension, t, will be determined using the case where ribbon cross-sectional area is limited by silicon breaking stress, σ_0. The cross-sectional area of four ribbons can be found by division of the force from a 1 kg suspension by silicon breaking stress as

$$w \times t = \frac{mg/4}{\sigma_0}. \tag{4.1}$$

The breaking stress of silicon at room temperature is 200 - 300 MPa, measured in tension [178]. Being conservative and using the lower stress value of 200 MPa with a margin for the strength of 3 × the breaking stress (with this margin known as the *safety factor*) this gives a minimum cross-sectional area of $3.68 \times 10^{-8}\,\mathrm{m}^2$. To maintain the same width-to-thickness ratio as used by Eddolls the values of $w = 323\,\mu\mathrm{m}$ and $t = 114\,\mu\mathrm{m}$ were chosen for the model. The resonant modes of this suspension can be seen in Figure 4.3 and Table 4.1.

For simplicity at this stage the silicon fibres are modelled as cylindrical with no shaping and no weld or bond region. The addition of these details are a large design task in themselves, relying on active development of fibre technology. Their inclusion is recommended for future design of the cryogenic suspension and is an active research area in the GW community, but is beyond the scope of this first investigation. The minimum diameter of fibres, d, which can be used to suspend the mass without breaking under the load can be calculated by considering the force from a 1 kg suspension over the area of four cylindrical fibres as

$$d = 2 \times \sqrt{\frac{mg}{4\pi \times \sigma_0}}. \tag{4.2}$$

Considering the same silicon breaking stress as used in Equation 4.1, Equation 4.2 outputs a strength-limited fibre diameter of 216 µm. To model this in Mathematica the fibre dimensions and final stage damping factor calculations were changed from the rectangular cross-section to circular, with the silicon fibre model outputting the resonant modes shown in Table 4.1 and length noise in Figure 4.3. The ribbon or fibre attachment positions on the UM and TM were kept consistent between the models.

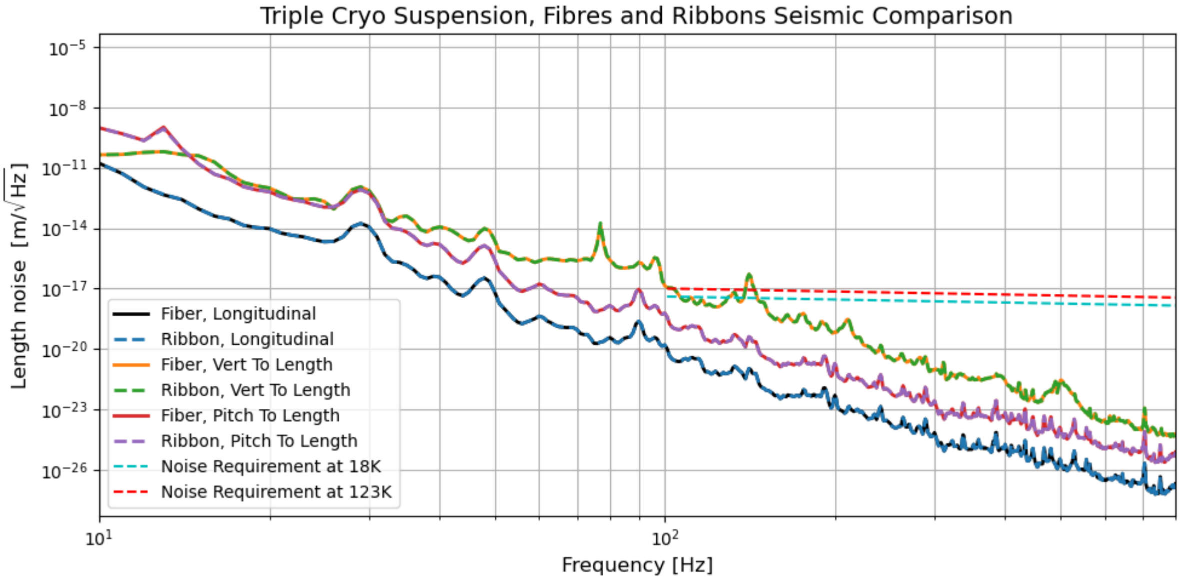

Figure 4.3: Longitudinal length noise plotted against frequency comparing a triple-stage suspension with a fibre final stage and a ribbon final stage. The fibre-type suspension noises are shown by a solid line with black for direct longitudinal coupling, orange for vertical-longitudinally coupled motion, and red for coupling from optic pitch motion. For a ribbon type suspension the lines are dashed and shown in blue, green, and purple for the longitudinal, vertical, and pitch coupling respectively. Coating length noise requirements shown in dashed red for 123 K and dashed cyan for 18 K.

4.2.2.1 Results of Fibre and Ribbon Comparison

Figure 4.3 shows the length noise for the two triple-stage suspension models with four $w = 323\,\mu\text{m}$, $t = 114\,\mu\text{m}$ ribbons, and for four $d = 216\,\mu\text{m}$ fibres in the final stage. The longitudinal length noise is calculated by multiplying the horizontal seismic displacement measured inside the cryostat by the ground-longitudinal to optic-longitudinal suspension transfer function, and is shown in dashed blue (ribbon model) and solid black (fibre model). The optic length noise resulting from the vertical motion first requires the measured vertical seismic displacement to be multiplied by the modelled ground-vertical to optic-vertical suspension transfer function. A 0.1% coupling of this vertical motion to longitudinal motion is assumed to give the length noise, which is shown in dashed green (ribbon model) and solid orange (fibre model). The Mathematica model also outputs the transfer function of ground-longitudinal motion to pitch motion of the optic. Ideal pitch motion with a laser beam perfectly centred on the optic would produce no length noise, however a vertically misaligned beam would couple optic pitch into length noise. This effect is considered by multiplying the horizontal seismic displacement by the longitudinal-to-pitch transfer function. For a 1 mm beam offset the length noise is calculated via trigonometry to give a 2% coupling. This alignment to length effect is shown in Figure 4.3 in dashed purple (ribbon model) and solid red (fibre model). Finally, the longitudinal, vertical, and pitch effects can be combined in quadrature for a root-mean-square (RMS) length noise for each Mathematica model.

It is clear from Figure 4.3 that the longitudinal length noise is almost unchanged when considering fibres or ribbons for use in the monolithic stage. This is expected as the fibre diameter is the same order of magnitude as the ribbon dimension in the x-plane, therefore all longitudinal modes of the suspension, as well as any modes from the unchanged upper stages, should not vary much between the models. These longitudinal modes are observed at 0.64 Hz, 1.46 Hz, and 2.46 Hz in both the ribbon and fibre models.

In the y-plane the ribbon is 2.83$\times$ larger than in the x-plane leading to a stiffness dependent on (x, y) direction - unlike in the fibre case - therefore the non-longitudinal resonant modes of the final stage differ slightly between the two models. The y-direction modes appear at 0.65 Hz, 1.46 Hz, and 2.46 Hz for the ribbon model and 0.65 Hz, 1.47 Hz, 2.46 Hz for the fibre model, showing that this $^{w}/_{t} = 2.83$ ratio is not large enough to push the y-modes a large amount in the frequency space.

Suspension Mode Types	**Ribbon Modes** [Hz]	**Fibre Modes** [Hz]
x-longitudinal	0.64, 1.46, 2.46	0.64, 1.46, 2.46
y-longitudinal	0.64, 1.46, 2.46	0.65, 1.47, 2.46
yaw	0.91, 2.04, 3.33	0.92, 2.05, 3.33
pitch	1.63, 3.75, 13.18	1.77, 3.77, 13.17
vertical	2.11, 15.45, 77.34	2.11, 15.45, 77.16
roll	3.09, 20.05, 93.60	3.14, 20.06, 93.38

Table 4.1: Suspension resonant modes compared for rectangular cross-section silicon ribbons and for circular cross-section silicon fibres. The resonant modes were output from a Mathematica model of the triple-stage suspension and are shown for the six translational and rotational degrees of freedom.

The vertical modes present at 2.11 Hz and 15.45 Hz are from the motion of the UM and IM and therefore are expected to remain the same between the models, which is observed here. The relative vertical motion between the IM and TM is resonant at 77.35 Hz with ribbons and 77.16 Hz with fibres. As shown in Section 2.2.2 the vertical mode is influenced by the vertical spring constant, k_{v}, which varies between ribbons and fibres due to the geometrical dependence of I, the second moment area of the suspending wire. For fibre geometries see Equation 2.20, while for ribbon geometries

$$I_{\mathrm{ribbon}} = \frac{w\,t^3}{12}. \tag{4.3}$$

[58]. Due to the matching of the cross-sectional areas this is not a large influence on the vertical suspension resonances. Note that the peaks in the vertical data at 96 Hz and 140 Hz are present in the seismic data, possibly resonances from the cryocooler plate, and are not suspension resonances.

The pitch modes are predicted at 1.63 Hz, 3.75 Hz, and 13.18 Hz for the ribbon suspension, and 1.77 Hz, 3.77 Hz, and 13.17 Hz for the fibre suspension. The resonance of the pitch modes slightly depend on the vertical spring factor, k_v, shown above to vary between the models. The thicker x-dimension of the fibre produces a slightly higher cross coupling of the 1.63 Hz and 13.18 Hz pitch modes to longitudinal optic motion. At 1.63 Hz the transfer function from longitudinal motion of the ground to longitudinal motion of the optic has a $^\mathrm{m}/_\mathrm{m}$ magnitude of 3 for the fibre and 1 for the ribbon. The 13.18 Hz pitch mode peak has a magnitude of 4×10^{-6} for the fibre compared to 1.1×10^{-6} for the ribbon. This cross-coupling effect is so small that it is not noticeable in the suspension's length noise output. Decreasing ribbon thickness only aids length noise when this dimension is far less than fibre diameter: a ribbon thickness of $t = 10\,\mu\mathrm{m}$ is about 20× thinner than the fibre and only gives a length noise performance increase of 3× at 100 Hz. Despite the 3× performance increase, the ribbon width would need to increase to $w = 3.68\,\mathrm{mm}$ to keep the same cross-sectional area. The brittle nature of crystalline silicon makes a ribbon with this extreme $^w/_t = 368$ ratio likely too difficult to work with.

The tiny deviations in resonant frequency of the Mathematica models do not affect the resultant suspension length noise plots - length noise is indistinguishable between the ribbon and fibre models. Based on seismic isolation performance alone there is no preference between a monolithic stage built with fibres or ribbons. The GCIF triple suspension monolithic stage design will therefore be based on other factors, such as thermal noise performance, engineering difficulty, results of future research and development, or heat extraction performance. In the remainder of this section fibres are modelled in the suspension for their increased thermal noise and heat extraction performance as discussed in Section 4.3.

4.2.3 Number of Blade Spring Stages

Cantilever blade springs are frequently used in pendulum suspensions to provide additional vertical seismic isolation to the suspension chain [15, 66, 160]. A diagram of a blade spring and the use in a suspension can be seen in Figure 3.8. The models in Section 4.2.2 contained one set of two stainless steel blade springs at the mounting point of the wires suspending the UM, but as seen in Figure 4.3 the suspension length noise is close to the experimental noise requirement at 100 Hz. In this section the addition of multiple blade spring stages is explored to reduce length noise in the sensitivity band >100 Hz.

Four models are considered, shown in Figure 4.4. The first model, named '1 Stage', is the same as used in Section 4.2.2 with one stainless steel blade spring attached to each of the steel wires between the suspension frame and UM. In the second model, named '2 Stages', four steel blades are considered in addition to the two modelled for '1 Stage' - these are mounted to the UM, attached to the top end of the four steel wires between the UM and IM. This gives the TM two stages of blade spring isolation through 120 mm long blade springs which can be attached to the suspension. The third model, named '2 Stages LUB' refers to a model with the same two stages of steel blades as the '2 Stages' model, however the two top-stage steel blades have been increased to a 500 mm length (LUB = long upper blades). The 500 mm length is the longest which can fit in the cryostat area, with the longer blades having the advantage of a lower resonant frequency but the disadvantage of making the suspension installation more difficult and reducing the available suspension placements inside the cryostat. The final model, '3 Stages', is identical to the '2 Stages' model with an additional set of four blade springs added at the top of each silicon fibre in the final suspension stage. To maintain a monolithic suspension stage between the IM and TM these blades are assumed to be silicon with further analysis on these blades in Section 4.2.4. The addition of each blade spring stage introduces complexity and cost to the suspension build so a strong justification for the inclusion of each stage is necessary. For this reason the blade spring resonant frequencies were calculated with a direct analytical approach using Equation 3.13 before each of the four designs were separately modelled in Mathematica for an accurate representation of the transfer function of the suspension as a whole.

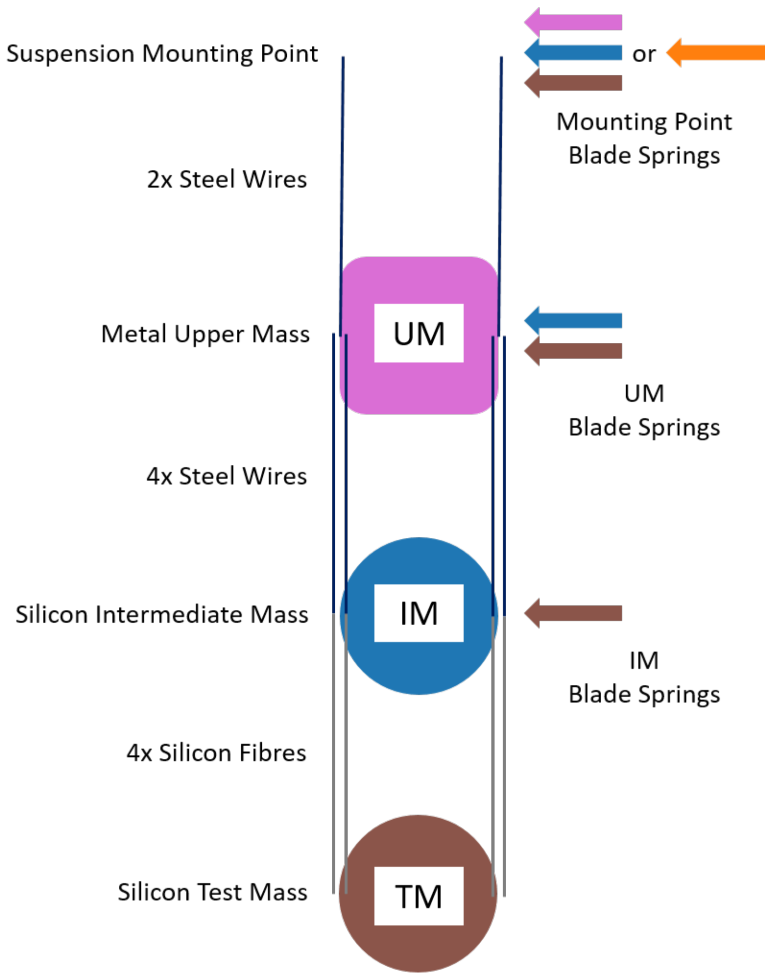

Figure 4.4: Diagram of the cryogenic triple-stage suspension design. The metal UM (pink), silicon IM (blue), and silicon TM (brown) are represented as blocks. Steel wires (navy) sit between the mounting point to UM, and UM to IM. There are silicon fibres (grey) at the monolithic stage between the IM and TM. The modelled blade spring attachment points are represented by arrows: the '1 Stage' attachment point at the pink arrow, the '2 Stages' attachment points at the blue arrows (with the large upper blades at the orange arrow), and the '3 Stages' attachment points at the brown arrows.

4.2.3.1 '1 Stage' Model

The blades modelled in Section 4.2.2 were based on blade springs already in use within the IGR, with a 120 mm length, 1 mm thickness, and an 18 mm to 5 mm taper along the length. With a 1.5 kg load per blade, the bounce frequency is 2.6 Hz. This result is within the <10 Hz resonant frequency aim, chosen to produce a resonance peak well before the experiments sensitivity band and to give a long $^1/_{f^2}$ drop off between the resonance and 100 Hz. For each additional set of blade springs added to the suspension, the length noise at 100 Hz can be improved - the blade springs benefit from a lower vertical spring constant than the wires/fibres alone and thus the vertical resonance modes are brought down in frequency, see Equation 2.8. The suspension length noise for this model is plotted in pink in Figure 4.5.

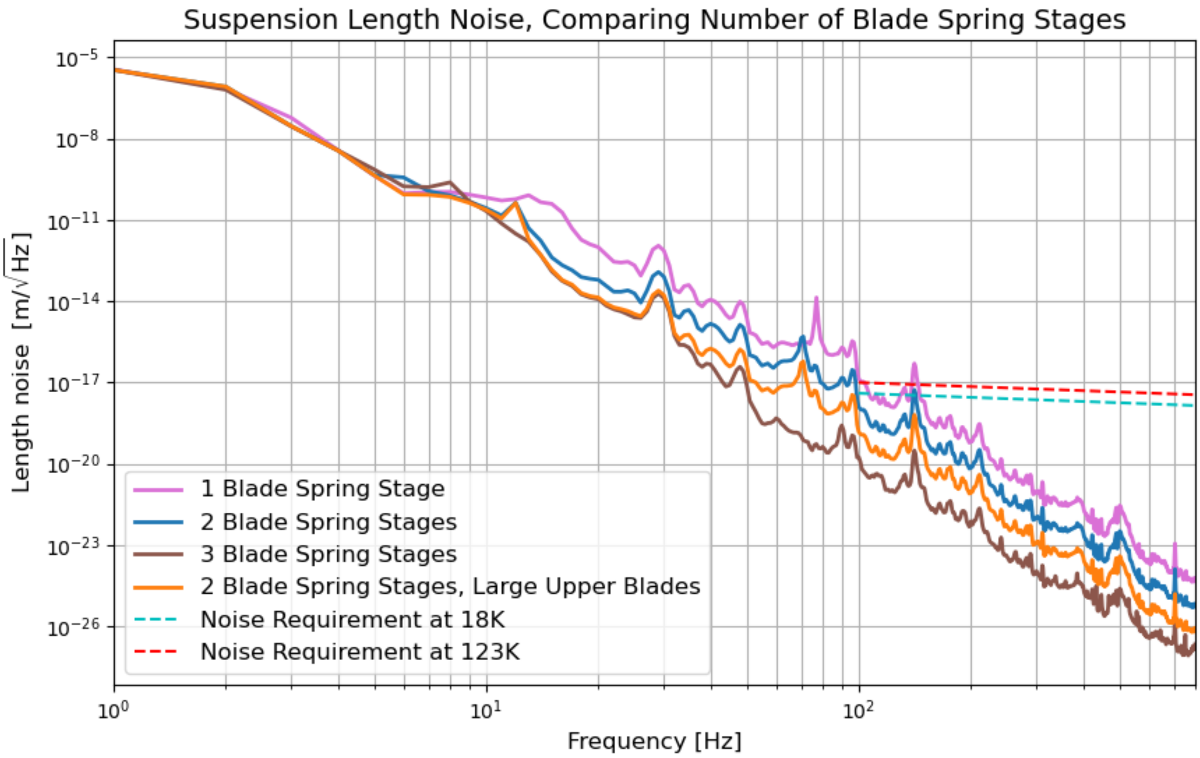

Figure 4.5: Cryogenic suspension length noise comparison of number of blade spring stages, with the resulting optic longitudinal motion with a single 120 mm steel blade spring stage (pink), two stages of 120 mm steel blade springs (blue), three stages of blade springs using silicon blades for the lower stage (orange), and two stages of steel blade springs with upper blades of length 500 mm (brown). Coating length noise requirements shown in dashed red for 123 K and dashed cyan for 18 K.

4.2.3.2 ‘2 Stages’ Model

For an additional four steel blades at the UM it was calculated that an increase in the number of blades coupled with the decrease in suspended load would result in a 4.7 Hz bounce frequency if using the same size of blade spring as above. Although 4.7 Hz is not unreasonable, the blade parameters were tweaked to a 120 mm length, 0.8 mm thickness, and 16 mm to 5 mm taper which resulted in a slightly lower 3.2 Hz resonance with a blade still of a suitable size. The blades at the UM were considered in addition to the two blades at the upper attachment point and so decreased the length noise of the suspension by one order of magnitude at 100 Hz when compared to the model with one blade spring stage. This is shown by the blue line in Figure 4.5.

4.2.3.3 '2 Stages Large Upper Blade' Model

At this stage in the design process, the seismic isolation system for the triple suspension is not yet known. Blade spring stages which are directly attached to the suspension or a seismic isolation platform installed inside the cryostat are both still viable options. The double-stage blade spring model of 120 mm blades (blue line) would be ideal in a compact system in which the triple suspension is attached to an independent platform inside the cryostat, such as the one described in Section 4.2.5. Despite the fact that a compact mirror suspension system allows for convenient reconfiguration of the optical set-up, lengthening of the blade springs pushes the vertical resonance of that spring stage down and therefore a model was created to explore the effect of maximising the length of the two steel upper-stage blades within the space allowed by the 1 m diameter cryostat. The blades were lengthened to 500 mm with a clamp-end width of 150 mm which tapers to 5 mm, and a 1 mm thickness. The blades between the UM and IM were kept at their original size due to the feasibility of adding long blades to a suspended mass which is only 100 mm in diameter. The long upper blades had a calculated bounce frequency of $f_s =$ 0.82 Hz and therefore are advantageous over the 120 mm length, $f_s = 4.69$ Hz design. The effect of the long upper blades are shown in orange in Figure 4.5, demonstrating a length noise between that of the 120 mm length double blade-stages and the triple-stage blades incorporating the silicon blades (discussed in Section 4.2.3.4). The length noise is below the design requirement >100 Hz while avoiding the build complications of silicon blades and is therefore an attractive option for the mechanical design of the suspension.

4.2.3.4 '3 Stages' Model

To test the maximum number of blade spring stages on the suspension a model was created which contained the 6 steel blade springs as above, with an additional four silicon blade springs at the IM. The silicon blades were included to give additional vertical isolation while keeping the final suspension stage entirely monolithic. The advantages of this would be a reduction in suspension thermal noise, and possible reduction of crackling noise as compared to steel blades [29], but the inclusion of a silicon blade spring in a suspension would be demonstration of a new technology and there would be significant design difficulties to overcome. Non-steel blade springs are a viable option in gravitational wave detector suspension design, with sapphire blade springs currently in use at the KAGRA detector, however making the shift to silicon blades would require research into bonding of the blades to the fibres and masses, and ensuring suitable heat flow from the TM through the blades and bonds.

4.2.3.5 Results of Blade Spring Investigation

The output transfer functions of the suspension models with one, two, and three stages of blade spring isolation could be combined with a seismic noise spectrum to produce the total length noise of each suspension. Length noise was considered from longitudinal motion, from a vertical to longitudinal coupling of 0.1%, and from a pitch to longitudinal coupling of 2%. These length noise sources were added in quadrature and plotted in Figure 4.5. The total length noise for the triple blade system is 840× lower than the 123 K noise requirement at 100 Hz. The double blade design with the large upper blades meets the laboratory's length noise requirement between 100 – 1000 Hz with a level of $1.25 \times 10^{-19}\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz and has the advantage of having a much more feasible design than the triple blade system. The double blade system with 120 mm upper blades has worse isolation performance than the 500 mm upper blades due to the increase in f_s of the shorter blades, however this compact design may fit better in combination with a seismic isolation platform and therefore should not be discarded at this stage. Without an isolation platform this design still meets the length noise requirements between 100 – 1000 Hz aside from the peak at 140 Hz. At 100 Hz the length noise of this model is $1 \times 10^{-18}\mathrm{m}/\sqrt{\mathrm{Hz}}$. The single blade design is not recommended as it does not meet the length noise requirements until 148 Hz and can be easily improved with an additional stage of steel blade springs which does not add too much complexity to the build. Currently, no monolithic silicon blade spring suspensions exist therefore the building of this suspension will depend upon future research and development, or use of a seismic isolation platform may be necessary in the cryostat.

4.2.4 Finite Element Analysis of Silicon Blade Springs

Section 4.2.3 investigated the effect of multiple blade spring levels on the seismic attenuation of the cryogenic suspension. The '3 Stages' model demonstrated the lowest length noise of the four models at frequencies >3 Hz but necessitated the use of silicon blade springs to maintain a monolithic final stage. Stainless steel blade springs have been previously used for room temperature suspensions in projects within the University of Glasgow, with the use of these blades well documented and understood. Silicon blades, however, would present a new challenge. Up to the point of writing, silicon blade springs have been designed, manufactured, and undergone testing of breaking stress, stress under load, and resonant frequency, but have not previously been used in the hanging of a suspension [95]. It was therefore necessary to explore the resonant modes and loaded stress of a silicon blade through finite element analysis.

At the University of Glasgow, a 2018 project tested different geometries of silicon blades which were laser cut from a 1 mm thick wafer, aiming to produce blades with a high breaking stress while keeping a low bounce frequency [95]. Triangular cantilever shaped blades consistently demonstrated uncoupled frequencies of around 10 Hz with a 100 g load (value dependent on blade geometry) while also having the highest breaking stresses of all the designs, with a highest demonstrated breaking stress of 162.6 MPa. Other tested blade shapes included various double- and triple-folded blades which aim to increase effective blade length while keeping the overall blade size compact. The improved breaking stress of the cantilever design was expected for a couple of reasons. Firstly, the decreasing width of the blade toward the end allows an even distribution of stress over the blade area and therefore a decrease in likelihood of breaking under a given load [98]. Secondly, the longer perimeter of the folded blades increases the area damaged by laser cutting, with silicon strength being highly dependent on surface quality [63]. For these reasons a triangular cantilever blade geometry was considered for use in the cryogenic triple suspension.

4.2.4.1 Results from Finite Element Analysis of Silicon Blade Springs

Equation 3.13 was used to calculate the resonant frequency of blade springs with different input parameters (blade length, thickness, wide end width, narrow end width). A 120 mm length, 18 mm to 5 mm width and 1 mm thickness silicon blade (matching the strongest geometry used in the 2018 experiment) was found to produce a bounce frequency of 5.53 Hz with a 250 g load while remaining small enough for feasible use in the cryogenic suspension. Following the method from Section 3.3.1.2 in Chapter 3, a computer-aided design model of this blade was created in Solidworks 2023 and imported into COMSOL for finite element analysis using the Solid Mechanics interface. In COMSOL the clamping of the blade was simulated by adding Fixed Constraints to the relevant boundaries and the 250 g load was added as a mass per unit length along the tip of the bottom edge of the blade. The finite element mesh size was chosen to balance short computation time with model accuracy. COMSOL calculated the bounce frequency as 5.35 Hz and a maximum stress of 120 MPa, with the output shown in Figure 4.6. The resonant frequency output is a 3% deviation from the analytical calculation which demonstrates good agreement between the methods. The small difference in values result from the simplicity of the analytical calculation in comparison to the COMSOL model, where the FEA includes bending point variations due to clamping and the attached load, and also calculates the higher frequency resonant modes of the blade such as lateral movement of the blade tip and twisting around the x-axis.

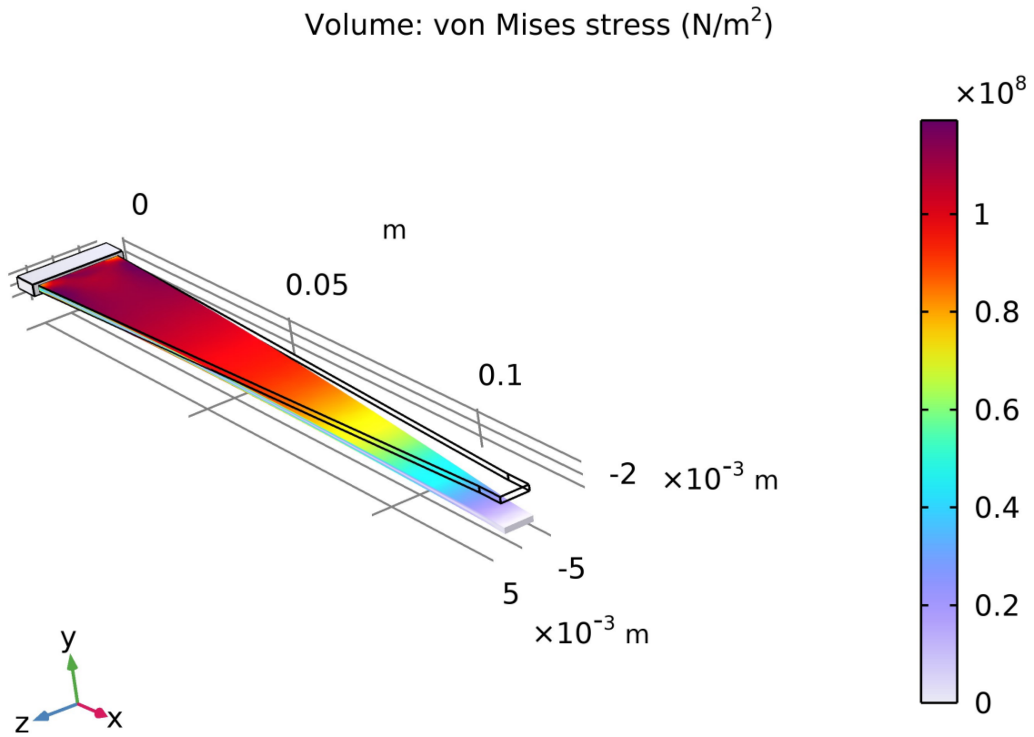

Figure 4.6: Image of silicon blade spring FEA model showing blade tip displacement for the 5.35Hz resonant frequency and material stress distribution.

The 5.35 Hz value was input to the Mathematica model as the bounce frequency of each of the four silicon blades. The four silicon blades were in addition to the steel blades at the mounting point and at UM discussed above, giving a total of three stages of blade spring isolation. From this the transfer function and length noise of a suspension with triple blade stages was found. Shown in Figure 4.5, the additional isolation stage decreases the suspensions length noise by a factor of 87 at 100 Hz compared to the '2 Stages' case.

Despite the obvious performance increase of three blade spring stages a silicon blade spring has not yet been built into a suspension. Research into silicon springs is still at early stages and a suitable blade with low bounce frequency and a suitable strength has not yet been found. The breaking stress of the silicon blades tested by R. Grevers does not allow the hanging of a 250 g load with a suitable safety factor – a loaded stress of 120 MPa with a measured maximum breaking stress of 162.6 MPa has a safety factor of only 1.36. In fact, many measured silicon blade springs broke with stresses below 120 MPa. In future, this could be possible through surface treatment of blades giving an increase to strength, and the analysis of other blade geometries, but that work is outwith the scope of this thesis.

4.2.5 Seismic Isolation Platform

Conceptual design work was undertaken for a seismic isolation system for the cryostat which could add additional vertical isolation at, or before, the suspension top-stage, eliminating the need for silicon blades while maintaining three stages of isolation. A seismic isolation platform would consist of a larger suspended mass which multiple lower elements could be hung from, thus saving space at the top of the suspension for longer blade springs. Seismic isolation is considered as the main design driver here, however common mode rejection for the cavity optics is another possible advantage of such a design. The cryostat chamber will house the 10 cm suspended cryogenic cavity and therefore must contain the two triple-stage suspensions for the cavity mirrors as well as the reaction chain for each TM suspension. Space constraints in the cryostat must therefore be carefully considered in the design of a seismic isolation platform as a shared upper stage.

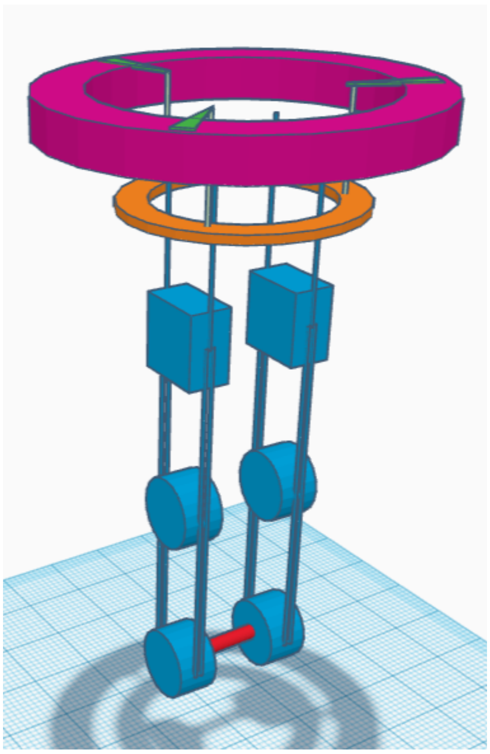

Figure 4.7: Conceptual design of possible set-up of a seismic isolation platform for the cryogenic cavity. Pink shows the outer platform as the mounting point for three (green) blade springs. The blade spring wires attach to the orange inner ring platform which is the mounting point for the suspension chains. The attachment point between the inner platform and suspension top wire is not shown. The two blue suspension chains form the cryogenic cavity, with the cavity laser beam depicted in red.

The first concept was a seismic isolation platform which would be built in to the cryostat and hung from the top plate. This design was based on the rubber and aluminium extrusion isolation stacks built in to the vacuum chambers previously used in the GCIF, and aimed to create a suspended stable optical platform from which the suspension could be built. Space constraints in the cryostat and the top-up cryostat opening mechanism make this a challenge since the platform is unable to be built from the floor as was previously done in the room temperature vacuum chambers. These constraints led to a ring-shaped outer platform which would be bolted to the cryostat top plate via beams (or could be suspended from blade springs or rubber). Blade springs would be installed on this platform with either three blades pointing radially inward or six blades in a T-pose arrangement. The blade springs would hold an inner platform ring from which the suspension would be hung, with a conceptual design shown in Figure 4.7. With this seismic platform an additional two stages of blade spring isolation could be introduced giving an additional $1/f^4$ drop off of vertical seismic motion coupling to the TM. This type of nested isolation system with inverted pendula is neat and compact, and provides potential for greater flexibility in general for the number of suspension stages, however this type of system can be hard to stabilise.

To test the feasibility of this concept, a CAD assembly was created in Solidworks to check that the design fitted in the available cryostat space. This gave a rough idea of the size and shape of the platforms, the heights at which they could be hung, and the blade spring dimensions which could be used. The outer platform was modelled as an aluminium hexagonal ring with 0.25 m sides. The three stainless steel blades are 90 mm long with a 20 mm – 5 mm taper, 1 mm thickness and a bounce frequency of 3 Hz. The inner platform is an aluminium ring with a 0.3 m outer diameter, 0.2 m inner diameter, and 0.01 m thickness and is suspended 45 mm below the outer platform. The suspension mounting points were set 50 mm above the inner platform to create the inverted pendulum stage.

Using the dimensions output from the CAD, a toy model of a double pendulum with a suspended platform was built in Mathematica which allowed exploration of the stability of this type of set-up. Mathematica modelling is powerful enough to calculate the concurrent motions of the suspension masses and the two platforms, including the inverted pendulum created by the suspended platform, and can produce realistic results of the stability of the set-up.

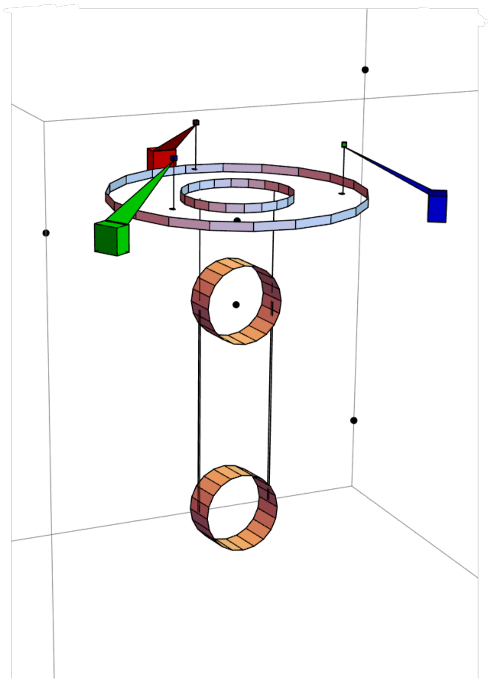

Figure 4.8: Output graphics from the Mathematica dynamical model of a suspended ring platform as the attachment point for a double suspension. This dynamical model was used to check the stability of such a set-up.

4.2.5.1 Results of Seismic Isolation Platform Modelling

The simulation produced realistic resonant frequencies of 0.6 - 32 Hz for the structure, with no complex outputs which would indicate instability in the design. The frequency range is believable, with the first three modes (0.61 Hz, 0.65 Hz, and 0.66 Hz) being the pitch, roll, and vertical modes of the ring platform, and the higher frequency modes being from the double suspension which are still present if the blade springs are immobilised. The set-up was found to be stable when the centre of mass (COM) of the suspension UM is below the COM of the ring platform. With the dimensions described above, the set-up is stable with the suspension mounting points sitting between the ring platform level and up to 0.15 m above the platform. The model predicts the set-up to topple if the suspension mounting point is higher than 0.15 m above the platform due to the UM COM approaching the platform to create an unbalanced inverted pendulum. The limit of this stability height depends on the masses of the suspension and platform, the distribution of mass around the platform, and the length of suspension wires, which are parameters which will evolve as the design of the triple suspension matures. This work demonstrated that a suspension chain could be successfully hung from an inverted pendulum suspended platform in a compact and stable set-up and provides a model from which constraints can be calculated for future mechanical design, such as tipping points, COM tolerance and wire attachment points.

4.2.6 Conclusions on Seismic Modelling

In this section of work a suspension was designed for the cryogenic cavity of the GCIF with the aim of meeting seismic isolation requirements for the laboratory. Space constraints and the availability of cryogenically compatible optic material necessitated a suspension of a 1 m total height with 1 kg optics. A triple suspension design was chosen, giving improved noise performance in comparison with a double suspension design. The seismic attenuation from a strength-limited silicon fibre was compared with a strength-limited silicon ribbon for the suspension of the final stage, producing nearly identical length noise outputs. The down selection of the fibre geometry for the suspension design was decided by thermal noise and heat extraction benefits, discussed in Section 4.3. Mathematica suspension models were created to compare the effect of adding blade springs in four different configurations. The '3 Stages' model produced the highest level of seismic attenuation, however the '2 Stages LUB' design is recommended for use in the laboratory for the optimal balance of attenuation performance and build feasibility. The silicon blade springs necessary for the '3 Stages' model could be an addition in a future iteration of the experiment after validation of the technology. A design was proposed for a compact isolation platform to further increase seismic attenuation at the TM. A nested pendulum configuration provides the most isolation within the space constraints imposed by the cryostat and was proven as a stable set-up via Mathematica modelling, however further development of this concept will be necessary once the suspension top stage design is more mature.

4.3 Heat Extraction Modelling

For 3G gravitational wave detectors operating in vacuum at cryogenic temperatures the cavity optics will need to be constantly held at a specified temperature while being heated by the incident laser beam. To avoid direct coupling of seismic noise to the TM there can be no heat links or heating elements directly attached to the TM: for a suspension with a monolithic final stage the temperature of the TM can only be varied through radiative heating/cooling, or through conduction along the suspension chain. It is clear that the crystalline optical materials must be chosen to have high thermal conductivity, both to conduct heat away from the TM to the higher stages of the suspension chain, and to prevent heating of the TM surface leading to warping through local regions of thermal expansion or contraction.

Silicon has a characteristic zero crossing in thermal expansion at both 123 K and <18 K therefore these are the operational temperatures being considered for 3G detectors with silicon optics. Multi-material suspension designs are being discussed which provide the possibility of combining a silicon optic (with the advantages of the availability of large, high-purity crystals and the thermoelastic nulling) with sapphire fibres (with advantages of the more advanced welding technology, and their previously demonstrated use at KAGRA). In this section, a monolithic silicon suspension is considered: the domination of radiative test mass cooling at 123 K results in the fibre material having little effect in this temperature regime, while in the 18 K temperature case the lower thermal conductivity of silicon (when compared to sapphire) makes a silicon fibre the limiting case for the GCIF. Additionally the heat extraction of sapphire fibres is already well studied within the KAGRA community [117]. The thermal noise of a silicon and sapphire multi-material suspension is considered in Section 4.4.

In a vacuum system thermal energy can only be transferred via thermal radiation and thermal conduction as there is no air for convective heat transport. Thermal radiation is described by the Stefan-Boltzmann Law where the radiant flux from an object is proportional to temperature as

$$q_{\text{radiation}} = \epsilon\,\sigma\,T^4 \tag{4.4}$$

for material emissivity ϵ and Stefan-Boltzmann constant σ [217]. Fourier's Law describes the heat flux conducted through a material as

$$q_{\text{conduction}} = -\kappa\,\Delta T \tag{4.5}$$

with κ the material's thermal conductivity and ΔT the temperature gradient across the material in K/m [217]. The high silicon thermal conductivity at 18 K combined with the low temperature means thermal conductivity is the dominant method of heat transfer in this temperature regime, whereas at 123 K heat transfer will be via both radiative and conductive methods. These two temperature scenarios were therefore separately modelled to explore how to effectively cool the TM within the two heat transfer regimes. This was done with finite element analysis, using the Heat Transfer Module in COMSOL which allows simulation of heat transfer mechanisms in materials and free space, and was used to model the temperature of the test mass in both stationary and time-dependent conditions [56]. To further the suspension design, it is important to understand how to cool the test mass to the desired temperatures, to map the thermal gradients present, and to calculate the thermal noise present in the suspension.

4.3.1 Consideration of Ribbons

Ribbon and fibre geometries were both initially considered for suspension of the TM from the UM, with a seismic study performed for both cases (Section 4.2.2) and displaying no significant difference to seismic isolation at the optic. Initially a ribbon geometry seems to have a potential advantage due to the small dimension in the x-direction for low seismic coupling while maintaining a large cross-sectional area for heat transport. Throughout the course of this work ongoing research into crystalline ribbons and fibres resulted in a preference for the fibre geometry. Publication of [74] demonstrated the loss of strength caused by the manufacture of silicon ribbons. Current ribbon samples are laser cut from a silicon wafer and damage from the laser results in a poor surface quality which impacts both the sample's strength and thermal conductivity. Treatments improving surface quality were shown to greatly increase material strength and may improve thermal conductivity but need further development before being a truly viable option [74]. Laser cutting ribbons from wafers also adds difficulty to neck shaping, causing a higher energy distribution in the bonding region and resulting in a suspension thermal noise dominated by HC bond loss as observed in [74]. Additionally, LIGO originally considered fused-silica ribbons for their 'advanced' upgrade however later moved to fused-silica fibres when initial promising gains in ribbon dilution factor proved too difficult to manufacture, with the transition between the central region and neck being the main issue [58].

Crystalline fibre development advancements (as outlined in Section 2.4.1.1) increased confidence that fibres of the necessary lengths and diameters would potentially be produced in time for 3G detectors and led to the decision to model fibres in this work. Research into crystalline materials for cryogenic 3G GW detectors is still in the early stages and future advancements will drive the specific design concepts for these detectors, particularly around the fibre, neck, and ear jointing regions.

4.3.2 Cryostat Heat Loads

A cryostat is not a perfectly closed thermal system - there will be various heat sources impacting the final temperature of the suspension chain. Firstly, the initial test mass will absorb some amount of the light which travels through the optic to reach the cavity, P_{input}, as a factor of the substrate absorption, α_{TM}, and the optic thickness, t_{TM}. A second term gives the heating dependent upon the laser power resonating in the cavity, and the absorption of the coating on the TM. The maximum heat deposited into the TM from the incident laser is therefore

$$P_{\text{laser}} = P_{\text{input}}\, \alpha_{\text{TM}}\, \text{t}_{\text{TM}} + P_{\text{cav}}\, \alpha_{\text{C}} \tag{4.6}$$

for P_{input} the incident laser power on the initial TM, assumed to be 100 mW, α_{TM} silicon absorption taken as the conservative value of 300 ppm/cm [67, 68], $\text{t}_{\text{TM}} = 55$ mm the test mass thickness, P_{cav} the laser power in the resonant cavity, assumed to be 200 W, and α_{C} the coating absorption, assumed to be 10 ppm [8]. This gives a deposited power of $P_{\text{laser}} = 2$ mW which is dominated by the coating absorption term.

The laser beam will need to enter the cryogenic cavity from the 10 m laser stabilisation cavity, which sits at room temperature. To allow this there is a 25 mm radius viewport to the cryostat's inner shield which effectively acts as a disk of 300 K radiatively heating the TM. In cryogenic systems such as these, a cold shield, or 'snout' is often used. This snout is a cooled cylinder which extends from the cryostat inner shield into the room temperature arm, lessening the heating on the TM by reducing the solid angle of the 300 K disk that has a direct line-of-sight to the TM. The heat load on the TM from this viewport, P_{viewport}, can be calculated as in [8] where

$$P_{\text{viewport}} = \sigma\, T^4\, \pi\, r_{\text{TM}}^2\, \frac{\pi\, r_{\text{snout}}^2}{4\, \pi\, L_{\text{snout}}^2}. \tag{4.7}$$

T is the 300 K of the viewport, and r_{snout} and L_{snout} are the radius and length of the cold snout. In the cryostat configuration planned for the GCIF the cold snout length is limited by the physical constraints of the short section of beam tube between the cryostat and steering mirror tank (shown in Figure 1.11). A short snout of $L_{\text{snout}} = 200$ mm with radius equal to the cryostat viewport radius of $r_{\text{snout}} = 25$ mm gives a heat load of $P_{\text{viewport}} = 14$ mW on the TM face.

Finally in this model an ambient heat load is considered from unknown sources such as heat radiation from wiring, imperfections in the cooling system, heating through ports, or self-heating of electronics inside the cryostat. This type of heat load is specific to each experimental system and will be measured at the GCIF after the installation of the suspension system inside the cryostat. The ambient heat load measured at the University of Glasgow inside a similar Leiden Cryogenics cryostat was estimated to be $P_{\text{ambient}} = 0.1$ mW, with this value being used in this study [74].

Heat Source	**Symbol**	**Value** [mW]
laser	P_{laser}	2
viewport	P_{viewport}	14
ambient	P_{ambient}	0.1

Table 4.2: Considered cryostat heat loads.

4.3.3 Phonon-Boundary Scattering

In the upper metal stages heat is conducted through free electrons, however heat conduction in crystalline silicon is dominated by phonons - collective excitations of atoms in a lattice formation. Phonons travelling through a material can scatter from impurities or crystal defects, material boundaries, electrons, or other phonons, all of which reduce heat transfer through that material. In a monolithic suspension chain the silicon fibre diameter must be kept small to ensure a non-rigid link between the masses to reduce seismic coupling, but when this dimension becomes comparable to the phonon mean-free-path there can be significant phonon-boundary scattering which can severely impact thermal conductivity. Surface quality is also a factor in the severity of this reduction, where a smoother surface gives less scattering of phonons, leading to a higher rate of heat transport [194]. Treatments such as chemomechanical polishing, argon ion etching and wet chemical etching have been shown to increase silicon tensile stress through the removal of surface damage and could possibly be utilised in suspension fibre production [74]. To maintain a high thermal conductivity 3G detector design will utilise high-purity silicon to reduce impurities and defects, with a silicon surface treatment recommended to reduce the effects of boundary scattering.

Phonon mean-free-path, l_{m}, can be calculated using the temperature dependent thermal conductivity, $\kappa_{\mathrm{Si}}(T)$, the volumetric specific heat capacity, $c_V(T)$, and material sound velocity, v_{s} as

$$l_{\mathrm{m}} = \frac{3\,\kappa_{\mathrm{Si}}(T)}{c_V(T)\;v_{\mathrm{s}}} \tag{4.8}$$

[110]. At 123 K the specific heat capacity is $c_V(123) = 338\ \mathrm{J\,Kg^{-1}\,K^{-1}}$, the thermal conductivity is $\kappa_{\mathrm{Si}}(T) = 667\ \mathrm{W\,m^{-1}\,K^{-1}}$, while the sound velocity in silicon is $v_{\mathrm{s}} = 5660\,\mathrm{ms^{-1}}$ [1, 108]. The resulting mean-free-path is $l_{\mathrm{m}} = 1.05 \times 10^{-3}\,\mathrm{m}$, meaning that at 123 K any silicon samples with dimensions of the order of a millimetre must account for phonon-boundary scattering. At 18 K the specific heat capacity, $c_V(18) = 1.52\ \mathrm{J\,Kg^{-1}\,K^{-1}}$, and the thermal conductivity, $\kappa_{\mathrm{Si}}(18) = 4778\ \mathrm{W\,m^{-1}\,K^{-1}}$, produce a much larger phonon mean-free-path of $l_{\mathrm{m}} = 1.67\,\mathrm{m}$, showing that phonon-boundary scattering must be accounted for in the silicon elements across the entire suspension.

There is insufficient experimental data showing the phonon-boundary effect on silicon samples of different thicknesses. To model this effect with the current lack of data the recommended silicon thermal conductivity values from [194] were multiplied by a scaling factor of $t_{\text{fibre}}/t_{\mathrm{T}}$, for t_{fibre} the minimum dimension of the silicon fibre (here, the diameter of 216 µm) and t_{T} the thickness of the sample in [194] which replicated most closely the thermal conductivity of the recommended values. This scaling factor method of modelling phonon-boundary scattering was taken from [63].

4.3.4 Cooling to 123K

A first step in exploring cryogenic interferometry in the Glasgow 10 m facility will be cooling silicon suspensions to 123 K, with a future aim of eventually cooling to 18 K. The cooling of the test mass to 123 K was the first step of the thermal modelling process. The domination of radiative heat transfer at this temperature allows for a simpler thermal model which can be more easily compared to analytical calculations for verification of the FEA thermal model. The FEA model is then expanded by addition of the intermediate mass and fibres of the monolithic stage, with the effects of hydroxide-catalysis bonding, phonon-boundary scattering, and TM heat loads considered.

4.3.4.1 Consideration of a Suspended Test Mass

With the basic design of the suspension completed the number of pendulum stages, the shape of the masses, and the length of the wires are known parameters. Solidworks was used to build a CAD model of the suspension and cryostat. The CAD drawing was useful in exploring the layout and suspension placement inside the cryostat chamber however its main use was for importing the suspension parameters into COMSOL for thermal modelling with finite element analysis. In finite element analysis an object is ‘meshed’ by splitting the object into a series of finite elements, of a size defined by the user. For each of these elements a physical property can be calculated, with the results summed over the elements to give a result for the object as a whole.

For the initial verification of the FEA model a simple case of radiative heat transfer was considered with a cylindrical silicon mass suspended inside the inner cryostat chamber. A CAD file was made for the 1 m diameter, 1 m height cryostat chamber and 96.8 mm diameter, 57 mm depth silicon mass. This file was imported into COMSOL for use with the 'Heat Transfer in Solids' package. The cryostat inner chamber material was set as a copper cylinder, with the material imported from the default COMSOL material library. This chamber will have a gold coating on the inner surface with an emissivity set to $\epsilon_c = 0.05$. Emissivity was assumed to be temperature independent to match the analytical calculation in Equation 4.10. The temperature of the inner chamber was set to 7 K as per the specifications of the cryostat. The silicon material properties are discussed in-depth in Chapter 2. A custom COMSOL material was created of 111 anisotropic silicon and was used here to model the mass. This material had temperature-dependent properties input for thermal conductivity, coefficient of thermal expansion, and heat capacity at constant pressure. For this study a heat flux was defined over the whole surface of the silicon mass to model the heating from both the absorption of the incident laser, and from some heating element which will be used to maintain a constant TM temperature of 123 K. A Surface-to-Surface Radiation was chosen for the 'Physics' study to model the net radiation exchange between the mass and cryostat. As part of this study the area outside the cryostat chamber was set as an opaque, infinite void so there is no loss of heat to the area outside the inner chamber's surface. The finite element mesh size was chosen to balance short computation time with accuracy of results. A Stationary Study was computed to find the steady-state solution to the net radiative heat transfer between the cryostat and heated mass.

The analytical calculation of heat transfer between the cryostat chamber and suspended mass is as follows. In the case of a heated test mass suspended in a vacuum chamber cryostat with an incident laser beam on one surface, the steady-state heat transfer can be approximated as

$$P_{\text{radiation}} = P_{\text{laser}} + P_{\text{heater}} \tag{4.9}$$

for $P_{\text{radiation}}$ the net radiation exchange between the TM and cryostat (including heat radiated from the TM, and reflection of radiation of the cryostat), P_{laser} the heat deposited into the TM from the laser, and P_{heater} some heating element which is necessary to keep the TM at the desired temperature. To calculate net heat transfer between the TM and cryostat it is assumed that the TM surface area is much smaller than the cryostat surface area, and therefore the TM can be treated as a small body in a large enclosure. With this assumption, the radiation exchange between two grey surfaces (a body with non-perfect absorption which is temperature and wavelength independent) is written as

$$P_{\text{radiation}} = \frac{\sigma\, A_{\text{TM}}\, (T_{\text{TM}}^4 - T_{\text{c}}^4)}{\frac{1}{\epsilon_{\text{TM}}} + \left(\frac{1}{\epsilon_{\text{c}}} - 1\right)\left(\frac{A_{\text{TM}}}{A_{\text{c}}}\right)} \tag{4.10}$$

[116]. For $A_{\mathrm{TM}} = 0.032\ \mathrm{m}^2$ surface area of the test mass, $A_{\mathrm{c}} = 4.712\ \mathrm{m}^2$ surface area of the cryostat, $T_{\mathrm{TM}} = 123\,\mathrm{K}$ desired test mass temperature, $T_{\mathrm{c}} = 7\,\mathrm{K}$ the temperature of the cryostat's inner shield, $\epsilon_{\mathrm{TM}} = 0.5$ the emissivity of the silicon test mass, and $\epsilon_{\mathrm{c}} = 0.05$ the emissivity of the cryostat. For the desired TM temperature of 123 K the required P_{heater} can be found by combining 4.9 and 4.10 for

$$P_{\mathrm{heater}} = \frac{\sigma\, A_{\mathrm{TM}}\,(T_{\mathrm{TM}}^4 - T_{\mathrm{c}}^4)}{\frac{1}{\epsilon_{\mathrm{TM}}} + \left(\frac{1}{\epsilon_{\mathrm{cryo}}} - 1\right)\left(\frac{A_{\mathrm{TM}}}{A_{\mathrm{c}}}\right)} - P_{\mathrm{laser}} = 0.193\,\mathrm{W}. \tag{4.11}$$

This simple scenario was recreated in the COMSOL environment whereby a single test mass was suspended in a 7 K cryostat chamber with an incident heat flux to the TM of 0.193 W. The COMSOL predicted TM temperature was output as 122.91 K, therefore matching the analytical solution to within 0.07% and showing good agreement between the calculated and modelled results.

4.3.4.2 Consideration of the Monolithic Suspension Stage

At 123 K the heat exchange will be both radiative and via conduction through the silicon fibres, therefore the entire monolithic stage of the suspension must be considered. The modelling of the suspension chain can be simplified by ignoring the metal UM and the reaction chain, which can be conductively cooled to any temperature via soft heat links, discussed in Section 4.3.5. To limit elements from touching the TM the heating power in this case is applied to the IM which heats the TM through a mixture of radiation and conduction as seen in Figure 4.9. This scenario is too complex for an analytical result which is why FEA analysis is so valuable. The FEA model also allows a more realistic model of heat flow from the inclusion of the thermal conductivity of hydroxide-catalysis bonds, and the effects of phonon-boundary scattering in the fibres.

The power conducted through the fibres, P_{cond} is dependent on the silicon thermal conductivity $\kappa_{\mathrm{Si}}\,(T)$, the fibre length, l, and total cross-sectional area A_{f}, and the temperature difference between the IM and TM as [217]

$$P_{\mathrm{cond}} = \frac{\kappa_{\mathrm{Si}}\,(T)\ \times\ (T_{\mathrm{IM}}\ -\ T_{\mathrm{TM}})}{l} \times\ A_{\mathrm{f}}. \tag{4.12}$$

Using the fibre diameter of 216 µm calculated in Section 4.2.2 the cross-sectional area of four fibres is $A_f = 1.47 \times 10^{-7}\,\text{m}^2$, with fibre length $l = 0.25\,\text{m}$. Thermal conductivity is temperature dependent therefore the FEA software will use the value of conductivity corresponding to the temperature of each finite element. The IM heating power P_{heater} was chosen to produce a final TM temperature of $T_{\text{TM}} = 123\,\text{K}$.

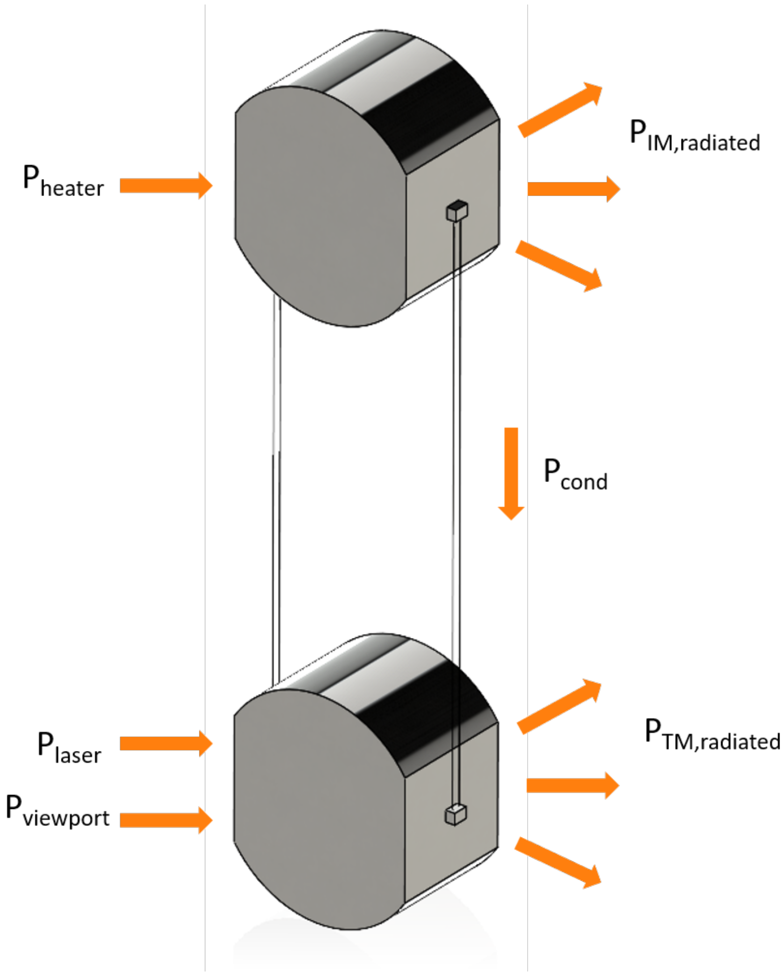

Figure 4.9: Computer-aided design of the monolithic suspension stage which was used in Section 4.3.4.2. The considered heat flows are represented by orange arrows.

To create the FEA model of the monolithic stage 55 mm × 60 mm bonding flats were added on each side of the IM and TM CAD models, bringing each mass to 983 g. To accommodate the addition of the fibres the IM and TM have a silicon ear hydroxide catalysis bonded on each flat, to which the silicon fibres will be attached. These ears were assumed to have a similar ear design as used in aLIGO but scaled down to give the maximum value aLIGO shear stress on the bond with the smaller GCIF masses. The hydroxide catalysis bonded area in the aLIGO suspension chain is $20 \times 60\,\text{mm}$ [201]. This gives a bonding area of $1200\,\text{mm}^2$ for a 40 kg mirror suspended with 2 ears, giving each ear-to-mass bond a shear stress of $1.64 \times 10^5\,\text{Pa}$. For a similar shear stress given by two ears suspending a 1 kg mirror the silicon ear dimensions were chosen to be $5 \times 7\,\text{mm}$ with a bond thickness of 650 nm

[104]. This stress value is well below the breaking stress of a silicon-silicon HC bond, with a safety factor of 183. The aLIGO ear ratio was slightly modified to give a larger surface area on the x-facing side of the ear which would be necessary for bonding if future work favours ribbons to be used in the monolithic stage. Ear design and bonding concepts for the triple suspension are still areas of active development. These ear dimensions were used in the thermal modelling for an estimate of heat flow through the monolithic stage rather than as a suggested concept for the mechanical design. The silicon fibres were attached to the ears by the flat cylinder faces of the fibres via HC bonding. Building of the real suspension will necessitate bonding or welding of the fibres to the ears which is an area of active research and beyond the scope of this investigation. In either a weld or a HC bond there will be a reduction in thermal conductivity compared to a body of crystalline silicon. Silicon welding is a developing technology and there is not yet measured data for thermal conductivity across a welded region. The thermal conductivity of a silicon-silicon HC bond has been measured between $10 - 300\,\mathrm{K}$ with these temperature-dependent values being used in the FEA modelling (example values at 18 K and 123 K have a thermal conductivity of 0.04 W/m/K and 0.14 W/m/K respectively) [104].

The FEA model from Section 4.3.4.1 was used, with the geometry of the monolithic suspension imported from the CAD. In addition to the Surface-to-Surface Radiation module used previously the Heat Transfer in Solids module was used to add a heating power, P_{heater}, to the IM and also for the laser and viewport heating on the TM, P_{laser} and P_{viewport}. The cryostat inner walls were kept at 7 K with the ambient heat flux, P_{ambient}, coming from these inner shield walls. The heating power was chosen to produce a TM temperature of $T_{\mathrm{TM}} = 123\,\mathrm{K}$.

4.3.4.3 Results of 123K Cooling

In the initial case, the monolithic silicon suspension stage was suspended in the cryostat with an incident heat flux in from the laser, ambient heat sources, viewport, and IM heater. The TM reached the desired temperature of $T_{\mathrm{TM}} = 123.0\,\mathrm{K}$ when $P_{\mathrm{heater}} = 3.8\,\mathrm{W}$ was incident on the IM. For a steady-state solution the temperature of the IM is $T_{\mathrm{IM}} = 260.6\,\mathrm{K}$. With the addition of HC bonds at the mass-ear and fibre-ear interfaces there is little change for the same heating power with $T_{\mathrm{TM}} = 123.0\,\mathrm{K}$ and $T_{\mathrm{IM}} = 260.6\,\mathrm{K}$. This is expected due to the small surface area and thickness of bonds coupled with the domination of radiative heat transfer at 123 K. In the case where the fibre thermal conductivity is lowered to account for phonon-boundary scattering in addition to the inclusion of HC bonds there is still only a small change in the mass temperatures. Here, $T_{\mathrm{TM}} = 123.0\,\mathrm{K}$ while $T_{\mathrm{IM}} = 260.7\,\mathrm{K}$

for the same input heater and laser powers. This study proves the dominance of radiative heat transfer at 123 K and shows that the IM must be 137.7 K hotter than the TM for the desired TM temperature in this scenario. This is not an ideal design as this creates a large thermal gradient across the silicon fibres, causing some regions of the suspension to sit far from the 123 K thermoelastic nulling point and increases temperatures in the monolithic suspension region, which increases thermal noise.

To improve the design a small heat shield was added around the monolithic stage with the aim of reducing the thermal gradient across the fibres. In this study the input heating was removed from the IM. Instead, the heat shield is to be heated via thermal links and will radiatively heat the entire monolithic stage. A 350 mm height inner diameter, 350 mm height cylindrical heat shield was modelled in CAD and imported into COMSOL. The cylinder thickness was modelled as 10 mm and should be kept as thin as possible to shorten cooling time of the payload. A cylinder of these dimensions sits close to the masses, increasing the efficiency of radiative heating, but also leaves space around the masses for the reaction chain or other equipment. The shield material was set as copper to match the cryostat, with an inner surface painted matte black for an emissivity of $\epsilon_{\text{shield}} = 1$. The heat shield was heated to 132.1 K which produced a TM temperature of 123.0 K and an IM temperature of 121.9 K, giving a 1.1 K thermal gradient across the fibres.

4.3.5 Cooling to 18K

As stated in Section 4.3, the relationship $q_{\text{radiation}} \propto T^4$ means radiative cooling is less effective at lower temperatures. At 18 K radiative cooling can be ignored and instead the TM must be cooled entirely through conductive heat extraction. This places a much greater importance on the thermal conductivity of the monolithic stage fibres, ears, bonds, and masses, and the conductive cooling of the entire suspension and reaction chains must now be considered.

4.3.5.1 KAGRA Cryogenic Design

KAGRA is the only GW detector currently operating at cryogenic temperatures, and as such the cooling method for the GCIF suspension used elements from KAGRA to inform the design. In this Japanese detector the mirror suspension chain is made up of the test mass, TM, intermediate mass, IM, and the upper mass known as the marionette, MN. The reaction chain follows suit with a mirror recoil mass, TMR, intermediate recoil mass, IMR, and marionette recoil mass, MNR. The suspension chain and reaction chain are both attached to a common suspended platform, PF. Figure 4.10 provides a diagram of this system. Two cooling bars are thermally connected to the 2nd stage of the 4 K cryocoolers, and are connected to the MNR using soft thermal links via a triple-stage vibration isolation system [216]. From the MNR, one set of thermal links cools the reaction chain from MNR to IMR to TMR. Another set of heat links from the MNR first links to the PF before moving down the suspension chain from PF to MN to IM. There can be no thermal links linking the IM and TM to avoid direct coupling of seismic noise to the TM, so the TM itself must be entirely cooled to 20 K via conduction through the sapphire suspension fibres. Soft thermal links were specifically designed for the KAGRA suspension chain. These links provide high thermal conductivity whilst being extremely flexible for reduction of vibration to the payload. Copper and aluminium are often used for thermal links due to their high thermal conductivities at cryogenic temperatures. Ultra-high purity aluminium (99.9999%) was chosen for use in KAGRA due to the lower spring constant of aluminium compared to copper. The lower spring constant creates a softer link therefore coupling less seismic noise between each stage. Many thin wires were twisted together to create each thermal link: the spring constant for multiple twisted wires is lower than for one thicker wire, even for the same total cross-sectional area [215, 216].

4.3.5.2 GCIF Thermal Design

At the current stage of the design process the GCIF suspension design differs from the KAGRA set-up. The GCIF suspension chain has an upper mass UM, intermediate mass IM, and test mass TM, while the reaction chain will have a reaction upper mass RUM, reaction intermediate mass RIM, and reaction test mass RTM. There is no upper platform common to both chains, although a common upper platform design may be investigated in future. In the thermal model of the GCIF payload it was decided to have a 7 K cold sink attached to the RUM. Thermal links would be attached between the RUM and RIM, and RIM to RTM. Unlike the KAGRA design a long thermal link would attach the bottom of the reaction chain, RTM, to the top of the suspension chain at the UM. A final thermal

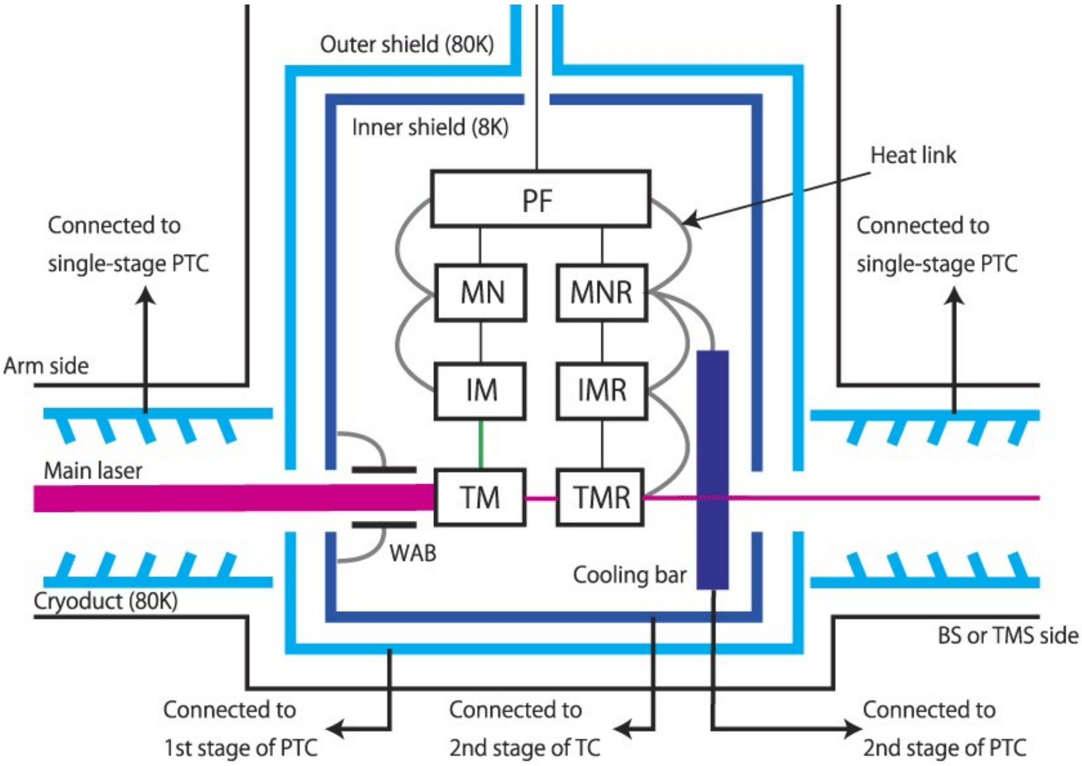

Figure 4.10: Side-view of the KAGRA cryogenic cooling system showing the heat link connections to the reaction chain, up to the suspended common platform, and down the test chain as far as the intermediate mass [197].

link would be attached between the UM and IM, and again the TM must be cooled via heat extraction through the crystalline suspension fibres. This cooling scheme was chosen to give the maximum possibility for seismic isolation of the TM from the cryocooler noise, and is shown in Figure 4.11.

The modelling of the heat extraction necessitated the addition of the reaction chain to the suspension model. At this stage the designs of the suspension UM and the reaction RUM have not yet been finalised, however the shape of each mass does not have a large effect on heat flow as only the steady-state solution is being considered. Mass size and shape will have more effect when calculating payload cool-down time, not considered here. The reaction chain for this study was duplicated from the suspension chain for simplicity. The suspension and reaction chain CAD models can be seen in Figure 4.9 (a). In the steady-state solution of the model each mass is considered to be held at a constant temperature. The cryostat inner shield is 7 K, and the TM must be held at 18 K. The heat flow through the thermal links is calculated, with heat link parameters modified to allow cooling of the TM to the desired temperature.

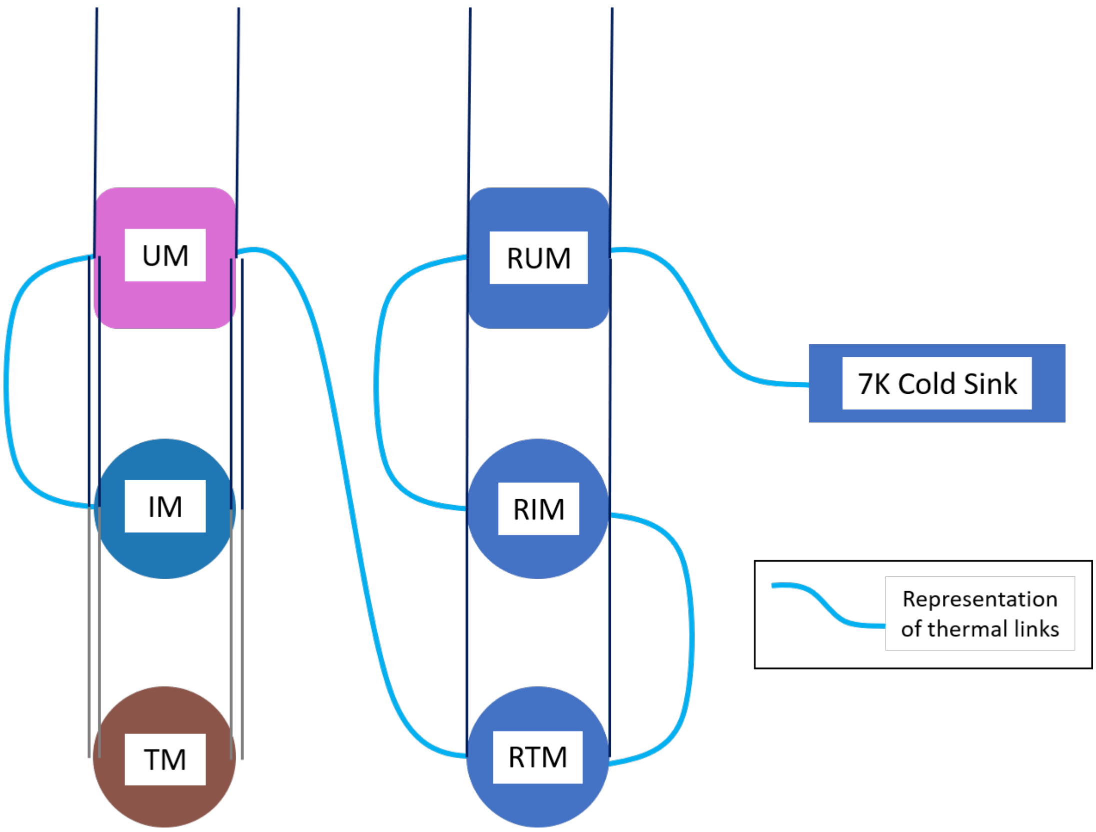

Figure 4.11: Diagram of the modelled GCIF cryogenic cooling system showing the test chain (left) and reaction chain (right) with the heat links shown in light blue.

The heat flux, $q_{\mathrm{m,n}}$, through a heat link attached between two stages, m and n (where temperature $T_{\mathrm{n}} > T_{\mathrm{m}}$), can be calculated from Equation 4.12 as

$$q_{\mathrm{m,n}} = \frac{\kappa_{\mathrm{hl}}\,(T_{\mathrm{n}} - T_{\mathrm{m}})}{l_{\mathrm{m,n}}} = \frac{P_{\mathrm{hl}}}{A_{\mathrm{hl}}}. \tag{4.13}$$

Assuming the same design of heat link is used at each stage, the heat link thermal conductivity, κ_{hl}, and heat link cross-sectional area, k_{hl}, will be constant through all the stages. Different lengths of heat link, $l_{\mathrm{m,n}}$, will be required based on the distance between stages m and n. P_{hl} is the power flowing across the heat links. In a steady-state solution where the TM is held at 18 K, this $P_{\mathrm{hl}} = 0.002\,\mathrm{W}$ at every stage as it must remove the heat flux incident on the TM from the laser heating.

Heat links are necessary for cooling as the suspension wires alone will not cool the suspensions sufficiently. The ultra-high purity heat links in use at KAGRA have a thermal conductivity of $\kappa_{\mathrm{hl,Al}} = 18500\,\mathrm{Wm^{-1}K^{-1}}$ at 10 K and a cross-sectional area of $A_{\mathrm{hl,Al}} = 6.06 \times 10^{-6}\,\mathrm{m^2}$ [216]. Despite the advantages of their high thermal conductivity and low spring constant these heat links are not available for commercial purchase. A more affordable and commercially available option is oxygen-free high thermal conductivity (OFHC) copper heat links such as [186]. These OFHC copper heat links have a thermal conductivity of $\kappa_{\mathrm{hl,Cu}} = 1500\,\mathrm{Wm^{-1}K^{-1}}$ at 10 K and a cross-sectional area of $A_{\mathrm{hl,Cu}} = 1.96 \times 10^{-5}\,\mathrm{m}^2$. The suspension wires used are made of stainless steel, with a 100 µm diameter. This gives the wires a thermal conductivity of $\kappa_{\mathrm{w}} = 25\,\mathrm{Wm^{-1}K^{-1}}$ at 10 K and a cross-sectional area of $A_{\mathrm{w}} = 7.85 \times 10^{-9}\,\mathrm{m}^2$ [74, 88]. Equation 4.13 can be rearranged to give the power flowing across the heat link or wire given the thermal conductivity, temperature gradient, length, and cross-sectional area of the heat link or wire. Assuming the maximum possible temperature gradient being 18 K to 7 K from the TM to cryostat, and a length of 0.25 m, the values above can be substituted into Equation 4.13 to give a maximum conducted power of 4.93 W for aluminium heat links, 1.29 W for copper heat links, and 8.64×10^{-6} W for the stainless steel suspension wires. This shows that the suspension wires alone are unable to conductively cool the TM, but either the aluminium or copper heat links would be sufficient to transport the 0.002 W of laser heating away from the TM. In the rest of this study the commercially available copper heat links were used. The heat link modelling in this section aims to provide a solution for cooling the test mass to the desired temperatures. Part of the implementation of this design would be future modelling to estimate (and minimise) any seismic coupling via the thermal links to the suspension chain.

4.3.5.3 Finite Element Analysis Cooling Model

The CAD model of the suspension and reaction chains with copper heat links was imported into COMSOL. The cryostat shield and monolithic stage materials were chosen as before, with phonon-boundary scattering and the effects of HC bonding included. Aluminium was imported from the COMSOL material library for the UM and RUM as well as stainless steel for the upper stage wires. The temperature-dependent copper used for the cryostat shield was also used to model the heat links. The inner surface of the cryostat shield was set to 7 K, with the $P_{\mathrm{ambient}} = 100\,\mu\mathrm{W}$ heat load coming from this surface. The TM face turned towards the cryostat viewport had an added heat load for laser beam heating and radiative viewport heating, $P_{\mathrm{laser}} + P_{\mathrm{viewport}} = 0.016\,\mathrm{W}$. A multiphysics study was conducted using the Heat Transfer in Solids and Surface-to-Surface Radiation modules which allows exploration of both the conducted and radiated heat throughout the suspensions.

4.3.5.4 Results of 18K Cooling

Strength limited fibre diameter.
The silicon fibres used up to this point had a diameter limited by fibre breaking stress (strength limited, SL) of $d_{\text{silicon,SL}} = 216\,\mu\text{m}$, as calculated in Section 4.2.2. These ultra-thin fibres severely limit the heat conducted away from the TM due to the small cross-sectional area and phonon-boundary scattering effects greatly reducing silicon thermal conductivity at these dimensions. The reaction chain masses cool to 7.76 K, 8.47 K, and 9.03 K for the RUM, RIM, and RTM respectively. The test chain UM and IM are connected via the thermal links and cool to 9.77 K and 9.92 K. The TM reaches a final steady-state temperature of 80.27 K where radiative cooling begins to dominate the heat transfer process. This demonstrates that the 216 µm diameter fibres are insufficient for cooling to 18 K, and that the fibres used in the GCIF will be limited by their thermal conductivity rather than their strength. A thicker silicon fibre was then modelled.

Thermal extraction limited fibre diameter.
The diameter of a thermal extraction limited silicon fibre, $d_{\text{silicon,HE}}$, was calculated which would allow cooling of the TM to 18 K through the conduction of the heat load on the TM. From Section 4.3.2 the combined heat load on the TM is $P_{\text{load}} = P_{\text{laser}} + P_{\text{viewport}} + P_{\text{ambient}} = 0.016\,\text{W}$. Equation 4.12 for thermal conductivity can be rearranged to output the fibre radius necessary to transport a given heat load:

$$d_{\text{HE}} = 2 \times \sqrt{\frac{P_{\text{load}}\, l}{n\, \pi\, \kappa\,(T)}} = 650\,\mu\text{m} \tag{4.14}$$

where $n = 4$ is the number of fibres. $\kappa\,(T)$ is integrated across the temperature range used and iteratively multiplied by the diameter-dependent photon-boundary scattering reduction factor. The new fibre diameters were modelled in CAD and imported to COMSOL, with the other model inputs unchanged.

With the four thicker silicon fibres between the IM and TM there is a smaller thermal gradient between the TM and the 7 K cryostat walls. The reaction chain masses cool to 7.01 K, 7.25 K, and 7.40 K for the RUM, RIM, and RTM. The suspension chain masses connected via thermal links cool to 7.74 K for the UM and 7.88 K for the IM. The TM, cooled only by conduction of the silicon fibres, reaches a steady state temperature of 17.65 K in the COMSOL model. COMSOL model estimations for payload cool-down time lie at 13 days, however this value is dependent upon the final thermal mass in the system and further modelling of fibre-mass attachments. This modelling demonstrates the feasibility of cooling the TM to 18 K in the GCIF cryostat, including factors such as realistic laser, ambient, and viewport heat loads on the mass, and the phonon-boundary scattering and hydroxide catalysis bonding effects on the monolithic stage thermal conductivity. A crucial step for the future work on this project would be modelling the seismic coupling to the reaction and test chain from the addition of the copper heat links. If the seismic coupling through the heat links is too great there will need to be amendments to the TM cooling design. This could include considering the use of soft aluminium heat links, addition of a heat link vibration-isolation system such as used at KAGRA, or design of an alternative cooling system.

4.4 Suspension Thermal Noise Modelling

Thermal displacement noise was calculated for the monolithic stage of the cryogenic triple suspension using the method outlined in Chapter 2. The contribution to thermal noise was considered for a 1 kg silicon test mass attached to four crystalline fibres of a circular cross-section and 0.25 m length. Both sapphire and silicon were considered as potential fibre material choices in this thermal noise study, consistent with the current research interests of the GW community. A study of both materials is important due to the material and geometry dependence of suspension thermal noise, and a comparison of the materials is shown at the end of this section. For both silicon and sapphire fibres the thermal displacement noise was calculated for operation at 123 K and at 18 K. At 123 K cooling is radiatively dominated so fibre diameters are limited by their strength, calculated using Eq. 4.2. For silicon this was previously shown to be $d_{\mathrm{silicon,SL}} = 216\,\mu\mathrm{m}$. Sapphire breaking stress was taken as $\sigma_0 = 800\,\mathrm{MPa}$, so with a safety factor of 3 the strength-limited diameter is $d_{\mathrm{sapphire,SL}} = 108\,\mu\mathrm{m}$ [34]. At 18 K the fibre diameters are limited by their thermal extraction capabilities as they must maintain the test mass temperature through conductive cooling. With the heat loads taken from Section 4.3.2, Equation 4.14 shows this to be $d_{\mathrm{silicon,HE}} = 650\,\mu\mathrm{m}$ for silicon. Sapphire thermal conductivity values between 2.3 K - 300 K were taken from heat conductivity measurements on sapphire fibres considered for use in KAGRA [117], and give a heat extraction limited fibre diameter of $d_{\mathrm{sapphire,HE}} = 194\,\mu\mathrm{m}$.

Material	Limitation	Diameter [µm]
silicon	strength	216
silicon	heat extraction	650
sapphire	strength	108
sapphire	heat extraction	194

Table 4.3: Silicon and sapphire fibre diameters in both strength and heat extraction limited cases.

4.4.1 Suspension Resonant Modes

The longitudinal, vertical, and violin modes are presented in Table 4.4 for the monolithic stage of the cryogenic suspension. Modes were calculated for the 123 K and 18 K cases for both silicon and sapphire fibres using the equations introduced in Section 2.2. The violin modes up to 1 kHz were calculated, with only the fundamental mode shown in Table 4.4 for brevity.

	Silicon Fibre Modes [Hz]		Sapphire Fibre Modes [Hz]	
	18 K	**123 K**	**18 K**	**123 K**
longitidinal	1.00	1.00	1.00	1.00
vertical	211.96	70.55	97.88	54.57
violin	113.37	340.59	288.82	518.02

Table 4.4: Monolithic stage longitudinal, vertical, and fundamental violin modes for the cryogenic triple suspension. Results are presented for the case of either silicon or sapphire fibres between the IM and TM, with modes calculated for both the 18 K and 123 K temperature regimes. For simplicity, only the fundamental violin mode is presented.

The longitudinal mode frequency depends only on gravitational acceleration and fibre length. These variables are independent of fibre material and temperature so the longitudinal mode is the same for each case considered.

The suspension vertical mode is proportional to fibre radius and to the square root of the material Young's modulus. Young's modulus is assumed to be temperature independent in this region [133], but is material dependent, with the Young's modulus of sapphire over 2× greater than that of silicon [148]. The stiffer sapphire would produce a higher vertical frequency than a silicon fibre of the same dimensions. However, the lower breaking stress of silicon requires that the strength-limited silicon fibre is twice the diameter of the strength-limited sapphire fibre, pushing the vertical resonance frequency of the silicon fibre above the resonance of the sapphire fibre. Similarly at 18 K the lower thermal conductivity of silicon requires that the silicon fibre diameter be 3× greater than the diameter of the sapphire fibre for the same heat load, and as such the vertical resonance for the silicon fibre is higher in frequency than for the sapphire fibre.

The violin mode resonance frequencies are inversely proportional to fibre radius and to the square of the material density, with sapphire density from [148]. The difference in silicon and sapphire fibre radii dominates over the material density effects, pushing the sapphire fibre violin modes higher in frequency than those of the silicon fibre in each temperature case.

4.4.2 Mechanical Loss Contributions

The four mechanisms contributing to suspension mechanical loss are material bulk loss, surface loss, thermoelastic loss, and loss of the weld or bonded region around the suspension fibre. These mechanisms were detailed in Section 2.3.1 with the silicon loss values presented in Section 2.4.1.2. Sapphire material properties were required to calculate suspension thermal noise for the case of a silicon test mass suspended by four sapphire fibres.

Sapphire bulk loss was taken as $\phi_{\mathrm{bulk}} = 5.6 \times 10^{-9}$ and assumed to be constant with temperature [148, 161, 196, 199]. Sapphire surface loss values were estimated from mechanical loss measurements of a 1.6 mm diameter sapphire fibre from $5 - 300\,\mathrm{K}$, made as a test fibre for the KAGRA mirror suspension [181]. It was assumed that surface loss would be the dominant loss mechanism for a fibre of that diameter to provide a conservative surface loss estimate. The losses were scaled to match the fibre diameters calculated for the GCIF suspension using Equations 2.14 and 2.15, giving $\phi_{\mathrm{surface,transverse}} = 8.25 \times 10^{-6}$ at 18 K and $\phi_{\mathrm{surface,transverse}} = 1.48 \times 10^{-4}$ at 123 K. Thermoelastic loss was calculated with Equation 2.16 using sapphire c_{p} and α values from [209]. Sapphire welding is an area of active research with sapphire fibre welding having been demonstrated as a proof of concept at the University of Glasgow [57, 172]. It is therefore assumed that the ear-fibre joint will be welded rather than HC bonded. The mechanical loss of sapphire-sapphire welds have not yet been characterised so a best-case and worst-case weld loss scenario have been assumed, with the real weld loss likely to lie between these values. The best-case scenario assumed the weld loss to be 10$\times$ that of the base material as approximately observed in current aLIGO fused-silica welds [61]. This is likely an optimistic assumption as the crystalline sapphire is predicted to produce a higher-loss weld than amorphous fused-silica through the difficulty of perfectly orientating the crystal lattice structure during the weld [57]. The worst-case scenario assumes the sapphire welded region to be as lossy as a sapphire-sapphire HC bond, taken from [101] as $\phi_{\mathrm{joint}} = 3 \times 10^{-4}$ at 18 K and $\phi_{\mathrm{joint}} = 8 \times 10^{-2}$ at 123 K. These loss contributions are then weighted by the energy distribution across the fibre and joint region, and are plotted in Figure 4.12 for a comparison of the silicon and sapphire loss values at each temperature.

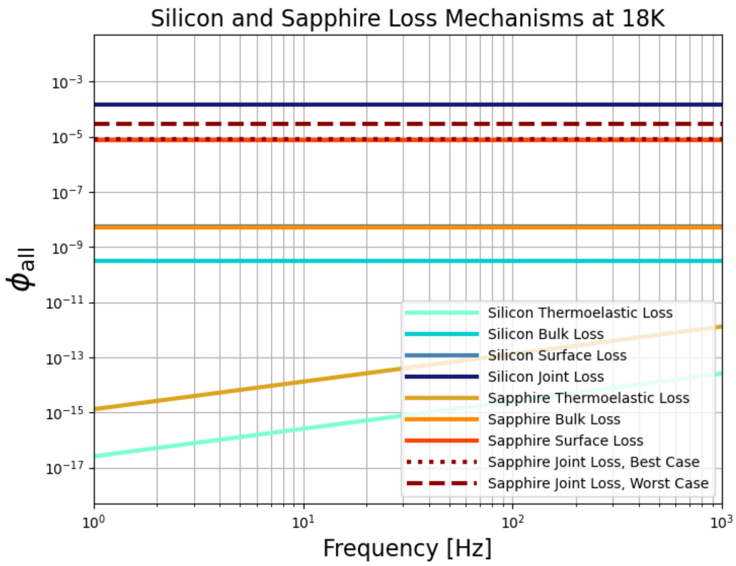

(a) Loss at 18 K

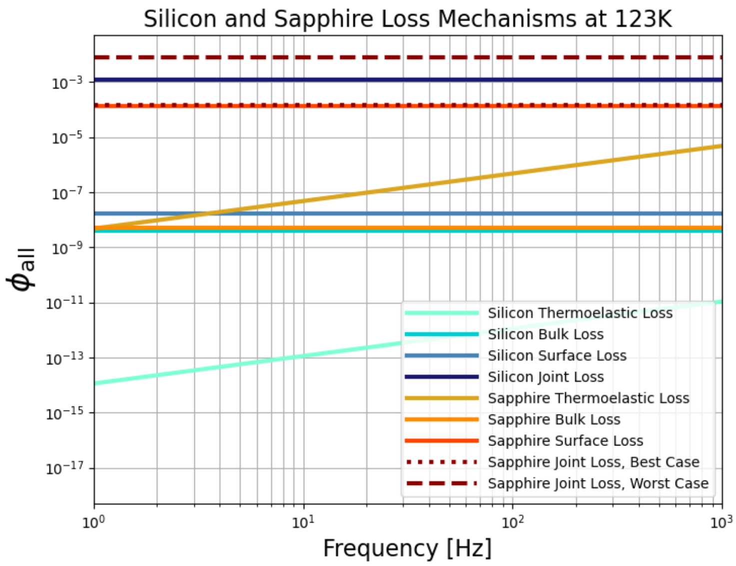

(b) Loss at 123 K

Figure 4.12: The thermoelastic, bulk, surface, and jointing contributions to the mechanical loss of a 1 kg suspension with either four silicon fibres or four sapphire fibres. Presented at both 18 K and 123 K, and weighted by the energy distribution across the fibre and joint.

The pendulum mode dilution factor is derived in Section 2.3.2, giving Equation 2.21. For both silicon and sapphire fibre scenarios the TM mass, m, constant of gravitational acceleration, g, fibre length, l, number of fibres, n, bending factor, c_{bend}, and fibre tension, F, are constant. Young's modulus, Y, is material dependent and the second area moment of the fibre, I, depends on fibre radius. Equation 2.21 was used to calculate the pendulum mode dilution factor for the case of a suspension with silicon fibres, and a suspension with sapphire fibres. At the time of writing the cross-sectional shaping of silicon and sapphire fibres along their length has not yet been demonstrated to be feasible and therefore straight fibre profiles were assumed. Both fibre scenarios were calculated for operation at 18 K and 123 K, with the results presented in Table 4.5. As discussed in Section 2.3.2 there is no dilution for the vertical suspension mode, therefore $D_{\mathrm{vert}} = 1$, and the violin mode dilution can be calculated from pendulum mode dilution as $D_{\mathrm{violin}} = \frac{1}{2} \times D_{\mathrm{pend}}$.

Fibre Material	D_{pend} **at 18 K** [Hz]	D_{pend} **at 123 K** [Hz]
silicon	10.2	92.4
sapphire	74.3	238.9

Table 4.5: Pendulum mode dilution values calculated for silicon and sapphire fibre scenarios at 18 K and 123 K.

With the resonant modes, mechanical loss mechanisms, and dilution factors known, these can be combined using Equations 2.11, 2.12, and 2.13. This gives the total mechanical loss of the pendulum, vertical, and violin modes, each weighted by the energy distribution across the suspension fibres.

4.4.3 Displacement Noise

The longitudinal mirror motion caused by suspension thermal noise can be calculated using Equation 2.22, substituting ϕ_{pend} and ω_{pend} for ϕ_{mode} and ω_{mode}. The resulting $x_{\mathrm{rms,pend}}(\omega)$ is shown in blue in Figure 4.13 for both silicon and sapphire.

As described in Section 2.3.3 the vertical mode's contribution to longitudinal mirror motion can be obtained by substituting ϕ_{vert} and ω_{vert} for ϕ_{mode} and ω_{mode} then multiplying by the vertical-to-horizontal cross-coupling factor. For the GCIF this factor is assumed to be 0.1%. This motion, $x_{\mathrm{rms,vert}}(\omega)$, is shown in orange in Figure 4.13.

The violin mode contribution to test mass longitudinal motion is calculated via Equation 2.25, which must be evaluated separately for each violin mode visible in the experiment's sensitivity band. For the GCIF suspension at 18 K this is violin modes 1-8 for a silicon fibre suspension, and modes 1-3 for a sapphire fibre suspension, with a smaller number of violin modes visible for both materials at 123 K. Violin modes are plotted in green in Figure 4.13.

Figure 4.13 shows example plots for the thermal displacement noise contribution from each suspension resonant mode. These are shown for the cases of a silicon fibre suspension at 18 K and a sapphire fibre suspension with weld-like joint loss at 18 K. Appendix A displays in full the thermal displacement noise plots - including resonant mode contribution - for each of the six suspension cases considered.

The total test mass thermal displacement noise from all resonant modes is found by summing $x_{\mathrm{rms,pend}}(\omega)$, $x_{\mathrm{rms,vert}}(\omega)$, and $x_{\mathrm{rms,violin,p}}(\omega)$ in quadrature (Equation 2.26). This $x_{\mathrm{rms,total}}(\omega)$ is evaluated separately for each of the six cases considered: a silicon fibre suspension at 18 K (pink solid line), a silicon fibre suspension at 123 K (pink dashed line), a sapphire fibre suspension with a best-case and worst-case weld loss at 18 K (blue solid line and orange solid line), and a sapphire fibre suspension with a best-case and worst-case weld loss at 123 K (blue dashed line and orange dashed line). Results are in Figure 4.14 with the above colour notation.

4.4.4 Results of Suspension Thermal Noise Modelling

4.4.4.1 Sapphire Thermal Displacement Noise

The effect of the weld-like and HCB-like joint losses were compared for sapphire at 18 K and 123 K. At 18 K the difference between the joint cases is small with a HCB-like join having a loss 3.6 $\times$ higher than a weld-like join. At 123 K there is a more significant 53 $\times$ higher loss for a HCB-like join compared to a weld-like join. With joint loss being the dominant loss mechanism for sapphire in both temperature regimes, the difference in the join cases is reflected in the total thermal displacement comparison. This is shown in Figure 4.14 where there is only a small difference in RMS thermal displacement noise between the sapphire jointing cases at 18 K and a much larger difference between the cases at 123 K. There is no joint loss contribution to the vertical mode displacement noise. This is observed as the vertical mode displacement noise is the same for sapphire between the two 18 K jointing cases, and between the two 123 K jointing cases. This results in a smaller overall vertical mode contribution to displacement noise for a HCB-like join compared to a weld-like join as the loss of the pendulum and violin modes increase with increasing joint loss. For an experiment operational at 18 K either jointing mechanism could be used, however at 123 K a weld-like loss is recommended if future research allows. Operation at 18 K results in a lower thermal displacement noise than operation at 123 K for any jointing case. In both temperature cases the weld-like join has a lower thermal displacement noise than a HCB-like join which makes the domination of jointing loss clear when compared to the surface, thermoelastic, and bulk losses.

4.4.4.2 Silicon Thermal Displacement Noise

Despite the large variation in silicon fibre diameters, both the pendulum mode and vertical mode have a lower displacement noise at 18 K than at 123 K. The contribution of the violin modes has a greater effect on the RMS displacement thermal noise. The thicker fibres needed for 18 K operation push the silicon fibre's violin modes both closer together and down in frequency space when compared to the thin fibres for 123 K operation. This results in more violin modes visible in the experimental sensitivity range at 18 K than for 123 K. In the < 100 Hz region where the pendulum mode dominates, the 18 K suspension has a lower thermal displacement noise than 123 K. However, the first of the violin modes appear for the 18 K fibres at 113.37 Hz and not until 340.59 Hz for the 123 K fibres. This causes a higher overall thermal displacement noise for the 18 K fibres between $100 - 300$ Hz. In the frequency space > 300 Hz where there are violin modes for both temperature cases, the thermal displacement noise levels are comparable. In the sensitivity range of the planned GCIF experiment there is not a significant difference between the thermal displacement noise levels of the 18 K and 123 K suspensions. In fact, having more violin modes present in the experimental sensitivity range may be an advantage for characterisation measurements, allowing the resonance frequencies and Q-factors of the physical suspension to be measured and compared to the model predictions.

4.4.4.3 Silicon and Sapphire Thermal Displacement Noise

Figure 4.13 shows a higher contribution of the vertical mode to sapphire's RMS thermal noise than seen for silicon. This is observed at both temperature regimes and for either sapphire jointing case. The vertical mode has only contributions from bulk loss and surface loss, with sapphire's higher surface loss being the cause of the difference in this case. As

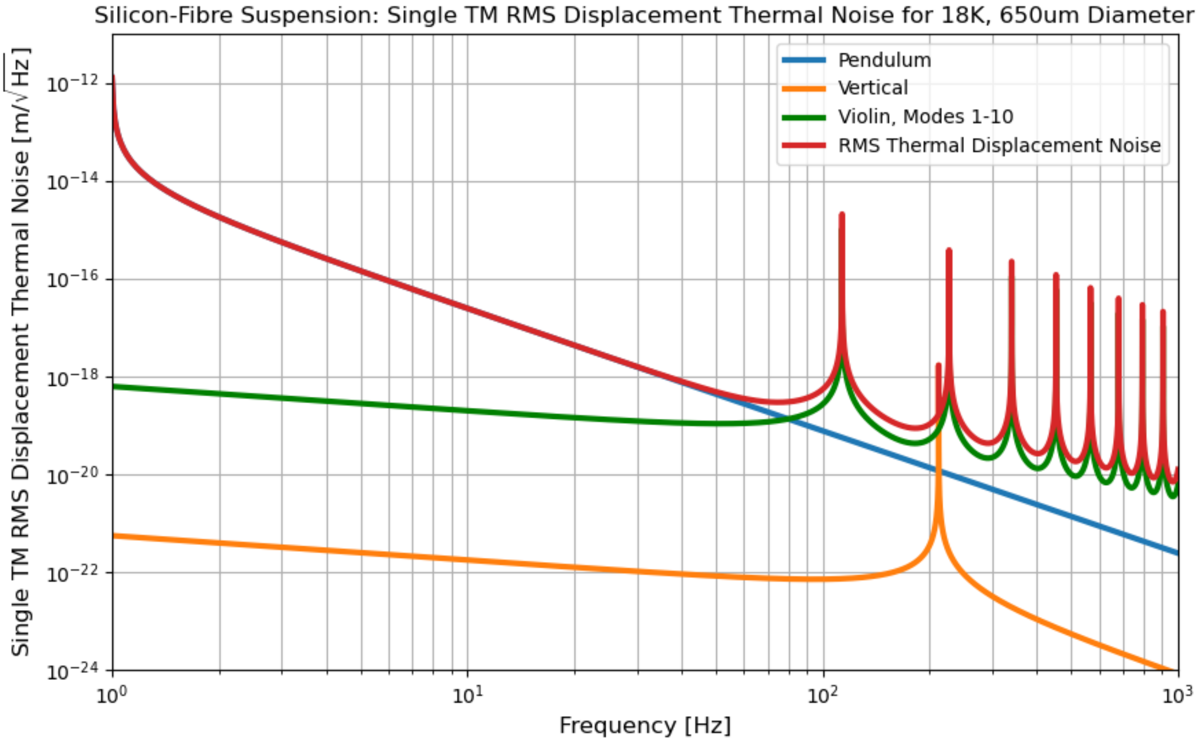

(a) Silicon 18 K

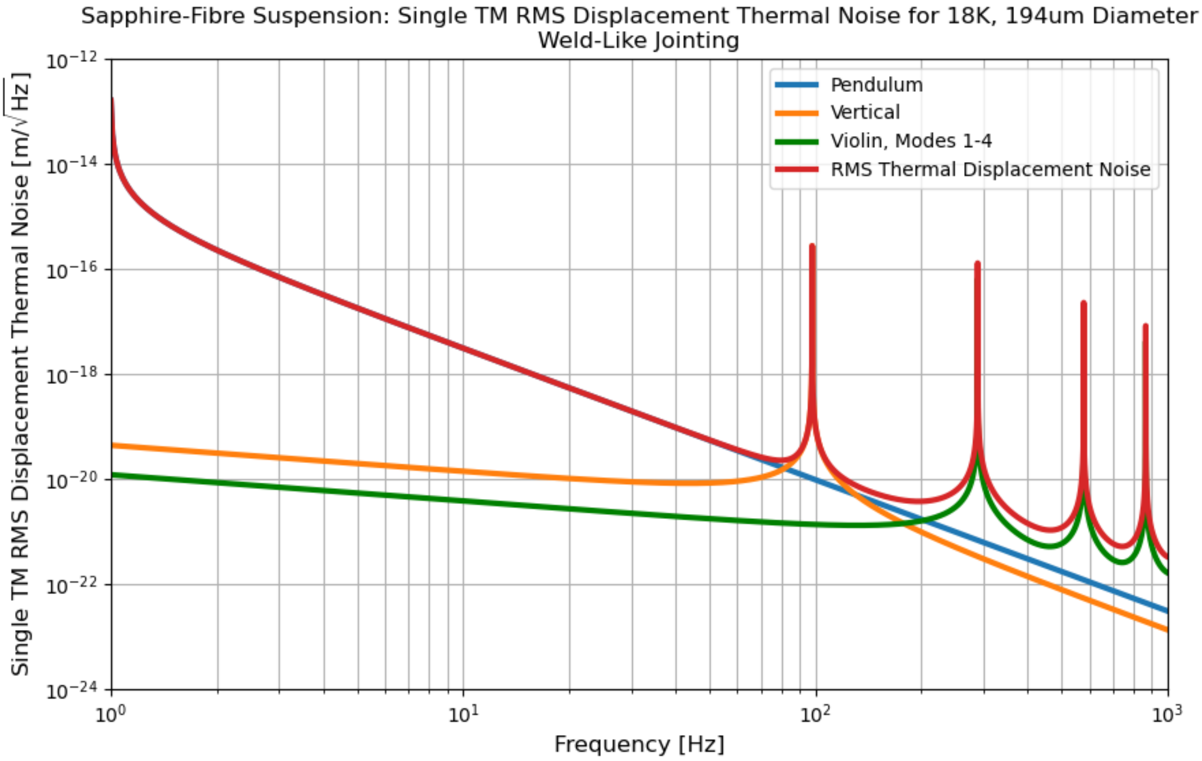

(b) Sapphire 18 K, weld-like joint.

Figure 4.13: The thermal displacement noise contribution from each suspension resonant mode with (a) a silicon fibre suspension at 18 K and (b) a sapphire fibre suspension with weld-like joint loss at 18 K. The pendulum mode contribution is shown in blue, the vertical mode in orange, and the violin modes in green. The total displacement noise from all resonant modes is plotted in red.

seen in Figure 4.12 sapphire's surface loss is 1300 × that of silicon's at 18 K, and is 8000 × larger at 123 K. The difference in material bulk loss is much smaller at only 1.3 × larger for sapphire at 18 K and 16.6 × at 123 K. Sapphire's higher surface loss is caused by both the differences in the conservative value taken for the material $h\,\phi_{\mathrm{s}}$ (which is lower for silicon) and by sapphire's lower fibre radius which increases surface loss as shown in Equation 2.14.

At 18 K silicon has a higher RMS displacement thermal noise across all frequencies when compared to sapphire, other than on the sapphire resonant peaks. The contribution of the pendulum mode dominates at lower frequencies and is higher for silicon due to the higher silicon joint loss. Although the vertical mode contribution is lower for silicon (as discussed), the lower frequencies of the silicon violin modes negate any potential gains from the vertical mode. The number of violin modes in the frequency range observed is higher for silicon than for sapphire, increasing the RMS noise at higher frequencies. Additionally, the noise level of each violin mode is higher for silicon, also caused by the higher silicon joint loss.

At 123 K Figure 4.14 shows the sapphire HCB-like join to have the highest thermal noise of the three cases. Silicon's thermal noise lies slightly below this, and the sapphire weld-like jointing case has the lowest thermal noise. This is due to the domination of joint loss in all these cases, with this pattern directly correlating with the joint loss pattern observed in the mechanical loss plots in Figure 4.12.

The building of a suspension without joint loss is not physically possible, however it is discussed here to examine what thermal noise limits might look like in a future case if not entirely dominated by joint loss. Figure 6 in Appendix A displays the no-jointing thermal noise curves for silicon and sapphire in comparison with the six cases previously discussed. In the case of no jointing contribution to mechanical loss, silicon is shown in brown at 18 K (solid line) and 123 K (dashed line), with sapphire shown in navy at 18 K (solid line) and at 123 K (dashed line). The other cases follow the convention of Figure 4.14.

Silicon with no jointing contribution is observed to have the lowest displacement thermal noise out of all cases considered. The 18 K silicon suspension has the lowest noise from 1 – 100 Hz, with the 123 K silicon suspension the lowest noise >100 Hz where the violin modes dominate. At 10 Hz (where there are no violin modes) the loss of silicon at 123 K is only 1.6 × higher than silicon at 18 K. Sapphire at 18 K with no joint loss is the next lowest case, with 14 × higher loss at 10 Hz than the no-joint 18 K silicon case. The weld-like joint and no-joint thermal noise curves are extremely close for sapphire fibres, showing

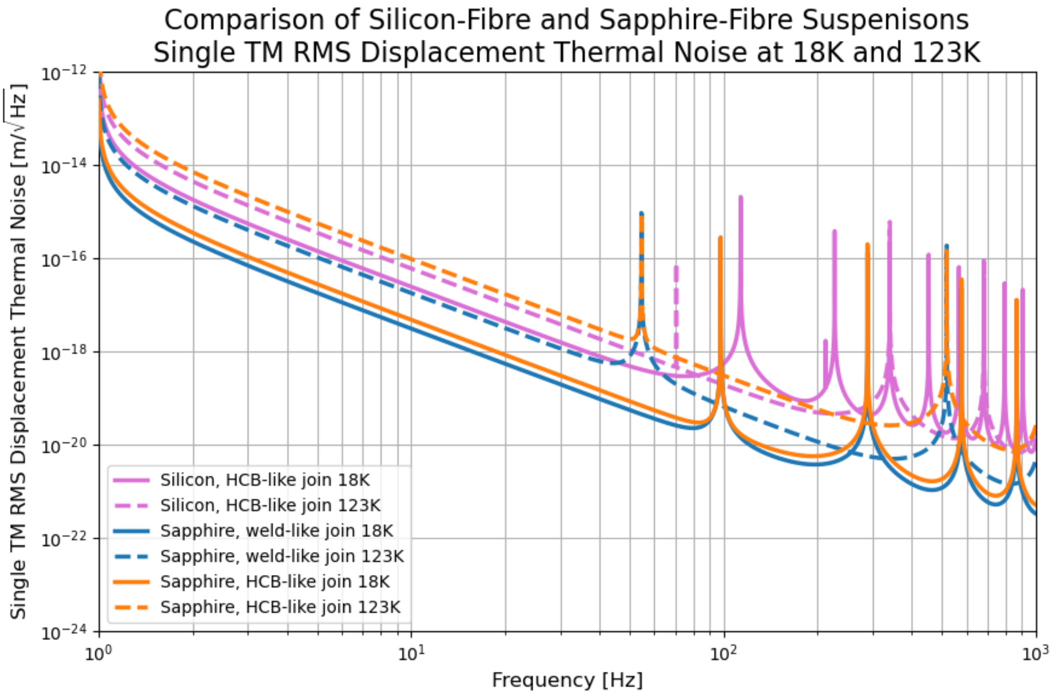

Figure 4.14: Silicon and sapphire RMS thermal displacement noise at 18 K and 123 K. Silicon thermal noise is shown in pink, sapphire with a weld-like join in blue, and sapphire with a HCB-like join in orange. 18 K thermal noise results are represented by a solid line and 123 K thermal noise results by a dashed line.

the importance of research into low-loss welding techniques. The 123 K case shows the largest difference, with the no-joint sapphire case 50 × larger than the no-joint silicon case. Again, the no-joint sapphire has a similar thermal noise to the weld-like jointing case. The 123 K loss for silicon fibres has a 230 × difference between the no-joint loss and joint loss cases which again proves the dominance of the jointing loss at this temperature.

Overall this study displays the complexity of thermal noise calculations for a monolithic suspension, and shows the number of factors which must be carefully considered in future design stages. There is a need for further mechanical loss measurements for both silicon and sapphire to greater increase the accuracy of thermal noise models such as this. This study also highlights the importance of joint loss considerations, and the finite element analysis of energy distribution across a fibre which affects the mechanical loss contributions. For future design stages of the GCIF cryogenic triple suspension it would be recommended to develop further research into welding and bonding techniques for silicon and sapphire fibres including the shaping of silicon and sapphire fibres, and to perform a full FEA analysis for the energy distribution across the fibre and jointing regions.

4.5 Conclusion

The investigations presented in this chapter produced a design concept for a triple-stage optical suspension suitable for use in a cryogenic cavity for direct measurements of coating thermal noise. This suspension meets the required seismic isolation and suspension thermal noise requirements. A cooling system was designed to give effective operation at both the 123 K and 18 K phases of the Glasgow Cryogenic Interferometer Facility and demonstrates that the silicon test mass can be cooled to the desired temperatures in each case.

The initial design of the cryogenic cavity suspension considered experimental sensitivity requirements, material availability and cost, ease of manufacture, and cryostat space constraints. This study yielded a triple-stage suspension with 1 kg silicon cavity optics, and wire/fibre lengths of around 0.3 m. A Mathematica model was built to model the pendulum dynamics of this triple-stage system. The transfer functions output from this model were combined with cryostat seismic measurements to predict the coupling of ground seismic noise through to the suspended optic for various scenarios. A comparison of circular cross-sectional fibres and rectangular cross-sectional ribbons showed almost no difference in seismic coupling levels.

Four different suspension models were considered to investigate the effect of blade springs on seismic coupling levels. Having only one stage of blade springs at the suspension mounting point does not provide sufficient seismic isolation to the optic. Two stages of 120 mm length blades just meets the noise requirements at 100 Hz but there are greater isolation gains made when increasing the length of the top blade springs to 500 mm. The longer blades push the blade resonant frequency lower therefore increasing seismic isolation at 100 Hz compared to the 120 mm blades. The blade springs installed on the UM remain at 120 mm to stay consistent with the expected UM dimensions. Three stages of blade springs provide the greatest seismic isolation to the optic, however maintaining a monolithic crystalline design of the UM and the optic would necessitate that the lower blade springs be manufactured from silicon. Silicon blade springs of a suitable strength have not yet been demonstrated and is a recommended area of future research. A double-stage blade spring design with long upper blades is recommended for the initial phase of the GCIF project.

To add additional seismic isolation without the need for silicon blade springs the design of a seismic isolation platform was considered. Such a platform must fit within the available cryostat space, allowing room for two test chain suspensions, two reaction chains, and auxiliary optics. An initial concept design was created of a compact nested pendulum design with an inverted pendulum. A Mathematica toy-model demonstrated the feasibility of this design which proved a stable set-up is possible with the correct placement of the upper mass against the suspended ring platform. This concept can be updated as the design for the suspension and reaction chains evolve.

A suspension cooling system was designed for the two planned stages of GCIF operation of 123 K and 18 K. The heat loads within the cryostat were considered alongside the dominant heat transfer mechanisms (radiative at 123 K and conductive at 18 K) to produce a cooling scheme for each scenario. At 123 K radiative cooling dominates and the test mass cools below 123 K sitting within the cryostat's 7 K inner chamber. A heater must be added to maintain test mass temperature at 123 K, with a heat shield being the recommended approach to reduce the thermal gradient between the IM and TM. For operation at 18 K thermal conduction is the dominant heat transfer mechanism. The incorporation of heat links and the conductive cooling across the monolithic stage become important. A heat link system for the GCIF suspension was designed, inspired by the system in use at the KAGRA detector which utilises the reaction chain to reduce direct linking of the cooling bar to the test chain. Either copper or aluminium heat links would allow the required cooling of the chains up to the IM. The strength limited silicon fibre diameters do not have a high enough thermal conductivity to cool the TM to 18 K so silicon fibres were considered which had a thermal extraction limited diameter, calculated as $d_{\mathrm{HE}} = 650\,\mu\mathrm{m}$. Using the heat extraction limited silicon fibres the TM was modelled to cool to 17.65 K, including bonding and phonon-boundary scattering effects.

Single test mass displacement thermal noise was calculated for six physical scenarios using the method outlined in Chapter 2. This study demonstrates the complexity of suspension thermal noise calculations and the need for future measurements of crystalline material loss, and finite element analysis of mass-fibre jointing possibilities for the GCIF suspension. With the model used in Section 4.4 sapphire fibres displayed a slightly lower thermal noise than silicon fibres at 18 K, and a comparable thermal noise to silicon at 123 K dependent largely on jointing method. The domination of joint loss on suspension displacement thermal noise is clear, leading to the study of a further four un-physical cases which neglected jointing loss. In this scenario silicon fibres displayed a significantly lower thermal noise than sapphire fibres for both temperature cases.

Chapter 5

CHARACTERISATION AND STABILISATION OF 2µm EXTERNAL CAVITY DIODE LASERS

5.1 2 µm Wavelength Lasers for Gravitational Wave Detection

Many 3G gravitational wave detector designs centre around the cooling of cavity optics to cryogenic temperatures to reduce thermal noise, the current limit to detector sensitivity. The fused-silica optics in use at the LIGO detectors exhibit broad peaks in dissipation at cryogenic temperatures, leading to the consideration of sapphire or silicon as the main mirror substrate in future detectors. Silicon's opacity below 1.1 µm requires a change in the wavelength of the main cavity laser from 1064 nm to the 1.4 – 2.1 µm range, and current 3G designs require a single-frequency continuous wave source of a few watts to hundreds of watts of power (with 3 W estimated for the cryogenic interferometer of ET, and 150 W for LIGO Voyager) [8, 182].

The seed lasers in use at current GW detectors are solid-state non-planar ring oscillator (NPRO) lasers with single-frequency 1064 nm wavelength emission. The characterisation of both free-running and stabilised noise has been well documented for these lasers, with a free-running relative intensity noise decreasing from $3 \times 10^{-3}\,1/\sqrt{\mathrm{Hz}}$ to $5 \times 10^{-7}\,1/\sqrt{\mathrm{Hz}}$ between 1 Hz – 10 kHz, and a free-running frequency noise of $100\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 Hz with a $1/f$ slope across the spectrum [121, 122, 210]. These values can be used at the performance 'gold standard' when comparing to developing laser technology.

1.55 µm laser sources are already well-developed for use in telecommunications, with inexpensive lasing sources widely available. Other advantages of this wavelength over 2 µm include the reduction in coating thermal noise due to thinner optical coatings (as coating thickness is proportional to wavelength), lower absorption in the test mass substrate, and photodetector technology already meeting the requirements for quantum efficiency [8, 143]. Despite this, 2 µm may be more suitable in a cryogenic cavity due to the vastly reduced coating absorption, which is strongly wavelength dependent in amorphous silicon coatings [175]. A lower absorption allows for easier maintenance of the cryogenic operating temperature and is considered by the LIGO Voyager design to be the driving factor in wavelength choice [8]. In addition, scattering from optic roughness scales as $1/\lambda^2$ and is therefore reduced for longer wavelengths [7].

With 2 µm a strong contender for use in future gravitational wave detectors there is a clear need for research into laser sources, materials, and optics compatible with this wavelength, such as high quantum-efficiency photodiodes, 2 µm squeezed light, and – most importantly – for optical coatings with low thermal noise. Existing sources between 1.8 – 2.1 µm are holmium- and thulium-based sources, and diode lasers. A holmium-based solid state laser is under development within the GW community [82], while thulium-based fibre lasers are available but exhibit high levels of noise [137]. This creates a current gap in technology for a low-noise and inexpensive 2 µm source which would allow for more accessible research and development into the above fields, as well as potential usage as a seed laser for 3G detectors. External cavity diode lasers (ECDLs) were chosen for investigation in this work as low-noise laser sources which are relatively inexpensive to manufacture and feature a compact, customisable design. These diode lasers are tunable in both wavelength and power, allowing a choice of output wavelength, and actuation for active noise reduction via a range of modulation options.

This chapter describes work undertaken on 2 µm ECDLs both at the University of Glasgow (UofG), and at the Centre for Gravitational Astrophysics at the Australian National University (ANU) as part of a 6-month research project. The ECDLs developed were based on the same design, but have different experimental applications.

The work undertaken at the University of Glasgow aimed to build a stabilised, low-noise 2 µm laser source to be used within the Glasgow Cryogenic Interferometer Facility. This source will be used for probing the silicon cavity to measure thermal noise properties of optical coatings, and to demonstrate control and stabilisation of an interferometer operating at 1.55 µm and 2 µm. The measurement of coating thermal noise at both 1.55 µm and 2 µm is of great interest to the GW community during this phase of development of 3G cryogenic detectors.

The 2 µm laser sources developed as part of this work at the Australian National University will be used to probe coating thermal noise using a folded optical cavity with a 2 µm source beam, cooled to cryogenic temperatures. Previous experiments have attempted to measure coating thermal noise using this method [90], however these experiments were performed at room-temperature and using a 1064 nm wavelength, as typical to room-temperature interferometers with fused-silica optics. Here, the characterisation of two further ECDLs is shown, and the development of stabilisation which demonstrates their suitability for directly measuring the thermal noise of optical coatings in a cryogenic cavity at 2 µm.

The development of three ECDL units provides a comparison between three independently built laser prototypes. This demonstrated the repeatability and stability of characterisation results but also the variability which arises between the different gain chips, and the difficulty of perfect alignment of the lasing cavity.

This chapter first provides an introduction to external cavity diode lasers in Section 5.2, containing a description of the ECDL design used in this work. The characterisation work of the three laser units is presented in Section 5.3 covering laser power output as a function of input current and wavelength, intensity noise measurements, frequency noise measurements, and the response of the laser to modulation signals. The intensity stabilisation work conducted at the University of Glasgow is presented in Section 5.4.1, and the frequency stabilisation work conducted at the Australian National University is presented in Section 5.4.2. A conclusion on the laser development work is given at the end of this chapter in Section 5.5.

5.2 External Cavity Diode Lasers

5.2.1 An Introduction to ECDLs

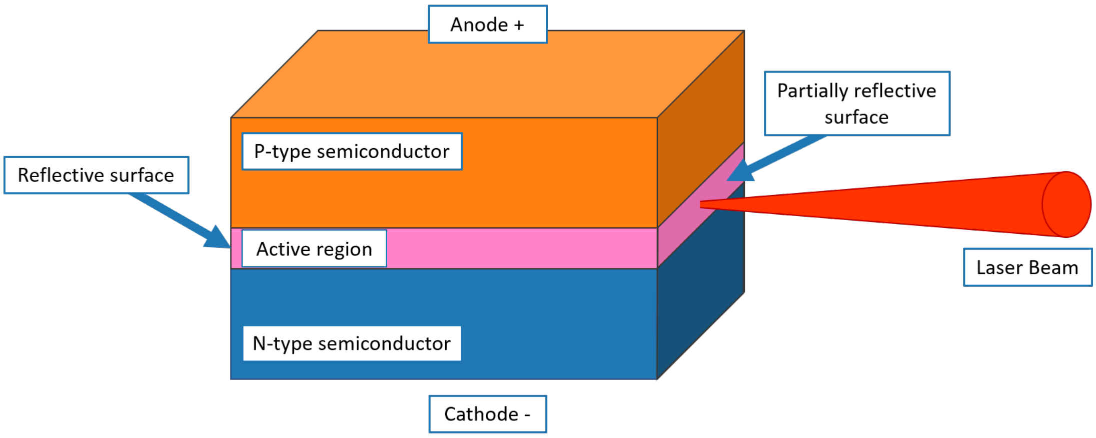

Figure 5.1: Illustration of a laser diode.

A laser diode diagram is shown in Figure 5.1. An undoped semiconductor material known as the *active region* is sandwiched between p-type and n-type semiconductor regions. A voltage applied across the junction pumps electrons and holes from the p-type and n-type regions into this active region where they recombine, releasing energy in the form of photons. The active region refractive index is chosen to create a waveguide, encouraging light propagation perpendicular to the applied current and reflective surfaces at each end of the active region create a Fabry-Perot cavity. The standing waves created by constructive interference in the cavity create the lasing conditions for stimulated emission in the active region. The semiconductor material choice decides the wavelength of the resulting beam, with the possible range of outputs covering infrared all the way to ultraviolet. Laser diodes have many uses in consumer electronics, product manufacturing, metrology, and medical applications, and are the most common type of laser manufactured. They are cheap and compact as well as being extremely durable.

External cavity diode lasers utilise laser diode chips for light generation, contained within a larger cavity with additional optics. In this case, an anti-reflective coating is used in place of one of the gain chip's reflective coatings. The output beam passes through a collimating lens to a diffraction grating and the laser cavity is formed between the diffraction grating and diode chip reflecting surface. The ECDLs used in this work are set-up in the first-order Littrow configuration, as shown in Figure 5.2, where the first-order diffraction from the grating is reflected back into the gain chip [113, 159]. The angular separation of diffraction maxima scales with the wavelength, and therefore retro-reflection only occurs for a narrow wavelength range. The diffraction grating angle can be adjusted with a piezoelectric transducer (PZT) to vary the wavelength of light reflected back to the gain chip, hence allowing wavelength tunability.

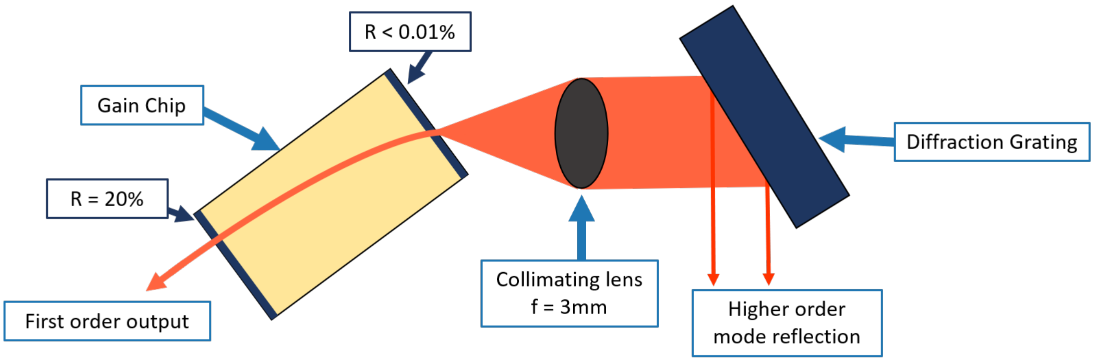

Figure 5.2: Illustration of an external cavity diode laser in the Littrow configuration.

Frequency adjustability is a major advantage of diode lasers, however this property also makes them more susceptible to mode hopping. Laser mode hopping occurs when modes of different frequencies are generated inside the gain medium and compete for the available energy. As one mode is favourably amplified by gain medium conditions and grating angle, this mode quickly grows in amplitude and extracts energy from the gain medium with a higher efficiency. This mode becomes the dominant output, suppressing the other modes. However, the gain medium conditions drift over time due to both environmental and physical factors. This may make conditions more favourable for a new mode, which is then able to compete with the previously dominant mode and cause the laser to switch rapidly (or 'hop') between modes. The laser input current, gain chip temperature, and grating angle can affect the chip's gain conditions. Temperature controllers are commonly used alongside gain chips to stabilise the output mode, with further stability possible through feedback to the laser input current.

5.2.2 The ECDL Design

The ECDL design used throughout this work is a modified version of [27] and was first developed by D. Kapasi at ANU [115]. The Glasgow laser was developed in collaboration with the ANU group and is based on this design, which shown in Figure 5.3 and described in [115] and below.

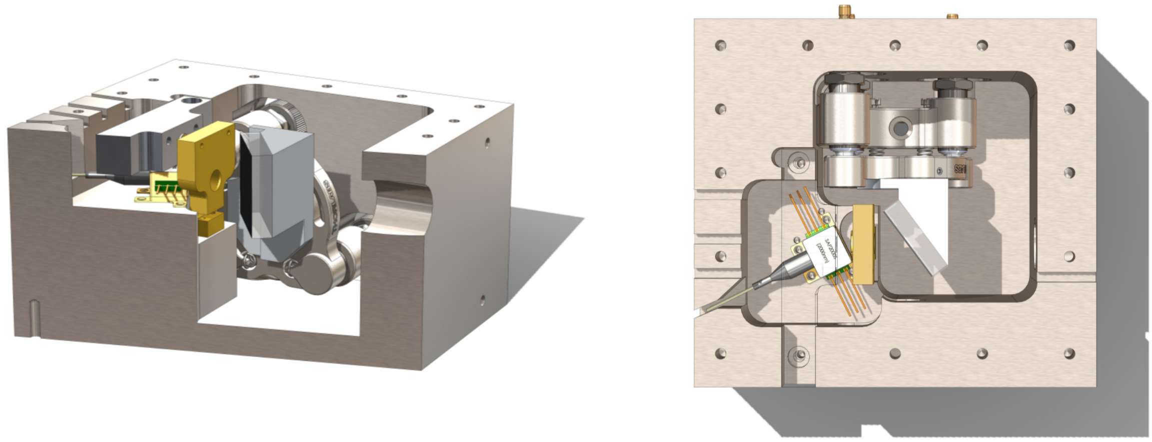

(a) Cut-out side view. (b) Top view.

Figure 5.3: Computer-aided design image of the ECDL layout showing the gain chip, lens mount, and diffraction grating configuration. CAD model by J. Eichholz.

A Thorlabs SAF1900S gain chip was chosen as the laser diode, featuring a curved ridge waveguide to reduce feedback from internal chip reflections into the gain medium, and has a fibre coupled beam output. The chip is wavelength-tunable around a centre wavelength of 1930 nm, with the exact centre wavelength found to vary between units, and the output power varies near linearly with input current between the ~ 100 mA current threshold and the 800 mA maximum operating current. A thermoelectric cooler and thermistor installed in the chip allow for fine control of the operating temperature, crucial for the mode stability of laser diode packages.

The beam exits the gain chip at a $26.5°$ angle and is then collimated by a Thorlabs 390093-D aspheric lens of focal length $f = 3$ mm. This lens is held close to the exit point of the diode in a custom brass mount. The lens mount attaches to the aluminium enclosure via two screws at its base and is adjusted along the optical axis by push-pull screws near its top surface. A Thorlabs GR13-0616 blazed diffraction grating is mounted after the lens to form one end of the cavity. The grating has 600 lines/mm and is highly reflective for the considered wavelengths, with a $\sim 90\%$ efficiency between $1.5 - 3$ µm for light which is perpendicularly polarised with respect to the diffraction lines. This grating is held on a highly stable Polaris kinetic mount with yaw and pitch control via PZT drivers to provide adjustment of the output wavelength.

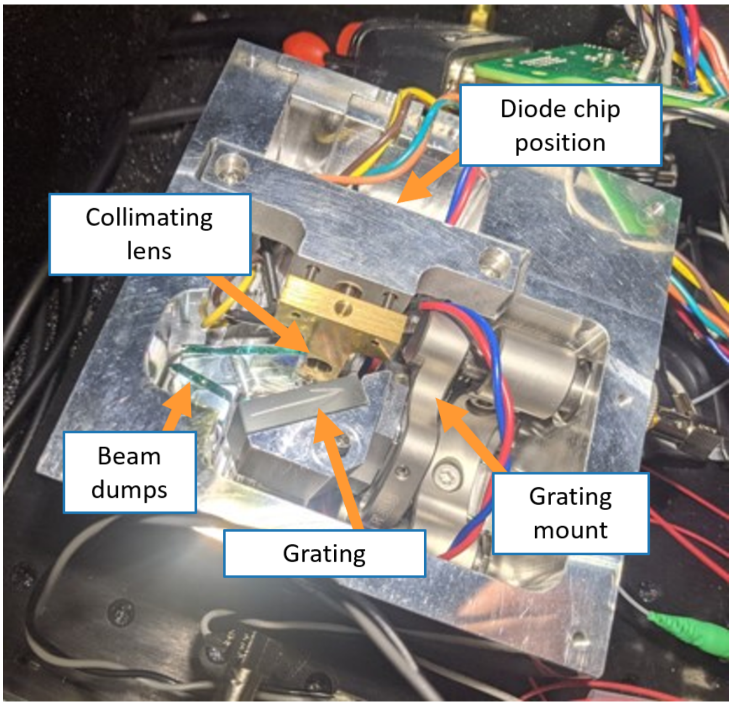

Figure 5.4: Photograph of the as-build ECDL layout for ANU ECDL1.

5.2.2.1 University of Glasgow Build

The ECDL at the University of Glasgow was set-up alongside a Thorlabs TED200C precision temperature controller for temperature stabilisation of the SAF1900S gain chip. A custom table-top frequency stabilisation servo (TTFSS) unit was designed by E. Oelker based on the design in use at the aLIGO detectors, then built and tested by E. Oelker and the author. This unit contains low-noise electronics designed for current drive, interface with computer controls, and response shaping with low noise electronics. An APE waveScan spectrometer was used to ensure the grating alignment produced an output wavelength of 1980 nm.

5.2.2.2 Australian National University Build

Two ECDL units were developed over the course of this project at the Australian National University. ECDL1 was first built by J. Eichholz and J. Jang in 2020 based on the prototype design of [115], with a polarisation maintaining fibre used in place of the single-mode fibre which was coupled to the gain chip during the first prototype build. The external cavity alignment of ECDL1 was modified by the author during the laser characterisation segment of this work, and other slight improvements were made such as the addition of beam dumps inside the laser enclosure to reduce the effect of stray light

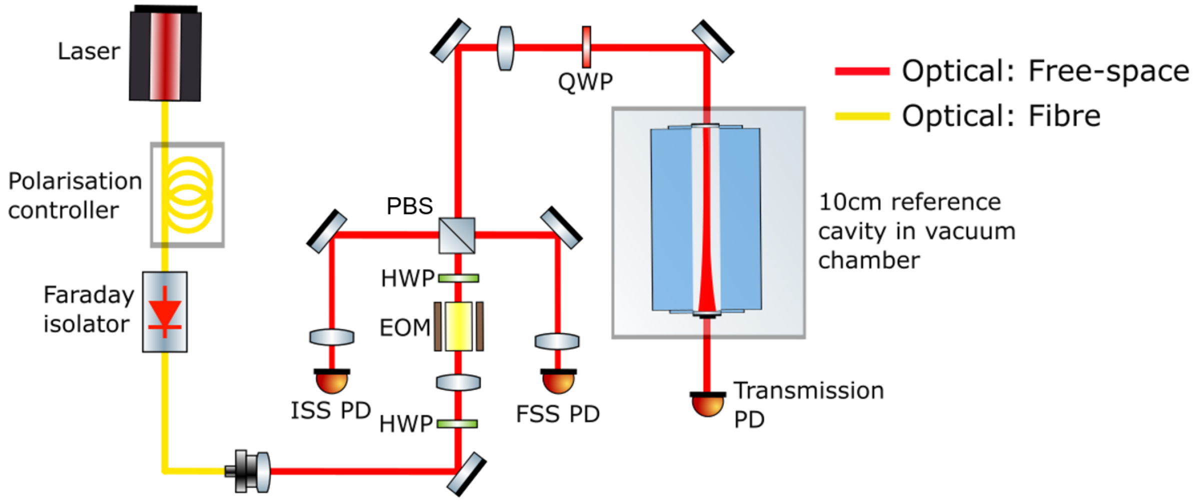

Figure 5.5: Table-top optical layout for laser characterisation and stabilisation at the University of Glasgow. HWP - half waveplate, EOM - electro-optic modulator, PBS - polarising beam splitter, ISS - intensity stabilisation servo, FSS - frequency stabilisation servo, PD - photodiode, QWP - quarter waveplate.

scattering as shown in Figure 5.4. ECDL2 was built by the author in 2023 with the help of J. Eichholz. The design of ECDL2 was similar to ECDL1, however utilised a non-built in PZT adjustor on the grating mount to reduce the assembly cost. For both laser units a Moglabs DLC502 Diode Laser Controller was used to control the gain chip temperature and injection current.

5.3 Laser Characterisation

5.3.1 University of Glasgow Experimental Set-Up

At the University of Glasgow a table-top optical layout was built for laser characterisation and stabilisation. This set-up will be used within the GCIF for a first stage of beam stabilisation before input to the 10 m reference cavity and the cryogenic test cavity. The entire optical set-up is shown in Figure 5.5, with such a set-up allowing the sensing and control of both beam intensity and frequency noise.

The laser output fibre shown in Figure 5.5 goes firstly through polarisation control then through a DK Photonics polarisation-maintaining isolator. After the isolator a 6.2 mm fibre collimator outputs the beam to free-space, where it travels through a QUBIG free-space electro-optic modulator (EOM) to reach a polarising beam splitter. The transmitted

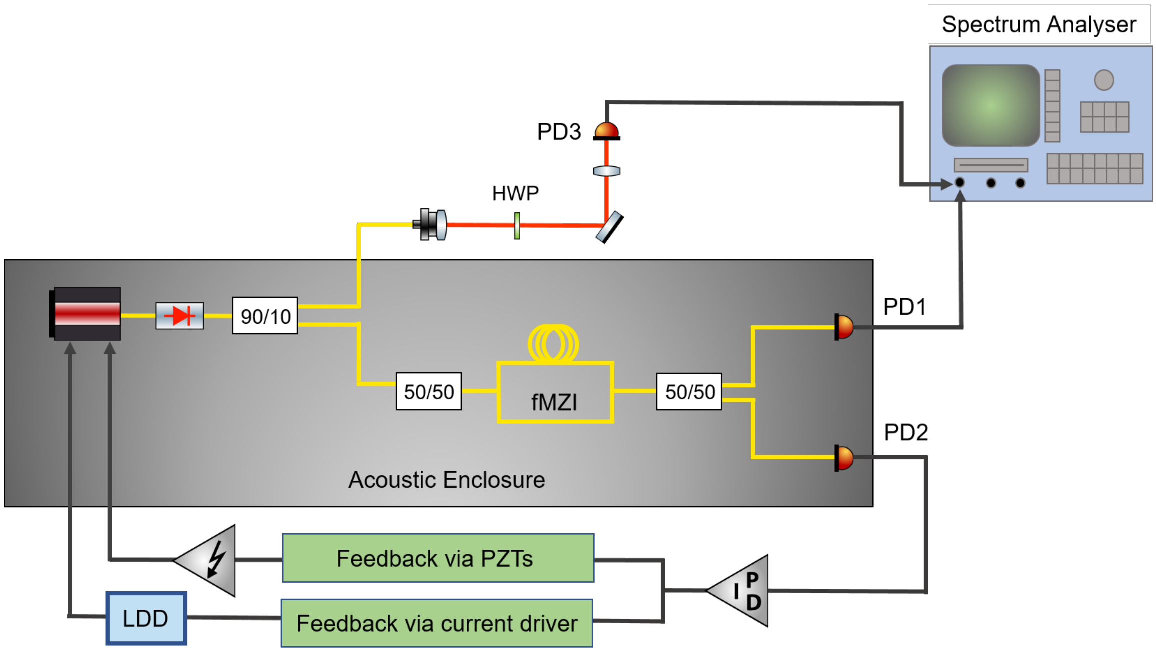

Figure 5.6: Table-top optical layout for characterisation at the Australian National University. An identical layout was created for each of the two lasers, with only one pictured here for clarity. HWP - half waveplate, fMZI - fibre Mach-Zehnder interferometer, PD - photodiode, LDD - laser diode driver

beam travels to the 10 cm optical cavity, and the reflected beam to an intensity stabilisation photodetector (ISS PD): a low-noise extended InGaAs Hamamatsu G12183-005K photodiode, designed for the sensing and stabilisation of intensity noise at this wavelength [152]. The optical cavity is held in a vacuum chamber to reduce air-absorption of the beam and to isolate the cavity from acoustic noise. The signal returning from the optical cavity reaches the beam splitter and is then diverted to the Hamamatsu G12182-003K frequency stabilisation photodiode (FSS PD). However, due to seismic and acoustic noise coupling into the laser significant mode-hopping was observed (as described in Section 5.2) which made it unfeasible to characterise the frequency stability via this method. The laser characterisation work aligned well with the research visit to ANU and therefore the frequency characterisation work was performed on that system. Further investigation of the UofG ECDL frequency noise will adopt the approach taken by ANU with implementation of an acoustic isolation box as well as the inclusion of a built-in fibre Mach-Zehnder interferometer for monitoring and pre-stabilisation.

5.3.2 Australian National University Experimental Set-Up

At the Australian National University an optical layout was built for laser characterisation measurements which is described here and shown in Figure 5.6. This initial layout was later expanded for the creation of a beat note between ECDL1 and ECDL2, and to allow frequency stabilisation. The expanded layout is described in Section 5.4.2. In the characterisation set-up both lasers were placed in an isolated enclosure for temperature stability and acoustic isolation. The diode outputs of both lasers were passed to fibre Faraday isolators to protect the laser cavity from back reflections, then to 90/10 fibre splitters. 90% of the lasers output power was diverted outside the isolation box to a fibre collimator and aligned to a G12182-010K photodiode (PD3) for measurements of laser intensity noise and amplitude response to modulation. The remaining 10% was passed to a fibre Mach-Zehnder interferometer (fMZI) with a 20 m path length difference which provided a frequency reference for the measurements. The signals from the fMZI paths were combined and again split 50/50 with one output leading to a Thorlabs PDA10D2 photodetector (PD1) for measurements of frequency noise and frequency response, while the second output signal (Thorlabs PDA10D2 photodetector PD2) was used in a low-bandwidth feedback loop which adjusted either the laser current or PZT grating angle using a Liquid Instruments Moku:Lab PID controller to keep the fMZI output at mid-fringe.

5.3.3 Laser Output Power

The optical power output was measured as a function of input laser current by attaching the output fibre of the UofG ECDL's Faraday isolator to a Thorlabs digital power meter fitted with a S148C sensor, and adjusting the input current on the laser control box. The results are shown in Figure 5.7 (a) for the lasers tuned to a wavelength of 1980 nm.

Performing measurements after the Faraday isolator protects the gain chip from back-reflected light re-entering the chip, but does introduce loss to the system. The single-stage polarisation-maintaining isolator contributes a 1.2 dB loss with a further 0.25 dB loss considered from each of the two FC/APC connector joints. The output power measured after the Faraday isolator is presented in pink, with dark blue representing the power measurements which are modified to compensate for the 1.7 dB introduced loss. The results discussed here are for the loss-compensated power measurements.

The input current was increased from 0 – 786 mA. Above the threshold current the power increased with increasing current up to a maximum of 9.98 mW at 786 mA which was limited by gain chip maximum input current. A straight line fitting in the 100 – 500 mA linear region resulted in a threshold current of 96.5 ± 4.1 mA with a slope efficiency of 16.7 ± 0.4 mW/A. This is a lower efficiency than observed for laser diodes of well-technologically-developed wavelength ranges, but consistent with other 2 µm diode efficiencies. The measured power agrees well with measurements taken of a Thorlabs TLK-1900M: a pre-aligned Littrow configuration which provides the predicted power against input current for the SAF1900S manufacture specifications, taken from [192]. The TLK-1900M is represented by the yellow line in Figure 5.7 (a). The agreement with the measured data implies a good alignment has been achieved between the laser diode, collimating lens, and diffraction grating of the external cavity for this laser unit.

The power dependence on wavelength was measured for both lasers at the Australian National University, with the results presented in Figure 5.7 (b). The 90% output of the fibre splitters were connected to a collimator and aligned in-air to the S148C sensor of a Thorlabs digital power meter and to a Bristol Instruments 671 wavelength meter, with a mirror on a 90° flip mount used to switch alignment between the two instruments. The grating angle was adjusted with the PZTs in pitch and yaw to maximise laser output power at 1984 nm with an injection current of 400 mA. The diffraction grating angle was then swept through yaw to vary the wavelength between 1840 - 2040 nm where the output power was measured at 5 nm increments. The data points include the loss of the Faraday isolator, fibre connectors, and 90/10 splitter.

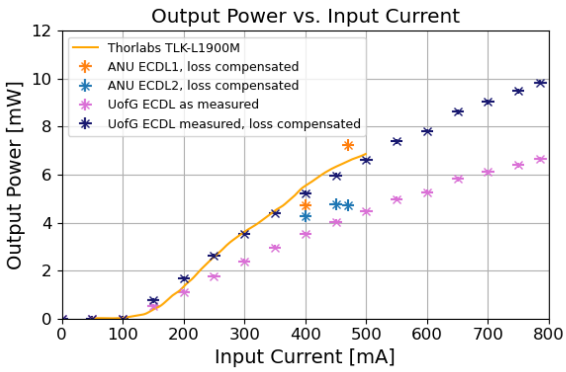

(a) Output power vs. input current.

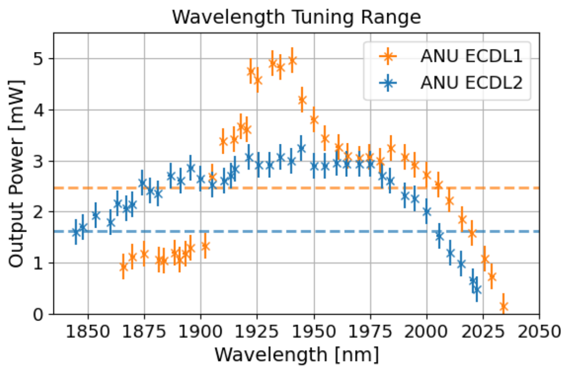

(b) Wavelength tuning range.

Figure 5.7: (a) Measured output power as a function of input current. Pink shows the measured data for the UofG ECDL, with dark blue showing the power compensated for isolator loss. The ANU lasers are shown in blue and orange, also loss compensated. The gain chip specifications measured on a pre-aligned Littrow configuration are shown in yellow [192]. (b) The wavelength tuning range measured for the ANU ECDLs, shown as measured power against wavelength. ECDL1 is shown in orange and ECDL2 in blue, with the dashed lines of each colour representing the FWHM of each data set.

The coarse wavelength tuning ranges of the two lasers span a similar range, however there is a large variation in output powers. ECDL1 has a maximum power of 4.95 ± 0.25 mW at 1940.5 nm, and a 15.2 dB drop between the peak and minimum measured power values. There is a peak in power output between 1920 – 1940 nm. ECDL2 has a maximum power of 3.24 ± 0.25 mW at 1944.7 nm and has a 8.2 dB variation between the maximum and minimum measured powers. There is less power variation across the ECDL2 tuning range, and no obvious peak appears. The ECDL1 peak indicates a possible misalignment of the grating pitch angle. This would affect the modes resonant in the cavity as the yaw angle is adjusted, causing mode hopping during the measurement and 'jumping' variations in power. ECDL1 has a full width at half maximum (FWHM) (-3 dB) tuning range of 103 ± 5 nm around a centre wavelength of approximately 1955 nm, while ECDL2 has a FWHM tuning range of 157 ± 5 nm with a centre wavelength of approximately 1925 nm. The FWHM are represented as dashed lines in Figure 5.7.

The output power of the ANU ECDLs are shown in Figure 5.7 (a) at 400 mA, 450 mA, and 480 mA input currents, with current output limited by the current driver. These data points were taken during measurement of the wavelength tuning range with the lasers operating at 1984 nm and were adjusted to account for the 1 dB Faraday isolator insertion loss, two fibre connections each of 0.25 dB insertion loss, and 90% splitter to allow comparison to the TLK-L1900M curve in this plot. It can be seen that the three lasers largely agree with the TLK-L1900M specifications but display varying power output values for a given input current. This output is dependent on gain chip variability and alignment of the external cavity. This fact is evident throughout the course of this work: each laser displays properties related to the external cavity alignment and the individual gain chip unit, with exact reproducibility between the units difficult to achieve.

5.3.4 Relative Intensity Noise

The intensity noise of the UofG ECDL was measured by aligning the in-air beam onto the ISS PD shown in Figure 5.5, with the DC output connected to a Stanford Instruments SR785 signal analyser. The half waveplates (HWPs) were adjusted to ensure correct polarisation, maximising the power incident upon the photodetector. The spectrum analyser was used to take FFT (Fast Fourier Transform) measurements from the DC photodetector output for frequency spans of 128 Hz – 102.4 kHz, 16 Hz – 12.8 kHz, and 2 Hz – 1.6 kH, both for photodetector dark noise and laser intensity noise. The measurement spans were later spliced together to form the full span measurement. Relative intensity noise (RIN) is calculated by normalising the photodiode output by the DC voltage of the photode-

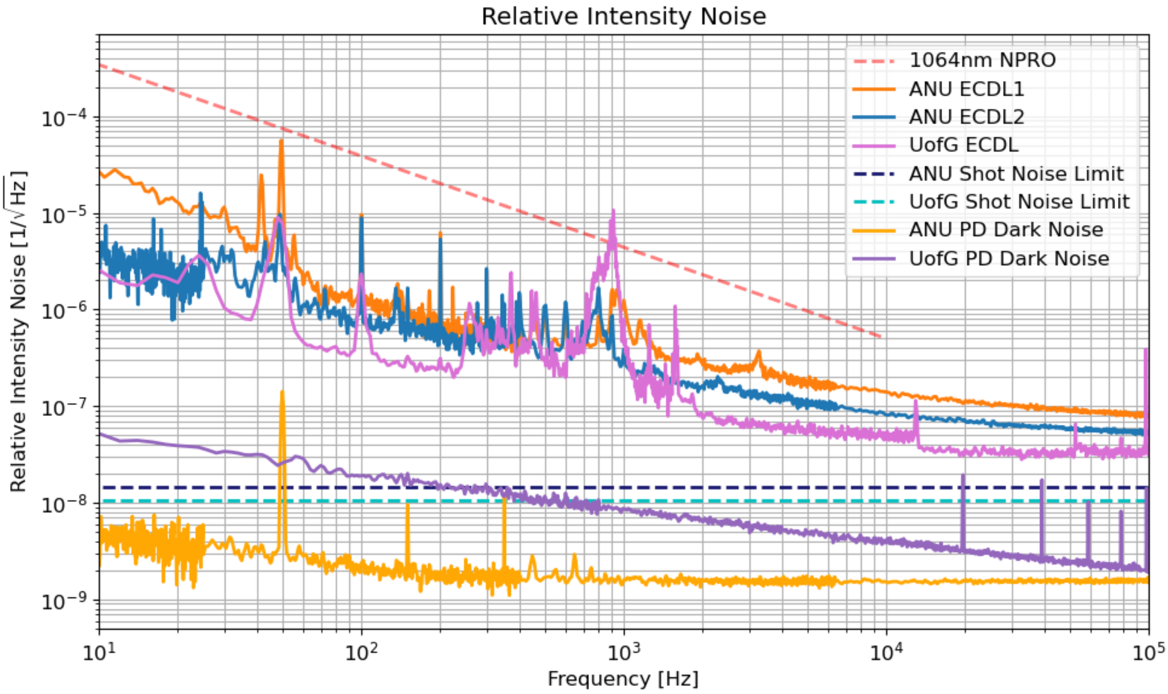

Figure 5.8: Relative intensity noise of the three ECDLs. The results for the UofG unit are shown in pink, ANU ECDL1 in orange, and ANU ECDL2 in blue. The instrument noise floor of the ANU measurement is shown in yellow and the dark noise of the UofG photodiode in purple. Shot noise calculated for each experiment is represented by dashed lines. The relative intensity noise of a 1064 nm NPRO laser is shown for comparison in dashed red.

tector for units of $1/\sqrt{\text{Hz}}$. The injection current of the ECDL was set to 450 mA and the laser tuned to a wavelength of 1980 nm. The shot noise limit of the measurement can be calculated as

$$h_{\text{shot}} = \frac{\sqrt{2\,e\,I_{\text{in}}}}{I_{\text{in}}} \tag{5.1}$$

for e electronic charge and I_{in} the photodiode current [19]. The shot noise limit of the measurement is $1.1 \times 10^{-8}\,1/\sqrt{\text{Hz}}$, represented as a cyan dashed line in Figure 5.8. In the same figure the results of the dark noise measurement are shown in purple and the UofG ECDL relative intensity noise in pink.

The characterisation set-up of Figure 5.6 was used for the relative intensity noise measurement of the ANU ECDLs. For each ECDL the output was coupled to a 90/10 optical splitter. The 90% output led to a collimator with the in-air beam aligned onto PD3 via a HWP, steering mirror, and focusing lens. The ECDLs were supplied with a 450 mA input current and set to a 1984 nm wavelength. A SR785 spectrum analyser took FFT measurements from the photodetector output for frequency spans of 102.4 kHz, 6.4 kHz, 400 Hz, and 25 Hz. These were divided by the DC voltage of the photodetector to form the resulting measurements shown in Figure 5.8 in orange for ECDL1 and blue for ECDL2. The two ANU ECDLs were measured separately by this method.

The comparison of relative intensity noise for the three ECDL units is shown in Figure 5.8. The three units display similar profiles showing intensity noise decreasing with increasing frequency, and a resonance peak around 1 kHz which arises from a mechanical resonance within the laser mount [115]. The smaller amplitude of this resonance peak for the ANU ECDLs is attributed to the acoustic enclosure around these units, which was not present in the UofG measurement. In both set-ups the instrumental noise floor and shot noise are not limiting the intensity noise measurement: each output is the intensity noise of the measured laser. The gradient of the intensity noise drop-off with frequency matches well between the three laser units as $1/\sqrt{\mathrm{f}}$ from 1 Hz – 10 kHz then flattening out from 10 kHz – 100 kHz.

A similar magnitude of noise is seen between the units across the measurement range, with the UofG ECDL displaying a slightly lower intensity noise across the span aside from larger peaks between 200 Hz - 1.1 kHz, and especially at the 1 kHz mechanical resonance. It is assumed that the addition of an acoustic enclosure would reduce the noise peaks in this region, as was observed at ANU. At 20 kHz (away from the mid-range peaks) the UofG unit displays a relative intensity noise 3.4× lower than ANU ECDL1, and 2.2× lower than ANU ECDL2. This difference is of the magnitude which could be caused by differences in the alignment of the external cavity, the individual gain chips, or by the dominant resonant mode in the cavity. The jump in noise of the UofG unit at 1.6 kHz and 12.8 kHz is attributed to such mode hopping occurring between the measurements. This section shows the first true intensity noise measurement for these 2 µm ECDLs not limited by photodiode dark noise [115]. All three units display a lower relative intensity noise than the 1064 nm NPRO lasers, which have a performance of $3 \times 10^{-3}\,1/\sqrt{\mathrm{Hz}}$ to $5 \times 10^{-7}\,1/\sqrt{\mathrm{Hz}}$ between 1 Hz – 10 kHz as mentioned in Section 5.1. This shows promise for future development of 2 µm ECDLs as interferometer seed sources.

5.3.5 Frequency Noise

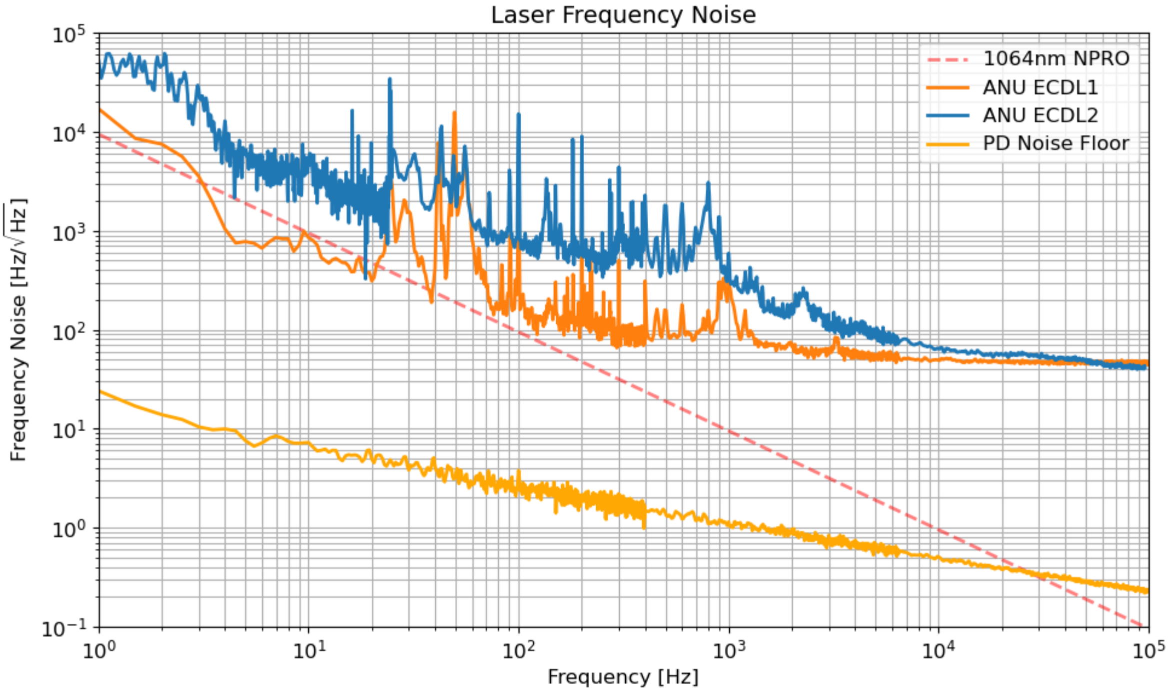

Figure 5.9: Frequency noise of the two ANU ECDLs. ANU ECDL1 is represented in orange, ANU ECDL2 in blue, and photodetector dark noise (Thorlabs PDA10D2) in yellow. The frequency noise of a 1064 nm NPRO laser is shown for comparison in dashed red.

Laser frequency noise was measured for the ANU ECDLs with the method outlined in [115] and the set-up shown in Figure 5.6. For each ECDL the output was coupled to a 90/10 optical splitter. The 10% output was fibre coupled to a 50/50 splitter which had both outputs leading to a fibre Mach-Zehnder interferometer with a 20 m path length difference. Using the output from PD2, the ECDL was locked with a low-bandwidth to the mid-fringe of the fMZI via a Moku:Lab PID controller providing feedback to the laser injection current. PD1 senses frequency shifts of the laser, $\delta\nu$, via differential phase fluctuations, $\delta\phi$, as

$$\delta\nu = \frac{10\text{MHz}}{2\,\pi}\,\delta\phi \tag{5.2}$$

with 10 MHz the fMZI free-spectral range [115]. The spectrum analyser took FFT measurements from the photodetector output for frequency spans of 102.4 kHz, 6.4 kHz, 400 Hz, and 25 Hz for ECDL2, and spans of 102.4 kHz, 6.4 kHz, and 400 Hz for ECDL1. The resulting outputs are shown in Figure 5.9 for ECDL1 in orange and ECDL2 in blue. Photodetector dark noise is shown in yellow in Figure 5.9.

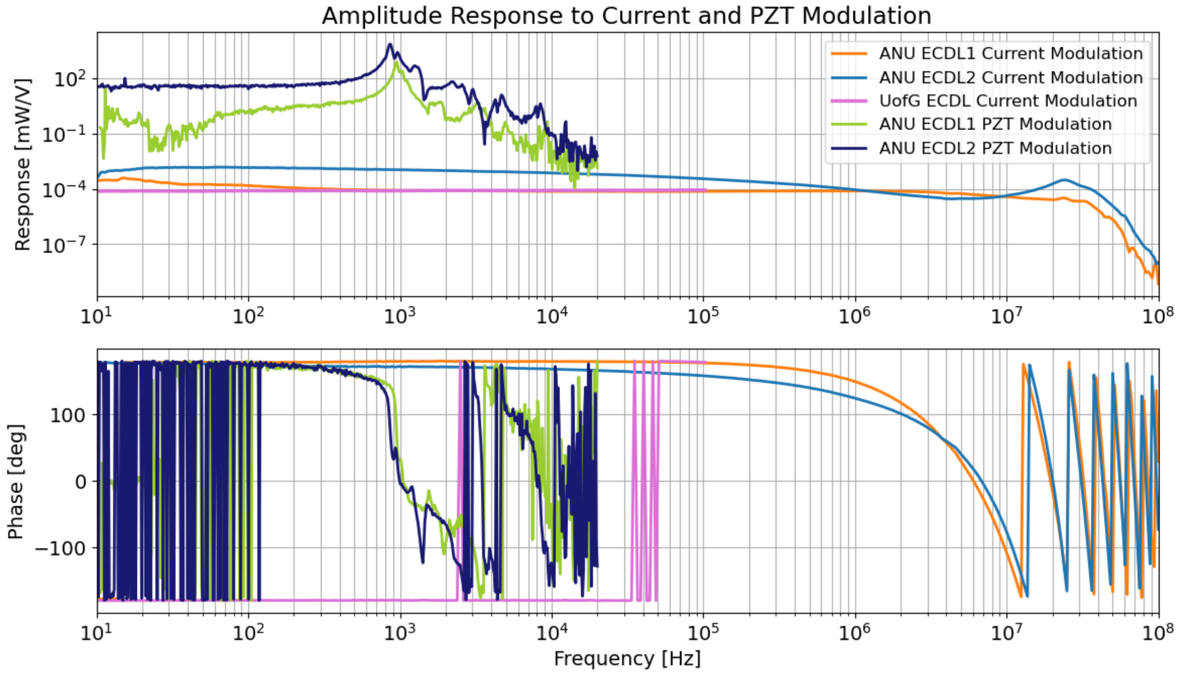

Figure 5.10: Amplitude response to current and PZT modulation. The response to current modulation is represented in pink for the UofG ECDL, in orange for ANU ECDL1, and in blue for ECDL2. The response to PZT modulation is shown in green for ANU ECDL1 and in dark blue for ANU ECDL2.

The frequency noise of both units is not limited by the noise floor of the photodetector. The magnitude of the ECDL1 frequency noise falls as $1/\sqrt{\mathrm{f}}$ from 4 Hz – 700 Hz and displays a flatter profile from 700 Hz – 100 kHz. ECDL1 has a lower frequency noise across most of the measurement span (1 Hz – 40 kHz) when compared to ECDL2. ECDL2 shows a $1/\sqrt{\mathrm{f}}$ drop in noise across a wider frequency span than ECDL1, from 4 Hz – 10 kHz before similarly flattening out as frequencies increase to 100 kHz. Above 40 kHz the noise levels of the two lasers tend to a similar value where ECDL1 displays a constant noise of 45.6 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ from 5 kHz – 100 kHz, and the frequency noise of ECDL2 slightly decreases with increasing frequency from 80.1 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 5 kHz to 40.2 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 kHz. From Section 5.1 the 1064 nm NPRO laser frequency noise is quoted as 100 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 Hz with a $^1/_f$ slope across the spectrum . The ECDL free-running frequency noise level is comparable to that of the NPRO lasers from 1 Hz – 100 Hz, however at higher frequencies the ECDL frequency noise flattens out to a constant value while the NPRO noise continues to fall as $^1/_f$. The ECDLs therefore have good performance at low frequency, but currently display excess high frequency laser frequency noise.

5.3.6 Modulation Transfer Functions

The frequency and amplitude output of the ECDLs can be controlled via modulation of the diode injection current, and by the angle of the grating forming the laser cavity. This control is necessary for stabilisation of the laser and therefore it is important to measure the effects of these modulations before the design of a noise suppression feedback loop. The modulation response and bandwidth must be noted, alongside the characterisation of any resonance peaks which could affect the loop stability.

At the Australian National University the amplitude response and frequency response functions of the ECDLs were measured for current modulation, and for grating angle modulation via the PZTs. Amplitude response was measured with the set-up of Section 5.3.4, using a Thorlabs PDA10D2 to record amplitude response. The frequency response was measured with the set-up of Section 5.3.5 which required a feedback loop to the fMZI with either current or PZT modulation holding the signal at the mid-fringe point. Feedback to the injection current was used to lock the ECDL to the fMZI during measurement of the response to PZT modulation, whereas feedback to the PZTs was used to lock the ECDL to the fMZI during measurement of the response to current modulation. The frequency response in both locking cases was measured using PD1. For both amplitude and frequency response measurements the SR785 spectrum analyser, with a maximum bandwidth of 102.4 kHz, was replaced with the frequency response analyser of the Liquid Instruments Moku:Lab to allow the measurements of the modulation response functions to be taken up to 120 MHz. Measurements of the response to modulation of the injection current was achieved through a 10 Hz – 100 MHz swept sine measurement from the frequency response analyser, with drive signal injected into the current driver. Measurements of the response to modulation of the grating PZTs was achieved through a 10 Hz – 20 kHz swept sine measurement from the frequency response analyser, with the drive signal injected before the high-voltage amplifier shown in Figure 5.6.

The results of the ECDL amplitude response to current and PZT modulation are presented in Figure 5.10, in units of mW/V as a measure of the power fluctuation over modulation drive voltage. For the amplitude response to current modulation, the dB response of the photodetector was converted to mW/V by consideration of the voltage-to-current response of the current driver, the current-to-power response of the laser unit, and the photodiode optical response of power-to-voltage. For the case of amplitude response to PZT modulation the input voltage to the PZTs was compared to the output photodetector

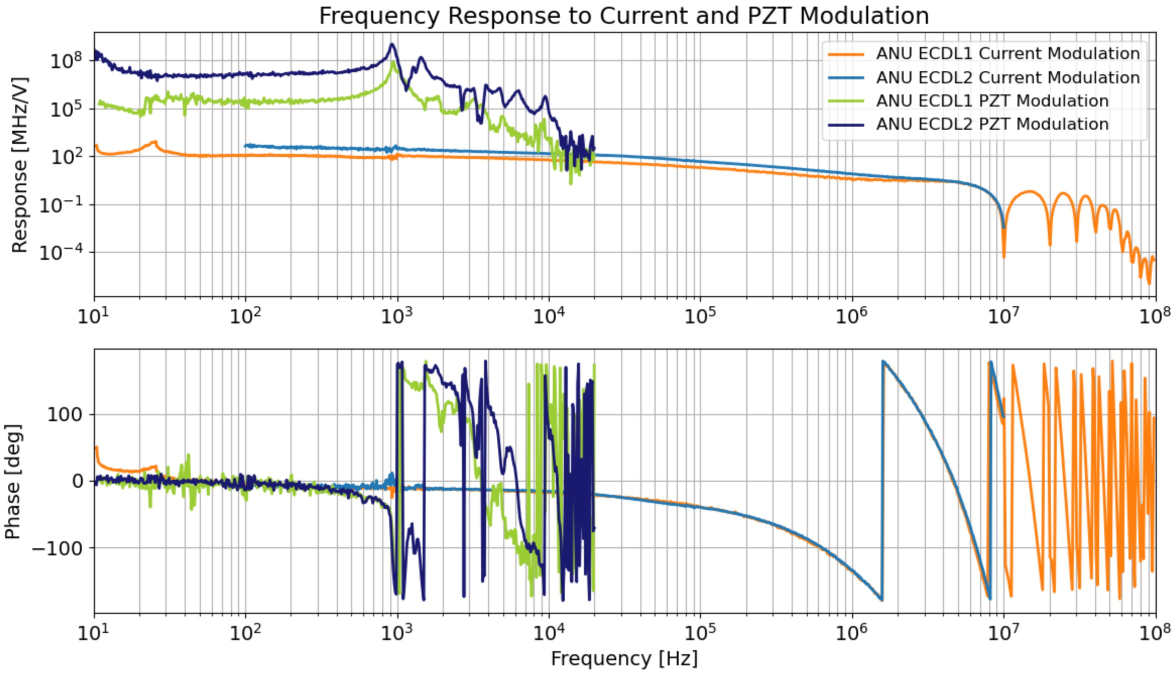

Figure 5.11: Frequency response to current and PZT modulation. The response to current modulation is represented in orange for ANU ECDL1, and in blue for ECDL2. The response to PZT modulation is shown in green for ANU ECDL1 and in dark blue for ANU ECDL2.

voltage observed, which was combined with the known power-to-voltage response of the photodetector to output units of mW/V. The orange and blue traces show the amplitude response to current modulation for ECDL1 and ECDL2 respectively, while the green and navy trances show the amplitude response to PZT modulation for each ECDL.

The results of the frequency response to current and PZT modulation are shown in Figure 5.11. Units of MHz/V display frequency fluctuations over modulation drive voltage, calibrated by the fMZI response. The same colour convention is used as for Figure 5.10, this time displaying the frequency response of the lasers.

At the University of Glasgow the ECDL's amplitude response to current modulation was measured, necessary for intensity stabilisation. The experimental set-up of Figure 5.5 was used with the amplitude response measured at the ISS PD. An SR785 signal analyser was used to perform a swept sine frequency response measurement across a 100 kHz span with the modulation signal input at the laser current controller. The output of the photodetector was converted to mW/V by consideration of the voltage-to-current response of the current driver, the current-to-power response of the laser unit, and the photodiode optical response of power-to-voltage. The resulting amplitude response is a flat 8×10^{-5} mW/V for the entirety of the 100 kHz span, with bandwidth limited by the SR785 signal analyser, and is represented by the pink trace in Figure 5.10.

The transfer functions of amplitude response to current modulation measured at ANU agree well between the two units, with 7.5×10^{-5} mW/V for ANU ECDL1, while ECDL2 has a slightly higher amplitude response of 7.8×10^{-4} mW/V. The two ANU ECDLs display a flat amplitude response to current modulation up to 30 MHz, limited here by the bandwidth of the photodetector.

The PZT modulation amplitude response of ANU ECDL2 has a flat profile of 39.8 MHz/V from 10 Hz up to the characteristic 1 kHz mechanical resonance. ECDL1 has a lower amplitude response than ECDL2, with a level of 1.5 MHz/V at 100 Hz. This unit also displays a noisy response between 10 – 60 Hz which, in combination with the tuning range output from Figure 5.7, suggests that there is some pitch or roll misalignment of the ECDL1 diffraction grating which impacts the mode resonant in the cavity as grating yaw angle is adjusted. The 1 kHz resonance sets the limitation on the usable bandwidth for PZT modulation and could be mitigated in future ECDL builds through a redesign of the laser mount. Modulation of the grating angle via the PZTs does induce a large response in the amplitude in both lasers, which may affect the method of joint intensity and frequency stabilisation chosen for future developments.

The frequency response to current modulation is 120 MHz/V for ECDL1 and 250 MHz/V for ECDL2 at 1 kHz. The response of both ECDLs is flat across the measurement span, with bandwidth limited by the 10 MHz free spectral range of the fMZI. The frequency responses to PZT modulation have similar profiles to the PZT modulation intensity response, again having a bandwidth limited by the mechanical resonance at 1 kHz. ECDL1 has less noise under 100 Hz than seen in the amplitude response, suggesting that the grating misalignment may be in pitch rather than in roll. ECDL2 has a similar frequency response profile to ECDL1, however it also has a higher response to input modulation, with ECDL1 having a 2.9×10^{5} MHz/V response at 100 Hz, while ECDL2 shows 1.6×10^{7} MHz/V at the same frequency. This could be due to the position of the grating inside the laser mount. If one ECDL has a grating rotational axis offset to the other unit, then sweeping the yaw angle of the grating with a constant input voltage would produce a variation in the response of the two units, changing length of the external laser cavity and the response of the laser to the sweep.

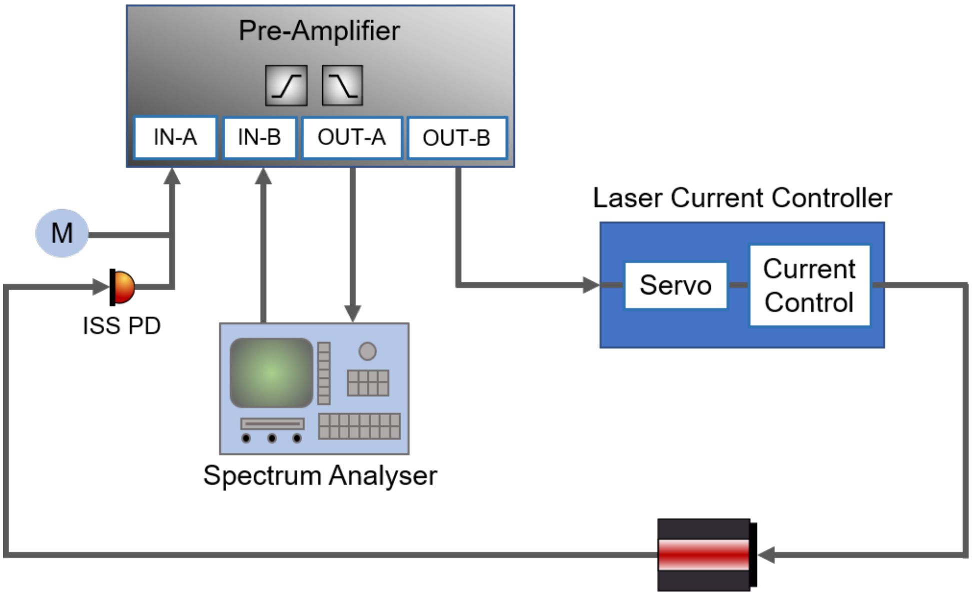

Figure 5.12: Feedback loop for measurements of stabilised intensity noise.

For both amplitude and frequency responses, the phase lag observed at higher frequencies is attributed to phase lag from electronic components within the digital sensing system. Similar transfer functions were observed for all three laser units, demonstrating the repeatability of results. However also evident is the variation between each ECDL unit which is dependent on the build and alignment.

5.4 Laser Stabilisation

5.4.1 Intensity Stabilisation

Intensity stabilisation of a 2 µm ECDL was demonstrated at the University of Glasgow. The table-top experimental set-up is as shown in Figure 5.5, with the feedback loop shown in Figure 5.12. The laser intensity noise is measured by the ISS PD shown in Figure 5.5. This PD was connected to a servo loop and the laser intensity was stabilised through current modulation. FFT measurements of free-running laser intensity noise were taken after the ISS PD at point 'M', as shown in Figure 5.12. To close the feedback loop and allow exploration of the gain required for stabilisation the output of the ISS PD was input to a Stanford Research Systems SR560 low-noise pre-amplifier. With the pre-amplifier set to A-B the spectrum analyser took a swept-sine measurement of the in-loop (servo locked) transfer function (M1) and the out-of-loop (servo unlocked) transfer function (M2) with the configuration shown in Figure 5.12 at a pre-amplifier gain setting of both 1000× and

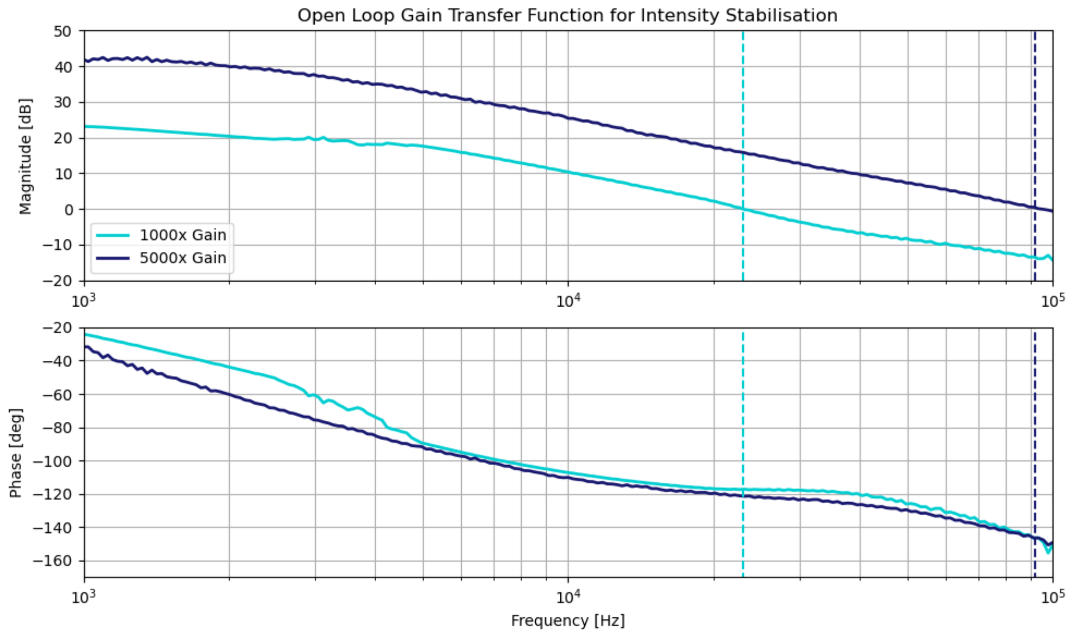

Figure 5.13: Open loop gain transfer function for intensity stabilisation, shown for a 1000× gain setting in cyan and a 5000× gain setting in dark blue. The unity gain frequencies are represented by dashed vertical lines following the same colour scheme.

5000×, with 5000× the maximum gain useable before overloading of the pre-amplifier. The in-loop and out-of-loop transfer function measurements of the servo were combined to produce the open loop gain of the servo as $\frac{\mathrm{M2}}{\mathrm{M1}} - 1$, with the gain and phase of the open loop transfer function shown in Figure 5.13 (the closed loop transfer function can be calculated via this method as $1 + \frac{\mathrm{M1}}{\mathrm{M2}}$). The transfer function of the 1000× gain setting displays a unity gain frequency (UGF) at 23 kHz with a phase margin of 62°, which increases to 92 kHz for the 5000× gain setting with a phase margin of 33°. The phase margin of the system determines the stability, defining the additional phase lag required to bring the system to instability. A system with a wider phase margin is more stable, with a phase margin of over 45° typically used and shows here that the 1000× gain setting has a higher stability [191]. With a 5× higher gain, the 5000× gain setting does not quite reach a UGF of 5× 23 kHz which can be attributed to the beginnings of slew-rate limiting in the SR560 pre-amplifier. The intensity response to this current modulation demonstrates a flat response for the frequencies of interest and the measured UGFs show there is an adequate bandwidth for the intensity stabilisation.

The spectrum analyser was used to obtain FFT measurements of the in-loop laser intensity noise with the set-up shown in Figure 5.12. Measurements were taken for frequency spans of 2 Hz – 1.6 kHz, 16 Hz – 12.8 kHz, and 128 Hz – 102.4 kHz for both pre-amplifier gain settings, with the measurement span limited by the bandwidth of the spectrum analyser. The results of the intensity stabilisation are shown in Figure 5.14, normalised by

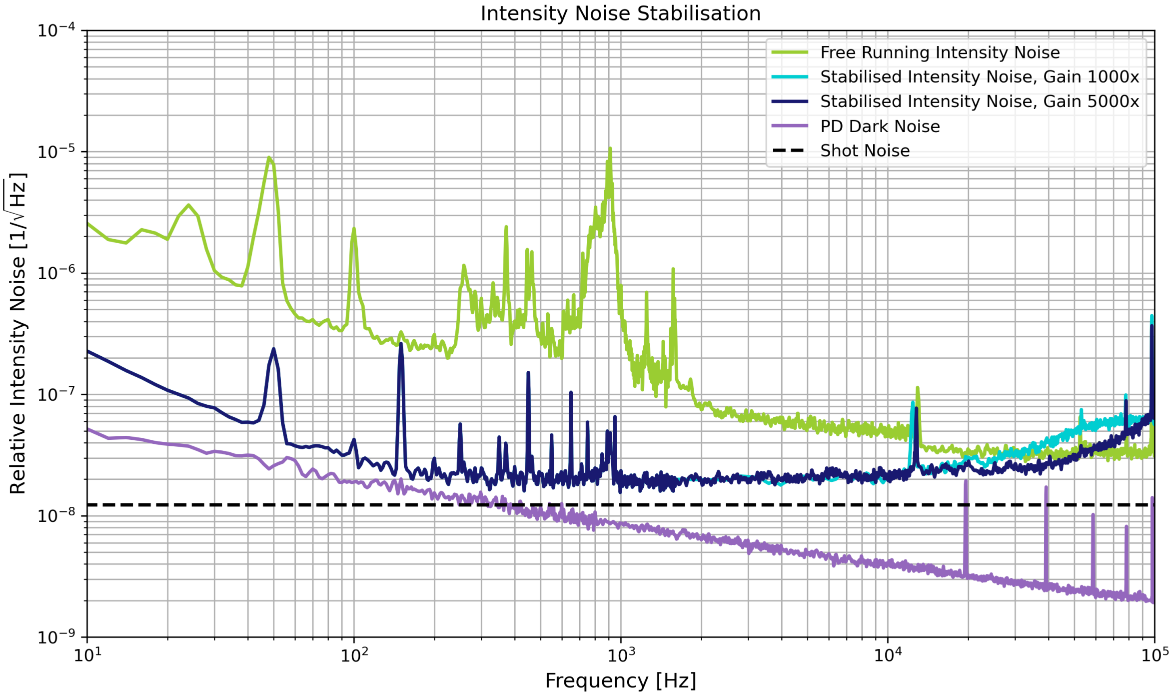

Figure 5.14: Stabilised intensity noise measurements taken at the University of Glasgow. The free-running laser intensity noise is shown in green, with the stabilised intensity noise displayed in cyan for a 1000× pre-amplifier gain and in dark blue for a 5000× gain. The shot noise is represented by a black dashed line and photodiode dark noise by the solid purple.

photodiode voltage to units of relative intensity noise ($1/\sqrt{\text{Hz}}$). The green trace displays free running intensity noise and photodetector dark noise as seen in Figure 5.8. The shot noise for the measurement is calculated from Equation 5.1 and is represented by the black dashed line. Stabilised intensity noise is shown for two pre-amplifier gain settings, showing the maximum possible gain setting of 5000× in dark blue, and a more stable 1000× gain in cyan which is plotted from 1 kHz – 100 kHz.

A reduction in intensity noise is demonstrated between 10 Hz – 50 kHz, achieved via feedback to the laser current. At low frequencies the measurement appears limited by photodetector dark noise from 5 Hz – 30 Hz. Above this the stabilised intensity noise is flat with frequency until 10 kHz, with a level of $2 \times 10^{-8}\,1/\sqrt{\text{Hz}}$. In this stabilised band between 40 Hz – 10 kHz the intensity noise lies only 1.8× higher than the shot noise limit. Below 10 kHz the two gain settings output the same level of stabilised intensity noise, which is not as expected from the magnitude of the two open loop gain transfer functions, and the noise suppression for both gain settings is lower than expected. This suggests there is a limitation to the noise suppression in the band between PD dark noise and the gain UGFs (30 Hz – 23 kHz for 1000× gain and 30 Hz – 92 kHz for 5000× gain). Possible origins for this noise could arise in feedback electronics, beam jitter in the laser enclosure,

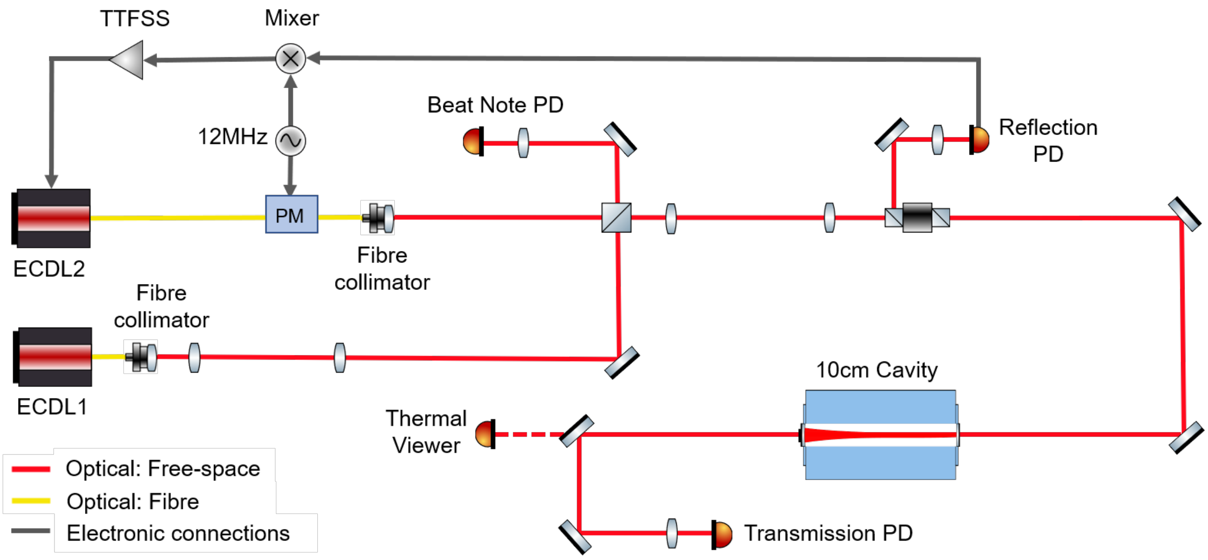

Figure 5.15: Table-top and feedback layout for the ECDL2 frequency stabilisation measurements. PM is the fibre phase modulator and the TTFSS the table-top frequency stabilisation servo. The red and yellow lines represents the beam path in air and fibre respectively, and the grey is servo electronic connections.

or polarisation noise in the fibres. Above the UGF of each gain setting 'gain bulging' is observed whereby the stabilised intensity noise is higher than the free running intensity noise. This is expected in noise suppression loops above the UGF where a low phase margin can create positive feedback, increasing noise output.

The in-loop stabilised intensity noise demonstrates proof of intensity noise stabilisation for a 2 µm ECDL with a maximum unity gain frequency of 92 kHz, close to the shot noise limit of the photodetector. The use of the SR560 pre-amplifier is the most likely source of electronic noise in the feedback loop. The use of the pre-amplifier in the loop was used for convenience in this proof-of-concept stabilisation experiment and was not intended for use in the final design. The servo electronics in the current controller box will be reconfigured in future to optimise for laser intensity stabilisation now that the proof of concept has been demonstrated. At this point, out-of-loop intensity noise can be measured to characterise the stabilised intensity noise which could be input to the Glasgow 10 m interferometer.

5.4.2 Frequency Stabilisation

Frequency stabilisation of a 2 µm ECDL was demonstrated for the first time at the Australian National University. The experimental set-up of Figure 5.15 was used to mode-match, align, and lock each ECDL to a 10 cm Fabry Perot optical cavity via Pound-Drever-Hall (PDH) locking: a high-performance frequency stabilisation method which locks the

ECDL to the optical cavity via phase modulation [69]. This method has a wider range of lock compared to the previously used fMZI mid-fringe locking, and allows high precision, narrow linewidth stabilisation. Locking of the laser to a cavity resonant mode would be simpler to implement than PDH, however the laser stabilisation would be more subject to environmental factors such as cavity thermal fluctuations, which are likely with the use of the table-top in-air cavity configuration.

The 10 cm optical cavity has a $\Delta\nu_{\mathrm{FSR}}$ =1.5 GHz free spectral range and a finesse of $\mathscr{F} = 10,000$. The predicted cavity linewidth can be calculated using the finesse and free spectral range as

$$\Delta\nu_{\mathrm{c}} = \frac{\Delta\nu_{\mathrm{FSR}}}{\mathscr{F}} = 150\,\mathrm{kHz} \tag{5.3}$$

to show the width of the frequency band which is resonant in the cavity [30]. For two mirrors of an identical reflectivity, R, cavity finesse is defined as [211]

$$\mathscr{F} = \frac{\pi\sqrt{R}}{1-R}. \tag{5.4}$$

A photodetector placed behind the cavity was used to sense the transmitted cavity light, with the cavity linewidth then measured as the FWHM of the transmitted cavity light, giving $\Delta\nu_{\mathrm{c}} = 136 \pm 4\,\mathrm{kHz}$ using ECDL2. Injected sidebands of a known frequency spacing were used to scale the measurement axes to produce this output. Using Equation 5.4, the difference between the predicted and measured cavity linewidth values could be caused by a difference in mirror reflectivity of only 29 ppm, which is within the manufacturing tolerance for mirror reflectivity, therefore the ECDL2 measured cavity linewidth matches well with the predicted value.

The measurement was also attempted using ECDL1 however due to mode-hopping the cavity peaks were not stable enough to allow a linewidth measurement. From the results of Section 5.3, this is likely due to a small misalignment of the ECDL1 grating pitch angle, or could be due to retro-reflections to the diode chip creating parasitic cavities in the system.

The reflected cavity signal from the Thorlabs PDA10D2 photodetector was used to lock the ECDL to the optical cavity via PDH locking with the feedback loop shown in grey in Figure 5.15, using a frequency mixer (Mini-Circuits Zx05-1MHW-S+), waveform generator (Rigol DG4102), and a modified aLIGO-style TTFSS [165]. A fibre phase-modulator (PM) was added to the fibre path of both lasers before the in-air collimator. The PM inputs a $\Omega = 12\,\mathrm{MHz}$ phase modulation to the laser frequency (where modulation frequency must be much greater than the cavity linewidth), creating a signal with a series of modulation sidebands separated by Ω from the laser carrier frequency, ω. This signal of $\omega \pm \Omega$ is incident upon the optical cavity where the sidebands (off-resonance from the cavity) are reflected from the first cavity mirror with no phase shift, while the carrier enters the cavity and receives a phase shift which is dependent on the difference in frequency between the carrier frequency and cavity resonance.

The superposition of signals reflected from the cavity creates a wave of frequency ω with an envelope which oscillates at a beat frequency dependent on Ω. The envelope oscillations are formed from interactions of the sidebands with the carrier light and with each other. These oscillations carry information of the laser detuning in their phase due to the dispersion of the carrier in the cavity near resonances, with the phase shift responding linearly to the carrier-cavity detuning within the linewidth around the resonance. The power of this signal is sensed by the Reflection PD of Figure 5.15, and the phase of this signal (which is dependent on laser frequency) is extracted by demodulation. During demodulation the photodiode signal is mixed with a reference signal of frequency Ω. The mixing outputs signals at both the sum and the difference of the two frequencies. Since the reference signal frequency was chosen to match the frequency of the signal of interest, Ω, then the output of the 'difference' mixing term produces a DC signal which can be isolated with low-pass filtering. This DC signal is proportional to the phase of the photodiode signal and therefore measures the difference between the carrier frequency and cavity resonance. The signal can thus be used as a feedback for active frequency stabilisation of the laser. The TTFSS was then used to lock the ECDL frequency to the cavity via feedback to the laser current, using the input of the generated error signal.

To maintain cavity lock, a servo was designed to feedback to the laser frequency via the injection current. The laser response to current modulation (Section 5.3.6), the photodiode response, and the cavity response were combined to create the transfer function of the combined plant. The servo was designed using a pole-zero model based on the combined plant response, with sufficient gain for the desired level of noise suppression and a UGF within the phase margin for a stable feedback loop. The closed loop gain transfer function was measured after the servo via a swept sine measurement. The Moku:Lab Frequency Response Analyser was used to allow a measurement up to 1 MHz. The open loop gain transfer function can be calculated from the closed loop gain, and is shown in Figure 5.16. This measurement displays a unity gain frequency of 340 kHz, with a 2 dB gain margin and 25° phase margin. The noise peaks under 10 kHz show a decrease in signal-to-noise ratio at lower frequencies which could arise from both the increase in laser noise at lower frequencies observed in Section 5.3, and from limitations of the drive signal at lower

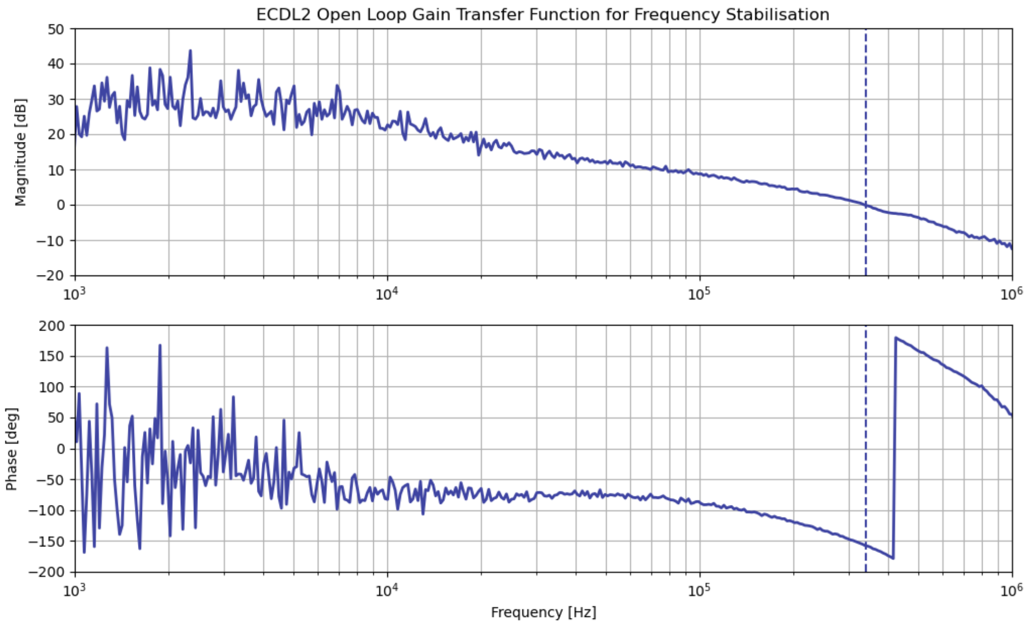

Figure 5.16: Open loop gain transfer function for frequency stabilisation of ECDL2. The unity gain frequency is represented by a dashed vertical line.

frequencies so as to not saturate the electronics. An LT Spice model of the servo transfer function and electronic noise was compared to the shot noise of the photodiode showing that servo electronics noise is only expected to be visible above photodiode shot noise at frequencies over 1 MHz.

Once the ECDL was frequency stabilised via locking to the 10 cm cavity, the stabilised frequency noise could be measured against the fringes of the fMZI. A Fabry-Perot scanning cavity used the cavity reflected light to ensure the laser was running single mode during the measurements. The results of the frequency stabilisation are shown in Figure 5.17. The free running laser frequency noise is represented in orange, measured before the stabilisation. The stabilised frequency noise is shown in dark blue, and the instrument noise in pink. The blue trace represents the free running laser noise divided by the open loop gain function as an estimate of the expected stabilised frequency noise.

Stabilisation of frequency noise has been demonstrated across the entire measurement span of 1 Hz – 100 kHz. Although the UGF of the open loop gain function sits at 340 kHz the final frequency noise measurements are limited to 100 kHz by the bandwidth of the spectrum analyser (the use of the SR785 was necessary due to the lower instrumental noise floor when compared to the Moku:Lab). A maximum suppression of 10,000$\times$ is observed at 50 Hz however this corresponds to the mains peak in the free running frequency noise data and is not a typical level of noise suppression. Ignoring noise peaks, the frequency

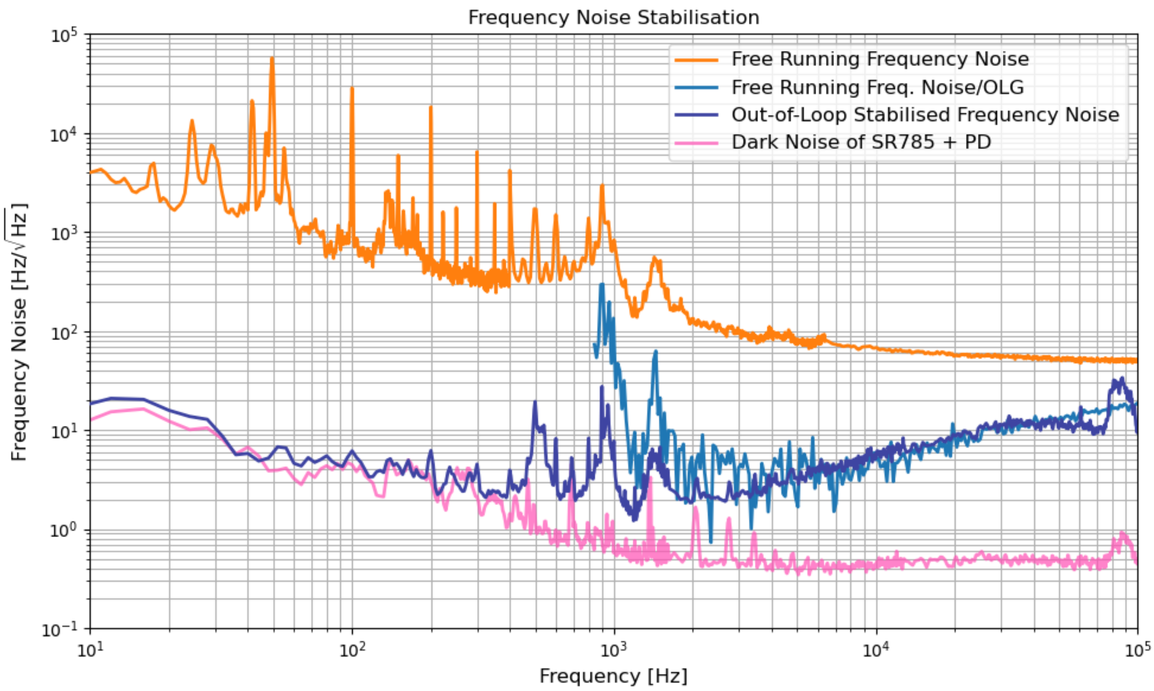

Figure 5.17: Stabilised frequency noise measurements taken at the Australian National University. The free-running laser frequency noise is shown in orange, with the out-of-loop stabilised frequency noise displayed in dark blue. The estimated expected frequency noise is shown in blue as the free running noise divided by the open loop gain. The measurement instrument noise is shown in pink as the combined signal analyser and photodetector dark noise.

noise suppression sits at 200× from 10 Hz – 1 kHz then decreases to 5× suppression by 30 kHz. The stabilised frequency noise performs as expected within the noise limitations of the experiment. At low frequencies (10 Hz – 40 Hz) the stabilised noise is limited by instrument noise of the SR785 spectrum analyser and the dark noise of the PDA10D2 photodetector, shown by the pink line in Figure 5.17. At frequencies above 3 kHz the stabilised frequency noise fits the expected noise limit calculated from the free running noise and feedback transfer function, represented by the blue line in Figure 5.17. The stabilised noise displays three prominent peaks at 0.5 kHz, 1 kHz, and 1.45 kHz. The 1 kHz and 1.45 kHz peak are both present in the free running frequency noise measurement taken on the same day, and in the ECDL2 frequency response to PZT modulation function in Figure 5.11, and are therefore thought to arise from mechanical resonances within the PZT mount. Above 2 kHz the noise rises from $2\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ to $10\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$, and is 5× below the free running frequency noise at the end of the measurement span at 100 kHz. The noise peak observed near 100 kHz is also present in the instrument noise plot, and was observed when using both the SR785 signal analyser and Moku:Lab Frequency Response Analyser, and is thought to be a transient artefact from the fMZI photodiode used to take the frequency response measurements.

The shot noise limit on the error signal sets the limit on the stabilised frequency, calculated in Hz/$\sqrt{\text{Hz}}$ as

$$h_{\text{shot,error}} = \frac{\sqrt{hc^3}}{8} \times \frac{1}{\mathscr{F} L \sqrt{\lambda P_{\text{c}}}} \tag{5.5}$$

for L cavity length and P_{c} the power in the carrier in the reflected signal, measured as 3 mW [30]. The resulting shot noise has a level of 1.96×10^{-4} Hz/$\sqrt{\text{Hz}}$ and therefore lies well below the frequency noise measurements of Figure 5.17.

This work has demonstrated a reduction in ECDL frequency noise via feedback to the diode injection current across a bandwidth of 10 Hz – 100 kHz. The noise suppression fits within predicted limits from instrument noise and the open loop gain, and demonstrates proof-of-concept for future 2 µm development work.

5.5 Conclusion

A background is presented on the need for 1.55µm and 2 µm laser wavelengths in 3G interferometry. The current status of 2 µm research is presented, identifying a gap in technology for a low-noise and inexpensive 2 µm source. External cavity diode lasers were chosen for investigation in this chapter across two research projects at the University of Glasgow and the Australian National University. Diode lasers and external cavity diode lasers in a Littrow configuration were discussed, and the ECDL design used throughout this work was then presented in detail.

Three ECDL units were studied over the course of this project. One was built at the University of Glasgow with the aim of becoming a low-noise seed laser for the Glasgow Cryogenic Interferometry Facility. Two further ECDL units were developed at the Australian National University which will be utilised in a folded optical cavity to take direct measurements of coating thermal noise at cryogenic temperatures. The laser units were characterised with measurements of output power, intensity noise, and frequency noise. Intensity stabilisation was then demonstrated at the University of Glasgow, while frequency stabilisation was demonstrated at the Australian National University.

Measurements of laser output power as a function of input current show that the outputs agree well with the diode chip predictions, with a threshold current of 96.5 ± 4.1 mA and a slope efficiency of 16.7 ± 0.4 mW/A measured for the UofG laser. Despite this, variation between the units is evident, with power output varying both with individual gain chip properties, and on the accuracy of the external cavity alignment. This fact is also evident in the output power against wavelength measurements taken for the two Australian National University lasers. A coarse wavelength tuning range of 103 ± 5 nm around a centre wavelength of approximately 1955 nm was found for ECDL1, while ECDL2 had a FWHM of 157 ± 5 nm with a centre wavelength of 1925 nm. The power variation of ECDL1 with grating angle suggests there may be a slight pitch misalignment of the diffraction grating.

Relative intensity noise was measured for the three units, with all units not limited by either photodetector dark noise or shot noise. The profiles measured for intensity noise with increasing frequency agree well between the ECDLs, with noise decreasing between 10 Hz – 2 kHz before flattening out up to 100 kHz. Resonance peaks near 1 kHz are observed in all three units, arising from a known mechanical resonance in the laser mount [115]. The lowest intensity noise measurement was measured for the University of Glasgow ECDL as $3 \times 10^{-8}\,1/\sqrt{\mathrm{Hz}}$ at 32 kHz.

Laser frequency noise was measured for the ANU lasers against the fringes of a fibre Mach-Zehnder interferometer. ECDL1 was found to have a lower overall noise than ECDL2 over the majority of the measurement range. Both lasers have frequency noise decreasing as $1/\sqrt{\mathrm{f}}$ from 4 Hz – 1 kHz and flatter from 1 kHz – 100 kHz. Above 40 kHz the noise levels of the two lasers tend to a similar value where ECDL1 has a constant noise of $45.6\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ from 5 kHz – 100 kHz, and the frequency noise of ECDL2 slightly decreases with increasing frequency from $80.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 5 kHz to $40.2\ \mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 kHz.

With the aim of creating noise suppression loops to each unit, the laser response to modulation was characterised. The amplitude and frequency response to current modulation was measured to be flat for a bandwidth of 10 Hz to over 10 MHz with a level of 1×10^{-4} mW/V for the amplitude response and 1×10^{2} MHz/V for the frequency response (varying between laser units). The PZT modulation is flat from 10 Hz up to the mechanical resonance peak at 1 kHz, with the level varying between the ANU ECDL units. ECDL2, used for the PDH locking during frequency stabilisation, displayed an amplitude response of 5 mW/V and a frequency response of 1×10^{7} MHz/V across the flat response region.

The first intensity stabilisation measurements were presented for a 2 µm ECDL, which used feedback to the injection current for laser intensity modulation. Intensity noise reduction was demonstrated across a span of 10 Hz – 50 kHz with a maximum noise reduction by a factor of 220 at the 900 Hz resonance peak. The intensity noise suppression is limited by photodiode dark noise between 5 Hz – 30 Hz, and above 30 Hz is limited by excess noise which has a flat response with frequency up to the UGF of the servo gain with a level of $2 \times 10^{-3}\,1/\sqrt{\mathrm{Hz}}$.

The first frequency stabilisation measurements were then presented for a 2 µm ECDL, which used feedback to the injection current to modulate laser frequency. Measurements of out-of-loop stabilised frequency noise were presented, demonstrating a reduction in frequency noise between 10 Hz – 100 kHz. Aside from noise peaks the frequency noise suppression sits at 200× from 10 Hz – 1 kHz then decreases to 5× suppression by 30 kHz, with the noise reduction as expected within the experimental noise limits.

This work demonstrates the full characterisation of three low-noise tunable external cavity diode lasers operating at 2 µm and presents the first results of intensity and frequency stabilisation. Future work within this field will involve concurrent intensity and frequency stabilisation. Implementation of this is possible via feedback to both injection current and grating angle alongside an EOM. The optical set-up of the Australian National University ECDLs was completed and in future will be integrated with a cryogenic folded optical cavity to take a direct measurement of coating thermal noise. A reduction in frequency noise was achieved via PDH locking to a 10 cm cavity. The University of Glasgow ECDL has demonstrated intensity noise reduction across the experimental sensitivity range of 100 Hz – 1 kHz and an adequate power output for the installation of a Tm fibre amplifier to increase power output to the 200 mW range. This would allow the 2 µm ECDL to be used as a seed beam for the GCIF interferometer.

Chapter 6

CONCLUSION

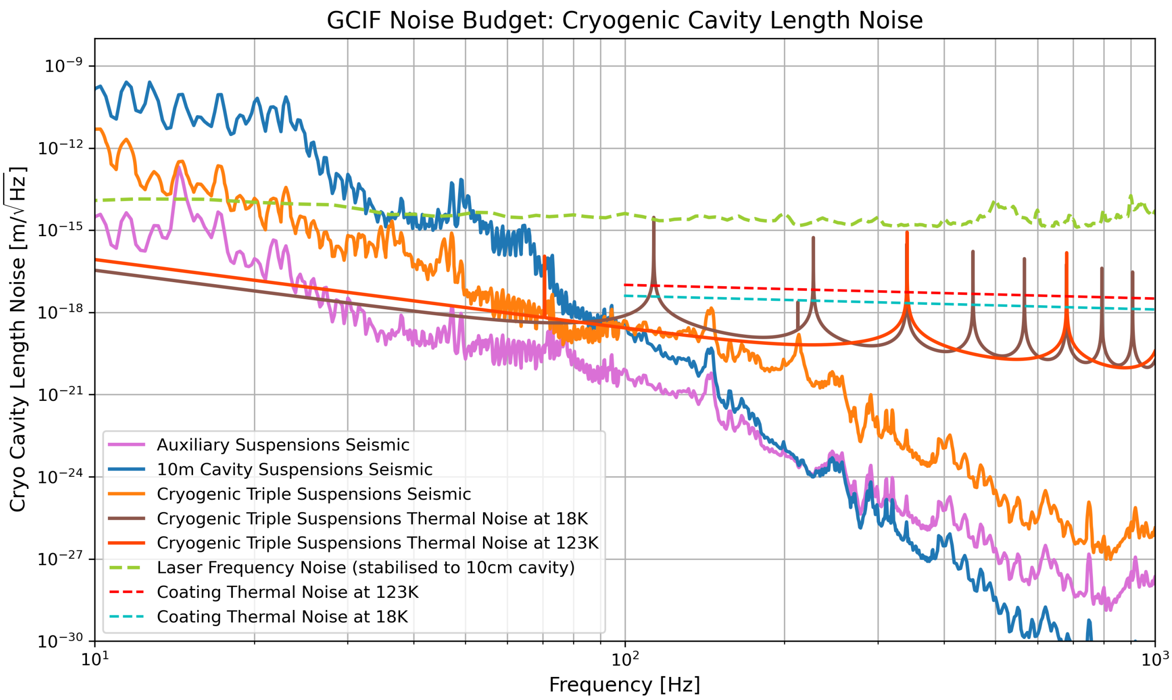

Figure 6.1: Cryogenic cavity length noise contributions from the 300g and auxiliary suspensions, the cryogenic triple-stage suspension, and 2 µm laser frequency noise, forming the noise budget for the GCIF from the sources considered in this work.

In 2015 the advanced LIGO interferometers made the first detection of a gravitational wave. This event opened a new field of astronomy, with discoveries having implications in scientific fields such as dense matter astrophysics, cosmology, nuclear physics, gravitation, and high energy physics – ultimately allowing a better understanding of our Universe. Since 2015 over 200 gravitational wave events have been observed using interferometric gravitational wave detectors. The design of these detectors uses laser light resonating in optical cavities to sense phase shifts in the beam caused by passing gravitational waves. To sense such a small distortion, these instruments must be extraordinarily precise with an extremely high sensitivity. This is achieved through the understanding and mitigation of detector noise sources.

The reduction in seismic noise is achieved though the hanging of suspended optics, isolating the cavity masses from ground motion. Quantum noise can be tuned through a choice of cavity laser power, with further reductions possible through the recent inclusion of squeezed light. Thermal noise is minimised through the design of the suspended optics by the use of low loss materials, and geometries which minimise the energy dissipation in high-loss regions. The majority of the current generation of interferometer detectors operate at room temperature. Several detectors (such as KAGRA, and planned third generation detectors such as ET) aim to further reduce thermal noise by operating at cryogenic temperatures. This change in temperature has a large impact on detector design, the most significant changes being to the mirror substrate material and laser wavelength. Several interferometer prototypes exist which are working on the development of necessary technologies for the next generation of detectors. One such interferometer prototype is the Glasgow Cryogenic Interferometer Facility (GCIF) which aims to study the noise performance of silicon optics within a crystalline suspension, take direct measurements of cryogenic coating thermal noise, and study the viability of 1550 and 2000 nm laser technology for use in future detectors, to name a few. A direct measurement of coating thermal noise can be made provided this is the limiting noise source of the experiment, thus giving a sensitivity requirement for the GCIF. This provides a baseline requirement for the seismic attenuation and thermal noise of optical suspensions, as well as for stabilised laser frequency noise. The work in this thesis is centred around the groundwork investigations into suspension design and laser optics which is required for the development of the GCIF.

The modelling, characterisation, and modification of two double-stage mirror suspensions is described in Chapter 3. These suspensions will be used for beam steering and as cavity optics in the room-temperature section of the GCIF and therefore must be carefully designed to contribute a low amount of length noise into the cryogenic test cavity. A detailed dynamical model of each suspension was developed which output displacement transfer functions and suspension resonances, and was used to predict the optic length noise resulting from seismic ground motion. This output was compared to the experimental sensitivity requirements and informed the design of suspension modifications which aimed to meet the noise requirements in the bandwidth of interest. The seismic-to-length-noise coupling for both suspensions was found to be dominated by the vertical mode contribution, therefore the designs for both suspensions centred around the a reduction in frequency of the vertical resonances, implemented using steel blade springs for the 300 g suspension, and rubber for the auxiliary suspension. The 300 g double suspension became the focus of this work based on the greater need for performance improvements. Experimental procedures were developed for the characterisation of the 300 g suspension for both force-displacement transfer functions and the vertical-vertical transfer function. These measurements were performed on both the unmodified and modified suspension versions to give a full characterisation of the builds. The desired reduction in resonant frequencies was demonstrated, and methods for the reduction of new suspension frame and blade spring resonances were investigated. Modelling of the modified suspension optic length noise showed both modified suspensions are expected to meet the experimental design requirement across the desired sensitivity range.

The work described in Chapter 4 produced a design concept for a triple-stage optical suspension with a fully crystalline final stage and silicon masses, suitable for use in a cryogenic cavity for direct measurements of coating thermal noise. The design of this suspension considered seismic isolation requirements, suspension thermal noise requirements, and the design of a cooling scheme at both 123 K and 18 K operation phases of the GCIF. This design concept brings together research at the cutting edge of the field and is an integral demonstration for larger cryogenic, crystalline suspensions such as required for ET. A triple-stage suspension design was chosen to meet seismic isolation requirements, which considered multiple stages of blade springs and a compact nested pendulum platform for further seismic isolation. A double-stage blade spring design with long upper blades was found to give the best balance of seismic isolation and reliability while minimising build complexity. Cooling schemes for each GCIF operation phase were modelled. These considered the cryostat heat loads and dominant heat transfer mechanisms of each phase. In the radiative regime at 123 K a heat shield around the monolithic stage is recommended to radiatively maintain a test mass temperature of 123 K, resulting in a negligible thermal gradient across the crystalline fibres to the intermediate mass. To operate a cavity temperature of 18 K with a seismically quiet, low thermal noise test mass, heat must be conducted away from the mass via the crystalline fibres of the monolithic stage. Silicon and sapphire fibres were both considered. The modelled fibre diameters were increased to facilitate the necessary conductive cooling, with bonding and phonon-boundary scattering effects accounted for in the model. A heat link system was designed, with required cooling of the chains up to the IM possible with either copper or aluminium heat links. Finally, the single test mass displacement thermal noise was modelled for six physical scenarios and four un-physical scenarios for the GCIF suspension. This work demonstrates the complexity of suspension thermal noise calculations, showing the need for the careful consideration of the geometries and jointing mechanisms of the GCIF suspension design. The domination of joint loss on suspension displacement thermal noise was clearly demonstrated. This work shows the need for future research and development of silicon and sapphire fibre production, shaping, welding, and bonding techniques, as all these factors affect the resulting thermal noise of the GCIF cryogenic suspension.

Changing of the optic substrate for cryogenic GW detectors necessitates also a change in the wavelength of the interferometer laser: silicon is opaque to the 1064 nm wavelength in use at current detectors and a move to longer wavelengths in the 1.55 – 2 µm range will be made. 1.55 µm and 2 µm wavelengths are both in consideration for 3G detectors, with the pros and cons to each option outlined at the beginning of Chapter 5. The GCIF plans to run as a dual-wavelength facility, demonstrating the technologies of both wavelengths. The work developed here concentrated on the development of 2 µm ECDL technology as a low noise, inexpensive, and tunable laser source with potential use both as a compact and customisable tool for research and development, or as a seed laser for 3G detectors. Three 2 µm units were developed, with investigation done as a collaborative effort between the University of Glasgow and the Australian National University. Measurements were taken of ECDL output power, intensity noise, frequency noise, and the laser response to modulation. The first intensity stabilisation measurements were presented for a 2 µm ECDL at the University of Glasgow, via feedback to the injection current. This work demonstrated a maximum noise reduction by a factor of 220 at the 1 kHz resonance peak, and an average reduction by a factor of 15 between 10 Hz – 1 kHz. Noise suppression was demonstrated across a frequency span of 10 Hz – 50 kHz. The first frequency stabilisation measurements were then presented for a 2 µm ECDL, demonstrated at the Australian National University. This was performed via feedback to the injection current and demonstrated a reduction in frequency noise across 10 Hz – 100 kHz, with a suppression factor of 200 between 10 Hz – 1 kHz. Future work will involve simultaneous intensity and frequency stabilisation via feedback to both grating angle and injection current, and the installation of a fibre amplifier at the University of Glasgow to allow the 2µm ECDL to be used as a seed beam for the GCIF interferometer.

Figure 6.1 brings together the cryogenic cavity length noise contributions from the sources investigated within this thesis. The seismic contribution to cavity length noise is considered from the modified 10 m cavity 300 g suspensions (blue) and the modified beam steering auxiliary suspension (pink) from Chapter 3. These results were previously shown in Figure 3.25 and have been converted here to cryogenic cavity length noise. The orange line represents the seismic contribution to cavity length noise for the cryogenic triple suspension, with the 650 µm diameter silicon fibre chosen as the limiting case. This assumes the '2 stage LUB' model with no additional seismic isolation platform. The red and brown solid lines represent the suspension thermal noise contribution to cavity length noise from the silicon fibre cases at 123 K and 18 K respectively. The case shown for 18 K is the 650 µm diameter silicon fibre with a HCB-like join, while the red line of the 123 K case is the 216 µm diameter silicon fibre with a HCB-like join. These results were previously shown in Figure 4.14 and have been converted to cryogenic cavity length noise.

The stabilised laser frequency noise shown in Figure 5.17 was intended to demonstrate proof-of-concept of frequency noise stabilisation via feedback to diode injection current, and was not designed at this stage to meet the noise requirements of the GCIF. The dashed green line in Figure 6.1 shows this laser frequency noise measurement, stabilised to the table top 10 cm cavity and then converted into cryogenic cavity length noise. The GCIF will utilise an aLIGO style nested loop stabilisation system, first with a pre-stabilisation stage using a fixed 10 cm table top cavity via feedback to the laser injection current (as demonstrated in Chapter 5). A second servo stage will lock the laser to the fixed 10 m reference cavity via feedback to the first stabilisation loop. A noise suppression of 200 was demonstrated in this work across the 100 Hz – 1 kHz range, by locking the laser to the 10 cm cavity. Comparing this pre-stabilised frequency noise to the cavity length noise requirements, a further suppression of 800 is required from the 10 m arm cavity to meet the sensitivity requirement of the experiment. To be comfortably below the experimental noise requirement (by a safety factor of 2.5) a suppression of 2000 is recommended for the second stabilisation loop, assuming that 200$\times$ suppression is already provided by the 10 cm cavity pre-stabilisation. With the experimental sensitivity bandwidth, this only necessitates a UGF of near 1 – 5 kHz, feasible with the available technology. In addition, the gain of the 10 cm cavity stage demonstrated the necessary technology, with gain ultimately limited by the sensing of the out-of-loop frequency noise measurement. The gain factor of this pre-stabilisation stage is not near the shot noise limit and can therefore be further increased, which would reduce the gain requirements on the 10 m servo stage.

The results presented in this work demonstrate a contribution to the research required for the next generation of gravitational wave detectors, producing a concept design for a cryogenic triple-stage suspension with silicon optics, and the advancement of 2µm laser technology. The importance of accurate dynamical and thermal noise modelling is highlighted, alongside the careful considerations which need to be made in the design of a low noise crystalline suspension which is able to maintain steady operation at cryogenic temperatures. This work also demonstrates the first intensity and frequency stabilisation of 2µm ECDLs for gravitational wave applications, as part of the development of low noise, tunable external cavity diode lasers at 2µm. This work supports the build of the GCIF as a research facility which is crucial for the advancement of technologies needed to build the next generation of gravitational wave interferometers.

Appendices

A Cryogenic Triple Suspension Thermal Noise Plots

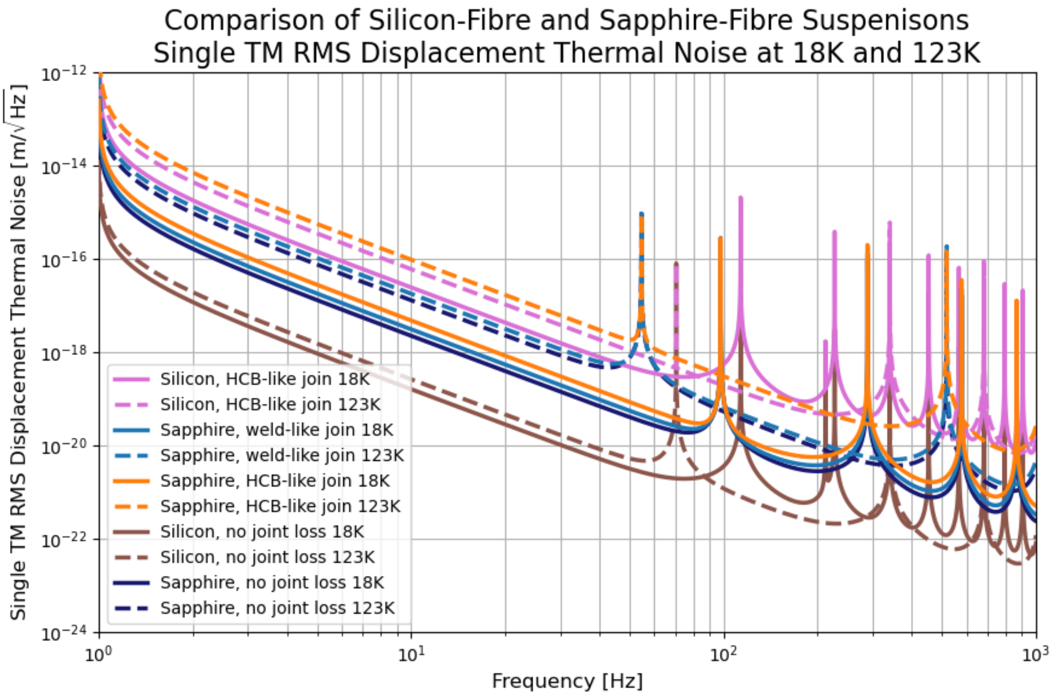

Figure 6: Single test mass RMS displacement thermal noise for all 10 suspension cases considered, including the unphysical no-joint cases.

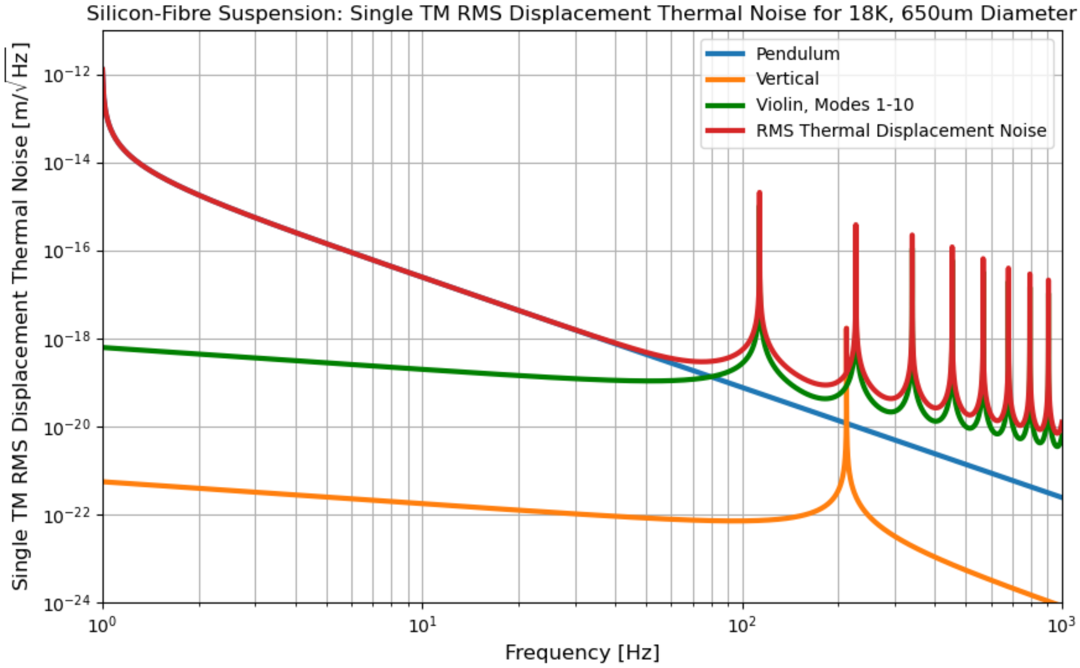

(a) Silicon fibre suspension at 18 K, 650 μm diameter.

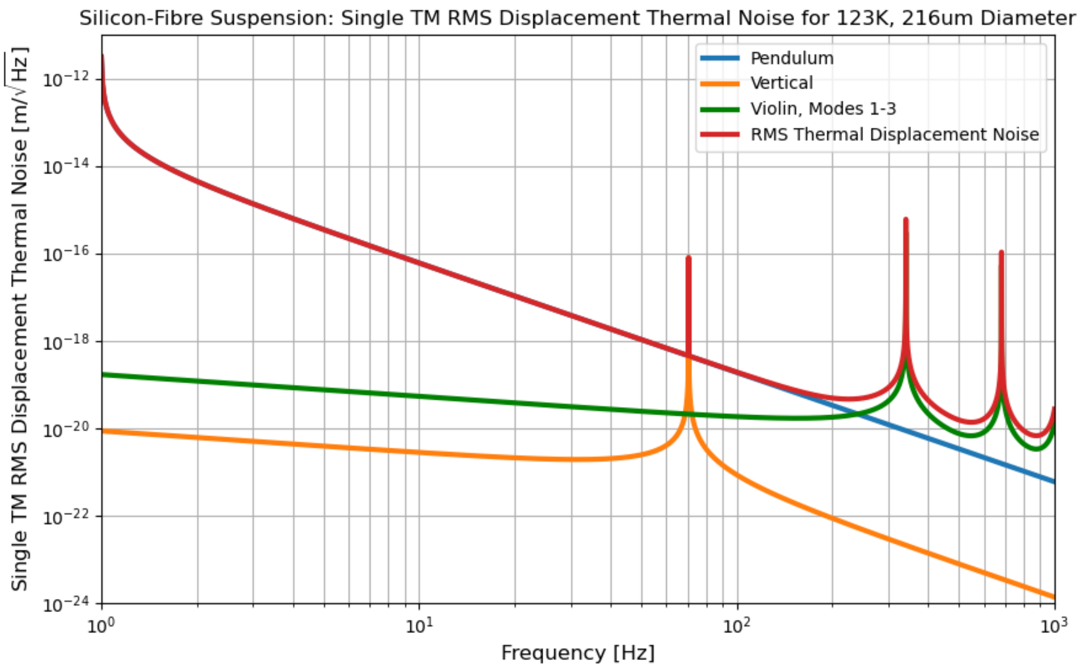

(b) Silicon fibre suspension at 123 K, 216 μm diameter.

Figure 7: Single test mass RMS displacement thermal noise for a silicon fibre suspension with HCB-like jointing, shown for 18 K and 123 K operation cases.

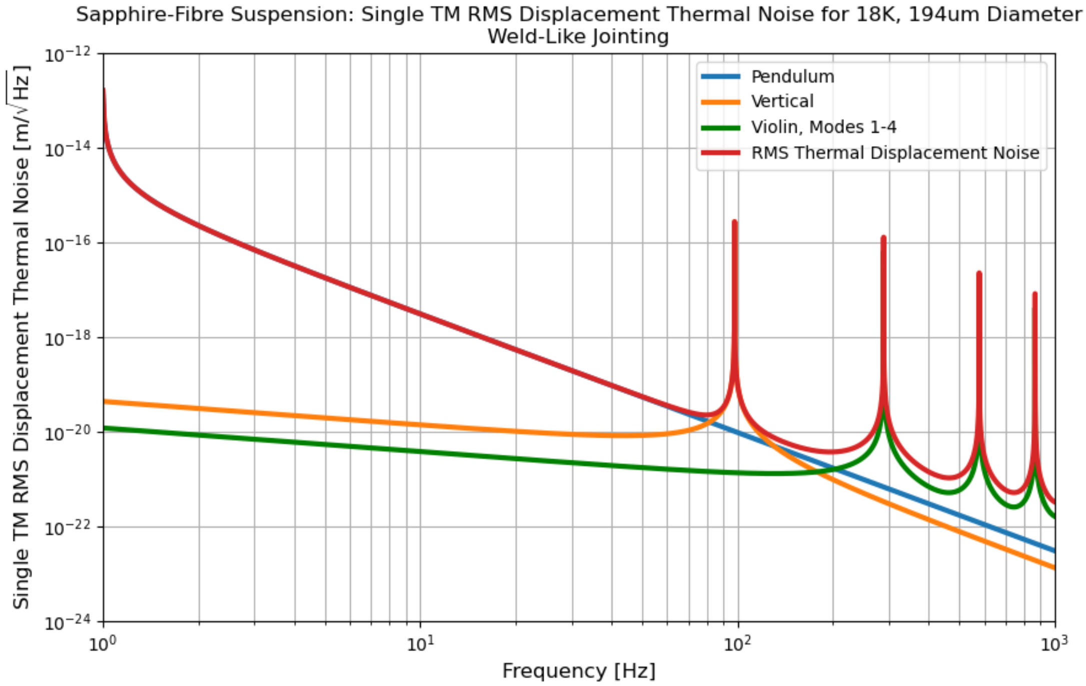

(a) Sapphire fibre suspension at 18 K, 194 µm diameter, with weld-like jointing.

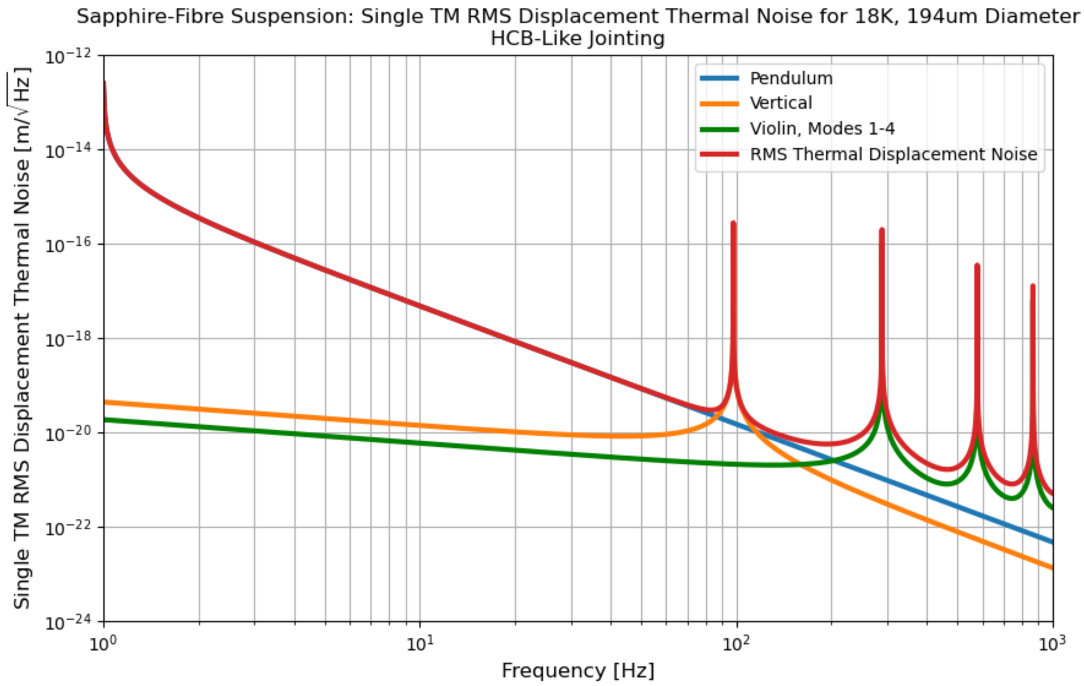

(b) Sapphire fibre suspension at 18 K, 194 µm diameter, with HCB-like jointing.

Figure 8: Single test mass RMS displacement thermal noise for a sapphire fibre suspension, shown for 18 K with weld-like joint and HCB-like joint cases.

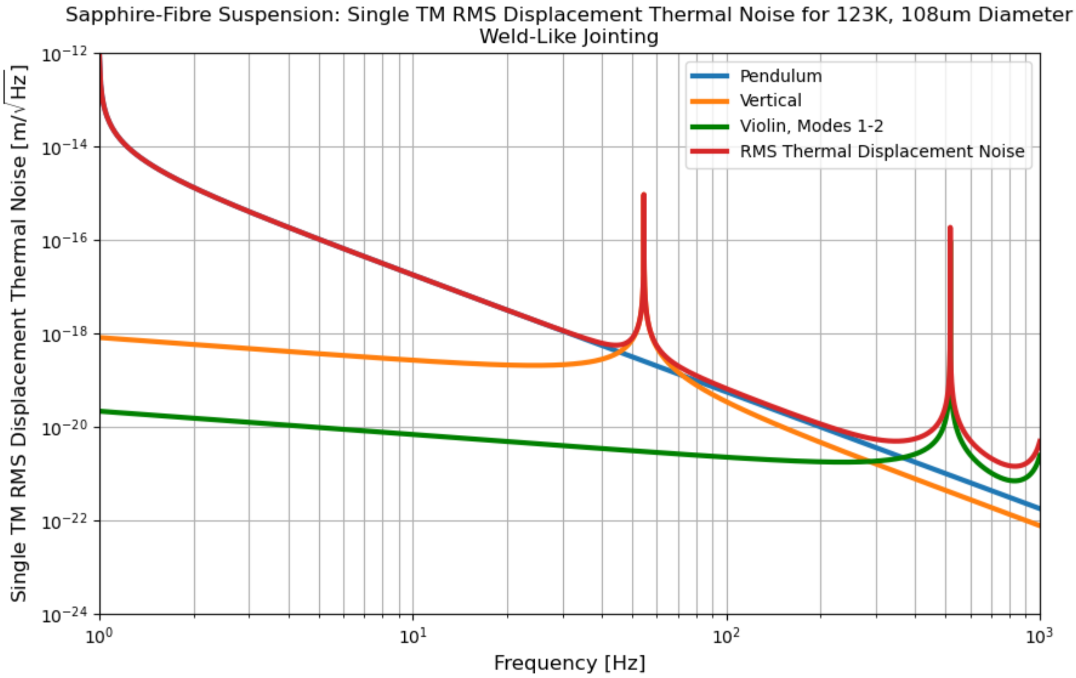

(a) Sapphire fibre suspension at 123 K, 108 µm diameter, with weld-like jointing.

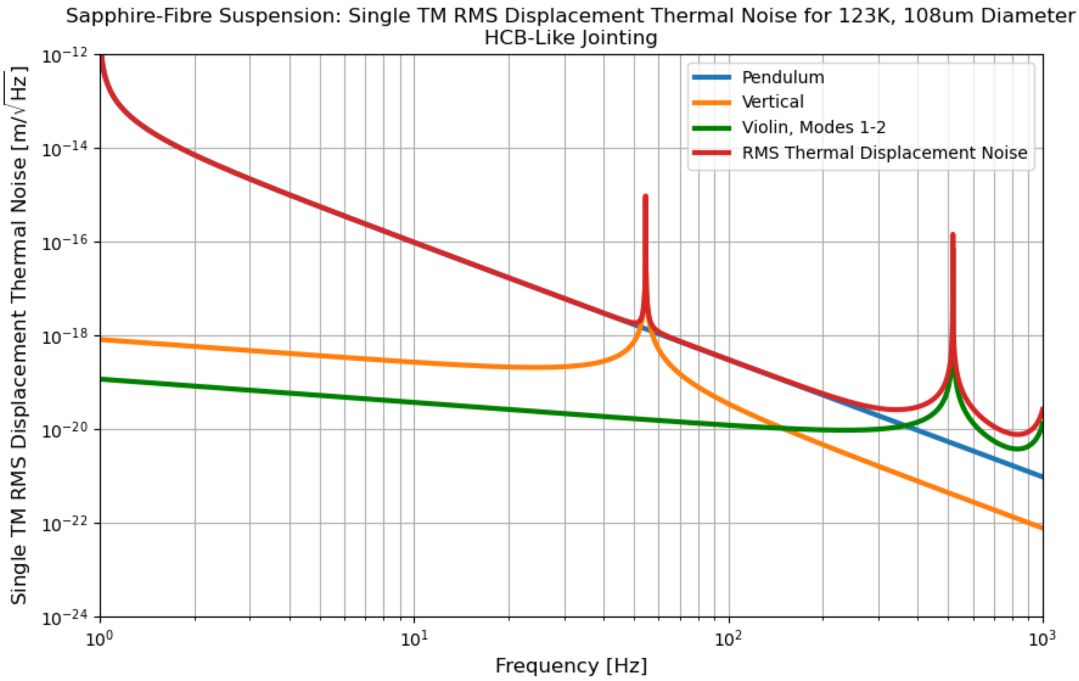

(b) Sapphire fibre suspension at 123 K, 108 µm diameter, with HCB-like jointing.

Figure 9: Single test mass RMS displacement thermal noise for a sapphire fibre suspension, shown for 123 K with weld-like joint and HCB-like joint cases.

Bibliography

[1] H. J. Maris A. K. McCurdy and C. Elbaum. 'Anisotropic heat conduction in cubic crystals in the boundary scattering regime'. In: *Phys. Rev. B* 2 (1970). DOI: https://doi.org/10.1103/PhysRevB.2.4077.

[2] B. P. Abbott et al. 'Observation of Gravitational Waves from a Binary Black Hole Merger'. In: *Phys. Rev. Lett.* 116 (2016). DOI: http://doi.org/10.1103/PhysRevLett.116.061102.

[3] A. Abramovici et al. 'LIGO: The Laser Interferometer Gravitational-Wave Observatory'. In: *Science* 256 (1992). DOI: https://doi.org/10.1126/science.256.5055.325.

[4] T. Accadia et al. 'The Seismic Superattenuators of the Virgo Gravitational Waves Interferometer'. In: *J. Low Freq. Noise Vib. Act. Control.* 30 (2011). DOI: http://doi.org/10.1260/0263-0923.30.1.63.

[5] F. Acernese et al. 'The Advanced Virgo detector'. In: *J. Phys.: Conf. Ser* 610 (2015). DOI: https://doi.org/10.1088/1742-6596/610/1/012014.

[6] J. Achenbach. *LIGO's success was built on many failures.* URL: https://www.washingtonpost.com/news/achenblog/wp/2016/02/12/the-many-many-many-unsung-heroes-of-ligo/. (accessed: 29.8.2024).

[7] R. Adhikari. *'Integrated detector commissioning' in Advanced Interferometric Gravitational Wave Detectors. Volume I: Essentials of Gravitational-Wave Detectors.* World Scientific Publishing Co. Pte. Ltd, 2019. ISBN: 9789813146082.

[8] R. Adhikari et al. 'A cryogenic silicon interferometer for gravitational wave detection'. In: *Class. Quantum Grav.* 37 (2020). DOI: https://doi.org/10.1088/1361-6382/ab9143.

[9] R. Adhikari et al. 'Length sensing and control of the Caltech 40m prototype'. In: *LIGO Document LIGO-G1300950* (2013).

[10] R. Adhikari et al. 'Mariner: The Cryogenic Upgrade of the 40m Prototype Interferometer'. In: *LIGO Document LIGO-G2301014* (2023).

[11] European Space Agency. 'LISA Definition Study Report'. In: *ESA Document ESA-SCI-DIR-RP-002* (2024).

[12] O. D. Aguiar. 'The Past, Present and Future of the Resonant-Mass Gravitational Wave Detectors'. In: *Research in Astron. Astrophys.* 11 (2010). DOI: `https://doi.org/10.1088/1674-4527/11/1/001`.

[13] D. Aisa et al. 'The Advanced Virgo monolithic fused silica suspension'. In: *Nuclear Instruments and Methods in Physics Research A* 824 (2016). DOI: `https://doi.org/10.1016/j.nima.2015.09.037`.

[14] T. Akutsu et al. 'First cryogenic test operation of underground km-scale gravitational-wave observatory KAGRA'. In: *Class. Quantum Grav.* 36 (2019). DOI: `https://doi.org/10.1088/1361-6382/ab28a9`.

[15] T. Akutsu et al. 'Overview of KAGRA: Detector design and construction history'. In: *Progress of Theoretical and Experimental Physics* 2021 (2021). DOI: `https://doi.org/10.1093/ptep/ptaa125`.

[16] W. G. Anderson et al. 'Excess power statistic for detection of burst sources of gravitational radiation'. In: *Phys. Rev. D* 63 (2001). DOI: `https://doi.org/10.1103/PhysRevD.63.042003`.

[17] S. M. Aston et al. 'Update on quadruple suspension design for Advanced LIGO'. In: *Class. Quantum Grav.* 29 (2012). DOI: `https://doi.org/110.1088/0264-9381/29/23/235004`.

[18] P. Astone et al. 'The gravitational wave detector NAUTILUS operating at T = 0.1 K'. In: *Astroparticle Physics* 11 (2010). DOI: `https://doi.org/10.1016/S0927-6505(97)00023-6`.

[19] B. W. Barr. 'Experimental Investigations into Advanced Configurations and Optical Techniques for Laser Interferometric Gravitational Wave Detectors'. PhD thesis. University of Glasgow, 2003.

[20] B. W. Barr et al. 'Optical modulation techniques for length sensing and control of optical cavities'. In: *Appl. Opt.* 31 (2007). DOI: `https://doi.org/10.1364/ao.46.007739.`.

[21] B. W. Barr et al. 'Translational, rotational, and vibrational coupling into phase in diffractively coupled optical cavities'. In: *Optics Letters* 36 (2011). DOI: `https://doi.org/10.1364/OL.36.002746`.

[22] L. Barsotti et al. 'The updated Advanced LIGO design curve'. In: *LIGO Document LIGO-T1800044-v3* (2018).

[23] M. Barton. *DualLite2NMB Mathematica Model.* URL: `https://svn.ligo.caltech.edu/svn/sus/trunk/Common/MathematicaModels/DualLite2DB/ASUS2L2DBModelDefn.nb`.

[24] M. Barton. 'Models of the Advanced LIGO Suspensions in Mathematica'. In: *LIGO Document LIGO- T020205-v2* (2014).

[25] M. Barton and C. Torrie. 'Suspension model comparisons'. In: *LIGO Document LIGO-T020011-00-D* (2003).

[26] M. Bassan. *Advanced Interferometers and the Search for Gravitational Waves*. Springer International Publishing, 2014. ISBN: 978-3-319-03791-2.

[27] S. Bennetts et al. 'External cavity diode lasers with 5kHz linewidth and 200nm tuning range at 1.55um and methods for linewidth measurement'. In: *Optics Express* 22 (2014). DOI: `https://doi.org/10.1364/OE.22.010642`.

[28] N. Beveridge. 'Characterisation of silicon-silicon hydroxide catalysis bonds for future gravitational wave detectors'. PhD thesis. University of Glasgow, 2012.

[29] R. Birney et al. 'Coatings and surface treatments for enhanced performance suspensions for future gravitational wave detectors'. In: *Class. Quantum Grav.* 34 (2017). DOI: `https://doi.org/10.1088/1361-6382/aa9354`.

[30] E. D. Black. 'An introduction to Pound–Drever-Hall laser frequency stabilization'. In: *Am. J. Phys.* 69 (2001). DOI: `https://doi.org/10.1119/1.1286663`.

[31] R. Brown. 'A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies'. In: *The Philosophical Magazine and Annals of Philosophy* 4 (1828). DOI: `https://doi.org/10.1080/14786442808674769`.

[32] I. Buchovska. Private Communication.

[33] A. Buikema et al. 'Sensitivity and performance of the Advanced LIGO detectors in the third observing run'. In: *Phys. Rev. D* 102 (2020). DOI: `https://doi.org/10.1103/PhysRevD.102.062003`.

[34] J. Callaghan. 'Development of Crystalline Suspension Fibres for Next Generation Gravitational Wave Detectors'. PhD thesis. University of Glasgow, 2024.

[35] H. B. Callen and R. F. Greene. 'On a Theorem of Irreversible Thermodynamics'. In: *Phys. Rev* 86 (1952). DOI: `https://doi.org/10.1103/PhysRev.86.702`.

[36] H. B. Callen and T. A. Welton. 'Irreversibility and Generalized Noise'. In: *Phys. Rev.* 83 (1951). DOI: `https://doi.org/10.1103/PhysRev.83.34`.

[37] B. Caron et al. 'The VIRGO interferometer for gravitational wave detection'. In: *Nuclear Physics B - Proceedings Supplements* 54 (1997). DOI: `https://doi.org/10.1016/S0920-5632(97)00109-6`.

[38] D. Chen et al. 'Vibration measurement in the KAGRA cryostat'. In: *Class. Quantum Grav.* 31 (2014). DOI: `https://doi.org/10.1088/0264-9381/31/22/224001`.

[39] N. Christensen. 'Stochastic gravitational wave backgrounds'. In: *Reports on Progress in Physics* 82 (2019).

[40] ET Collaboration. *Einstein Telescope Image Gallery*. URL: `https://www.et-gw.eu/index.php/etimages`. (accessed: 14.10.2024).

[41] ET Collaboration. 'Workshop Minutes'. In: *ET-LF Core Optics/Suspensions Workshop* (2018).

[42] KAGRA Collaboration. 'KAGRA: 2.5 generation interferometric gravitational wave detector'. In: *Nature Astronomy* 3 (2019). DOI: `https://doi.org/10.1038/s41550-018-0658-y`.

[43] LIGO Scientific Collaboration and Virgo Collaboration. 'A gravitational-wave measurement of the Hubble constant following the second observing run of Advanced LIGO and Virgo'. In: *The Astrophysics Journal* 909 (2021). DOI: `https://doi.org/10.3847/1538-4357/abdcb7`.

[44] LIGO Scientific Collaboration, Virgo Collaboration et al. 'A gravitational-wave standard siren measurement of the Hubble constant'. In: *Nature* 551 (2017). DOI: `https://doi.org/10.1038/nature24471`.

[45] LIGO Scientific Collaboration and Virgo Collaboration. 'GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral'. In: *Phys. Rev. Lett.* 119 (2017). DOI: `https://doi.org/10.1103/PhysRevLett.119.161101`.

[46] LIGO Scientific Collaboration and Virgo Collaboration. 'Tests of General Relativity with GW170817'. In: *Phys. Rev. Lett.* 123 (2019). DOI: `https://doi.org/10.1103/PhysRevLett.123.011102`.

[47] The LIGO Scientific Collaboration. 'Advanced LIGO'. In: *Class. Quantum Grav.* 32 (2015). DOI: `http://doi.org/10.1088/0264-9381/32/7/074001`.

[48] The LIGO Scientific Collaboration. 'Instrument Science White Paper 2022-2023'. In: *LIGO Document LIGO-T2200384v2* (2022).

[49] The LIGO Scientific Collaboration. *LIGO Surpasses the Quantum Limit*. URL: `https://www.ligo.caltech.edu/news/ligo20231023`. (accessed: 29.08.2024).

[50] The LIGO Scientific Collaboration. *LIGO's Interferometer*. URL: `https://www.ligo.caltech.edu/page/ligos-ifo`. (accessed: 27.10.2024).

[51] The LIGO Scientific Collaboration. *Stochastic Gravitational Waves*. URL: `https://www.ligo.org/science/GW-Stochastic.php`. (accessed: 27.10.2024).

[52] The LIGO Scientific Collaboration and Virgo Collaboration. 'GW150914: The Advanced LIGO Detectors in the Era of First Discoveries'. In: *Phys. Rev. Lett.* 116 (2016). DOI: `http://doi.org/10.1103/PhysRevLett.116.131103`.

[53] The VIRGO Collaboration. 'Virgo nEXT: beyond the AdV+ project'. In: *Virgo Document VIR-0497A-22* (2022).

[54] Virgo Collaboration. *The first Fused Silica monolithic suspension for Advanced Virgo*. URL: `http://public.virgo-gw.eu/the-first-fused-silica-monolithic-suspension-for-advanced-virgo/`. (accessed: 21.8.24).

[55] Inc COMSOL. *COMSOL Material Library*. URL: `https://www.comsol.com/material-library`. (accessed: 16.05.2024).

[56] Inc COMSOL. *Heat Transfer Module: Analyze Thermal Effects with Advanced Simulation Software*. URL: `https://www.comsol.com/heat-transfer-module`. (accessed: 25.06.2024).

[57] A. V. Cumming. Private Communication.

[58] A. V. Cumming. 'Aspects of mirrors and suspensions for advanced gravitational wave detectors'. PhD thesis. University of Glasgow, 2008.

[59] A. V. Cumming et al. 'Design and development of the advanced LIGO monolithic fused silica suspension'. In: *Class. Quantum Grav.* 29 (2012). DOI: `https://doi.org/10.1088/0264-9381/29/3/035003`.

[60] A. V. Cumming et al. 'Finite element modelling of the mechanical loss of silica suspension fibres for advanced gravitational wave detectors'. In: *Class. Quantum Grav.* 26 (2009). DOI: `https://doi.org/10.1088/0264-9381/26/21/215012`.

[61] A. V. Cumming et al. 'Lowest observed surface and weld losses in fused silica fibres for gravitational wave detectors'. In: *Class. Quantum Grav.* 37 (2020). DOI: `https://doi.org/10.1088/1361-6382/abac42`.

[62] A. V. Cumming et al. 'Proposed large scale monolithic fused silica mirror suspension for 3rd generation gravitational wave detectors'. In: *Phys. Rev. Applied* 17 (2022). DOI: `https://doi.org/10.48550/arXiv.2111.09119`.

[63] A. V. Cumming et al. 'Silicon mirror suspensions for gravitational wave detectors'. In: *Class. Quantum Grav.* 31 (2013). DOI: `https://doi.org/10.1088/0264-9381/31/2/025017`.

[64] P Saulson D Reitze and H Grote. *Advanced Interferometric Gravitational-Wave Detectors Volume 1: Essentials of Gravitational-Wave Detectors*. World Scientific Publishing, 2019. ISBN: 978-981-3143-07-4.

[65] A. Dari et al. 'Breaking strength tests on silicon and sapphire bondings for gravitational wave detectors'. In: *Class. Quantum Grav.* 27 (2010). DOI: `https://doi.org/10.1088/0264-9381/27/4/045010`.

[66] V. Datillo and the VIRGO Collaboration. 'The VIRGO suspensions: design and recent performance measurements'. In: *Phys. Lett. A* 318 (2003). DOI: `https://doi.org/10.1016/j.physleta.2003.07.012`.

[67] J. Degallaix et al. 'Bulk optical absorption of high resistivity silicon at 1550 nm'. In: *Optics Letters* 38 (2013). DOI: `http://doi.org/10.1364/OL.38.002047`.

[68] J. Degallaix et al. 'Measurement of the optical absorption of bulk silicon at cryogenic temperature and the implication for the Einstein Telescope'. In: *Class. Quantum Grav.* 31 (2014). DOI: `https://doi.org/10.1088/0264-9381/31/18/185010`.

[69] R. W. P. Drever et al. 'Laser Phase and Frequency Stabilization Using an Optical Resonator'. In: *Appl. Phys. B* 31 (1983). DOI: `https://doi.org/10.1007/BF00702605`.

[70] G. Dvali and S. Tye. 'Brane inflation'. In: *Phys. Lett. B* 450 (1999). DOI: `https://doi.org/10.1016/S0370-2693(99)00132-X`.

[71] S. Dwyer et al. 'Gravitational wave detector with cosmological reach'. In: *Phys. Rev. D.* 91 (2015). DOI: `https://doi.org/10.1103/PhysRevD.91.082001`.

[72] S. Dwyer. 'Quantum noise reduction using squeezed states in LIGO'. PhD thesis. Massachusetts Institute of Technology, 2013.

[73] S. Dwyer and S. W. Ballmer. 'Radiative thermal noise for transmissive optics in gravitational-wave detectors'. In: *Phys. Rev. D* 90 (2014). DOI: `https://doi.org/10.1103/PhysRevD.90.043013`.

[74] G. Eddolls. 'Design, build and characterisation of a prototype single crystalline silicon cryogenic suspension for 3rd generation gravitational wave detectors'. PhD thesis. University of Glasgow, 2022.

[75] A. Einstein. In: *Preussische Akademie der Wissenschaften* (1915).

[76] A. Einstein. In: *Preussische Akademie der Wissenschaften, Sitzungsberichte* (1916), pp. 688–696.

[77] E. J. Elliffe et al. 'Hydroxide-catalysis bonding for stable optical systems for space'. In: *Class. Quantum Grav.* 22 (2005). DOI: `https://doi.org/10.1088/0264-9381/22/10/018`.

[78] M. A. Fazio et al. 'Structure and morphology of low mechanical loss TiO2-doped Ta2O5'. In: *Opt. Mater. Express* 10 (2020). DOI: `https://doi.org/10.1364/OME.395503`.

[79] S. Hild G. Hammond and M. Pitkin. 'Advanced technologies for future ground-based, laser-interferometric gravitational wavedetectors'. In: *Journal of Modern Optics* 61 (2014). DOI: `https://doi.org/10.1080/09500340.2014.920934`.

[80] S. Waldman G. McIvor and P. Willems. 'Analysis of LIGO Test Mass Internal Modes as a Measure of Coating Absorption'. In: *LIGO Document LIGO-G070636-00-0* (2007).

[81] J. Gamez. 'Prototype Mirror Suspension for Cryogenic Interferometers'. In: *LIGO Document LIGO-T2200279v1* (2022).

[82] M. Ganija et al. 'Cryogenically cooled, Ho:YAG, Q-switched laser'. In: *Appl. Phys. B* 126 (2020). DOI: `https://doi.org/10.1007/s00340-020-07420-9`.

[83] J. Giaime et al. 'A passive vibration isolation stack for LIGO: Design, modeling, and testing'. In: *Rev. Sci. Instrum.* 67 (1996). DOI: `https://doi.org/10.1063/1.1146573`.

[84] U. Gibson. 'Crystalline silicon suspension fibre'. In: *Conference Presentation: ET-LF Core Optics/Suspensions Workshop* (2018).

[85] A. Gillespie and F. Raab. 'Thermal noise in the test mass suspensions of a laser interferometer gravitational-wave detector prototype'. In: *Phys. Lett. A* 178 (1993). DOI: `https://doi.org/10.1016/0375-9601(93)90861-S`.

[86] University of Glasgow. *500k bequest makes splash with Glasgow gravitational wave scientists.* URL: `https://www.gla.ac.uk/news/archiveofnews/2017/december/headline_566307_en.html`. (accessed: 29.08.2024).

[87] E. Goetz and K. Riles. 'An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems'. In: *Class. Quantum Grav* 28 (2011).

[88] N. A. Gordon. ‘Characterisation and control of coupled optical springs for future gravitational wave detectors’. PhD thesis. University of Glasgow, 2015.

[89] C. Gräf et al. ‘Design of a speed meter interferometer proof-of-principle experiment’. In: *Class. Quantum Grav.* 31 (2014). DOI: https://doi.org/10.1088/0264-9381/31/21/215009.

[90] S. Gras and M. Evans. ‘Direct measurement of coating thermal noise in optical resonators’. In: *Phys. Rev. D* 98 (2018). DOI: https://doi.org/10.1103/PhysRevD.98.122001.

[91] Max Planck Institute for Graviational Physics. *AEI 10m Prototype*. URL: https://10m.aei.mpg.de/project-overview/. (accessed: 27.08.2024).

[92] Max Planck Institute for Gravitational Physics. *GEO600 Gravitational Wave Detector*. URL: https://www.geo600.org/. (accessed: 28.08.2024).

[93] E. Green. ‘The Story of Q’. In: *American Scientist* 43 (1955). DOI: http://www.jstor.org/stable/27826701.

[94] A. M. Gretarsson and G. M. Harry. ‘Dissipation of mechanical energy in fused silica fibers’. In: *Review of Scientific Instruments* 70 (1999). DOI: https://doi.org/10.1063/1.1150040.

[95] R. Grevers. *Testing Silicon Blade Spring Designs for Application in Future Gravitational Wave Detectors*. University of Glasgow Summer Project Thesis. 2018.

[96] I. Gupta et al. ‘Characterizing Gravitational Wave Detector Networks: From A to Cosmic Explorer’. In: *Published to arXiv* (2024). DOI: https://doi.org/10.48550/arXiv.2307.10421.

[97] E. Gustafson et al. ‘LSC White Paper on Detector Research and Development’. In: *LIGO Document T990080-x0* (1999).

[98] G. Hammond. Private Communication.

[99] G. M. Harry et al. ‘Thermal noise from optical coatings in gravitational wave detectors’. In: *Appl. Opt.* 45 (2006). DOI: https://doi.org/10.1364/AO.45.001569.

[100] G. M. Harry et al. ‘Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings’. In: *Class. Quantum Grav.* 19 (2002). DOI: http://doi.org/10.1088/0264-9381/19/5/305.

[101] K. Haughian et al. ‘Mechanical loss of a hydroxide catalysis bond between sapphire substrates and its effect on the sensitivity of future gravitational wave detectors’. In: *Phys. Rev. D* 94 (2016). DOI: http://doi.org/10.1103/PhysRevD.94.082003.

[102] D. Heinert et al. ‘High-sensitivity tool for studying phonon related mechanical losses in low loss materials’. In: *Journal of Physics: Conf. Series* 92 (2007). DOI: https://doi.org/10.1088/1742-6596/92/1/012183.

[103] J. S. Hennig. ‘Mirror suspensions for the Glasgow Sagnac speed meter’. PhD thesis. University of Glasgow, 2018.

[104] M. C. Masso Herrera. 'Properties of bonded silicon for future generations of gravitational wave observatories.' PhD thesis. University of Glasgow, 2019.

[105] J. Hough et al. 'The development of long baseline gravitational radiation detectors at Glasgow University'. In: *Gravitation, Geometry and Relativistic Physics. Lecture Notes in Physics* 212 (1984). DOI: `https://doi.org/10.1007/BFb0012592`.

[106] R. Hull. *Properties of Crystalline Silicon*. Inst. of Engineering and Technology, 1999. ISBN: 978-0-85-296933-5.

[107] S. H. Huttner et al. 'Techniques in the optimization of length sensing and control systems for a three-mirror coupled cavity'. In: *Class. Quantum Grav.* 23 (2008). DOI: `https://doi.org/10.1088/0264-9381/25/23/235003`.

[108] A. V. Inyushkin et al. 'Ultrahigh thermal conductivity of isotopically enriched silicon'. In: *J. Appl. Phys.* 123 (2018). DOI: `https://doi.org/10.1063/1.5017778`.

[109] S. Rowan J. Hough and B.S. Sathyaprakash. 'The Search for Gravitational Waves'. In: *J. Phys. B: At. Mol. Opt. Phys.* 38 (2005). DOI: `https://doi.org/10.1088/0953-4075/38/9/004`.

[110] S. J. Poon J. W. Sharp and H. J. Goldsmis. 'Boundary Scattering and the Thermoelectric Figure of Merit'. In: *physica status solidi (a)* 187 (2001). DOI: `https://doi.org/10.1002/1521-396X(200110)187:2%3C507::AID-PSSA507%3E3.0.CO;2-M`.

[111] T. Ono J. Yang and M. Esashi. 'Energy Dissipation in Submicrometer Thick Single-Crystal Silicon Cantilevers'. In: *Journal of Microelectromechanical Systems* 11 (2002). DOI: `https://doi.org/10.1109/JMEMS.2002.805208`.

[112] M. Constancio Jr et al. 'Silicon emissivity as a function of temperature'. In: *International Journal of Heat and Mass Transfer* 157 (2020). DOI: `https://doi.org/10.1016/j.ijheatmasstransfer.2020.119863`.

[113] A. Steinbach K. B. MacAdam and C. Wieman. 'A narrowband tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb'. In: *Am. J. Phys.* 60 (1992). DOI: `https://doi.org/10.1119/1.16955`.

[114] T. Kajita. 'KAGRA - Status and Prospect'. In: *General Relativity and Beyond Symposium* (2024).

[115] D. P. Kapasi et al. 'Tunable narrow-linewidth laser at 2um wavelength for gravitational wave detector research'. In: *Optics Express* 28 (2020). DOI: `https://doi.org/10.1364/OE.383685`.

[116] S. Kasseb and G. Al-Hariry. *Electronics Cooling Part B: Heat Transfer Principals in Electronics Cooling*. URL: `http://pathways.cu.edu.eg/ec/`. (accessed: 26.6.24).

[117] A. Khalaidovski et al. 'Evaluation of heat extraction through sapphire fibers for the GW observatory KAGRA'. In: *Class. Quantum Grav.* 31 (2014). DOI: `https://doi.org/10.1088/0264-9381/31/10/105004`.

[118] K. M. Kim et al. ‘Maximum stable zone length in float-zone growth of small-diameter sapphire and silicon crystals’. In: *J. Appl. Phys.* 50 (1979). DOI: https://doi.org/10.1063/1.326410.

[119] X. Koroveshi et al. ‘Cryogenic payloads for the Einstein Telescope: Baseline design with heat extraction, suspension thermal noise modeling, and sensitivity analyses’. In: *Phys. Rev. D* 108 (2023). DOI: https://doi.org/10.1103/PhysRevD.108.123009.

[120] R. Kubo. ‘The fluctuation-dissipation theorem’. In: *Reports on Progress in Physics* 29 (1966).

[121] P. Kwee et al. ‘Stabilized high-power laser system for the gravitational wave detector advanced LIGO’. In: *Optics Express* 20 (2012). DOI: https://doi.org/10.1364/OE.20.010617.

[122] P. Kwee and B. Willke. ‘Automatic laser beam characterization of monolithic Nd:YAG nonplanar ring lasers’. In: *Appl. Opt.* 47 (2008). DOI: https://doi.org/10.1364/AO.47.006022.

[123] S. Xuan L. Yan and X. Gong. ‘Shock isolation performance of a geometric anti spring isolator’. In: *Journal of Sound and Vibration* 413 (2018). DOI: https://doi.org/10.1016/j.jsv.2017.10.024.

[124] Caltech/MIT/LIGO Laboratory. *LIGO Technology Development and Migration*. URL: https://advancedligo.mit.edu/sus.html. (accessed: 17.05.2024).

[125] L. A. Ladino and H. S. Rondon. ‘Determining the damping coefficient of a simple pendulum oscillating in air’. In: *Phys. Educ.* 52 (2017). DOI: https://doi.org/10.1088/1361-6552/aa6431.

[126] P. Leaci, the LIGO Scientific Collaboration and the Virgo Collaboration. ‘Searching for continuous gravitational wave signals using LIGO and Virgo detectors’. In: *Journal of Physics: Conf. Series* 354 (2012).

[127] J. Levine. ‘Early Gravity-Wave Detection Experiments, 1960-1975’. In: *Physics in Perspective* 6 (2004).

[128] LIGO. *2017 Nobel Prize in Physics Awarded to LIGO Founders*. URL: https://www.ligo.caltech.edu/page/press-release-2017-nobel-prize. (accessed: 27.4.2021).

[129] LIGO. *Image Gallery, LIGO Hanford Observatory*. URL: https://www.ligo.org/multimedia/gallery/lho.php. (accessed: 18.08.2024).

[130] K. L. Dooley for the LIGO Scientific Collaboration. ‘Status of GEO 600’. In: *J. Phys.: Conf. Ser* 610 (2015). DOI: https://doi.org/10.1088/1742-6596/610/1/012015.

[131] Virgo Collaboration LIGO Scientific Collaboration and KAGRA Collaboration. ‘Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3’. In: *Phys. Rev. X* 13 (2023). DOI: https://doi.org/10.1103/PhysRevX.13.011048.

[132] D. Lindley. 'A Fleeting Detection of Gravitational Waves'. In: *Phys. Rev. Focus* 16 (2005).

[133] Z. Liu. 'Temperature-Dependent Elastic Constants and Youngs Modulus of Silicon Single Crystal'. In: *JACoW* MEDSI2020 (2021). DOI: `http://doi.org/10.18429/JACoW-MEDSI2020-WEPC09`.

[134] J.E. Logan, J. Hough and N.A. Robertson. 'Aspects of the thermal motion of a mass suspended as a pendulum by wires'. In: *Phys. Lett. A* 183 (1993). DOI: `https://doi.org/10.1016/0375-9601(93)91161-W`.

[135] J. Macarthur. 'Towards surpassing the standard quantum limit using optical springs'. PhD thesis. University of Glasgow, 2014.

[136] M. Maggiore. *Gravitational Waves Volume 1: Theory and Experiments*. Oxford University Press, 2008. ISBN: 978-0-19-857074-5.

[137] G. L. Mansell. 'Squeezed light sources for current and future interferometric gravitational-wave detectors'. PhD thesis. Australian National University, 2018.

[138] I. Martin. 'Studies of materials for use in future interferometric gravitational wave detectors.' PhD thesis. University of Glasgow, 2009.

[139] D. V. Martynov et al. 'Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy'. In: *Phys. Rev. D* 93 (2016). DOI: `https://doi.org/10.1103/PhysRevD.93.112004`.

[140] F. Matichard et al. 'Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance'. In: *LIGO Document P1200040-v55* (2019).

[141] E. Mauceli et al. 'The Allegro gravitational wave detector: Data acquisition and analysis'. In: *Phys. Rev. D* 7 (1997). DOI: `https://doi.org/10.1103/PhysRevD.54.1264`.

[142] D. F. McGuigan et al. 'Measurements of the Mechanical Q of Single-Crystal Silicon at Low Temperatures'. In: *Journal of Low Temperature Physics* 30 (1978). DOI: `https://doi.org/10.1007/BF00116202`.

[143] M. Mehmet et al. 'Squeezed light at 1550 nm with a quantum noise reduction of 12.3 dB'. In: *Optics Express* 19 (2011). DOI: `https://doi.org/10.1364/OE.19.025763`.

[144] Y. Michimura et al. 'Mirror actuation design for the interferometer control of the KAGRA gravitational wave telescope'. In: *Class. Quantum Grav.* 34 (2017). DOI: `https://doi.org/10.1088/1361-6382/aa90e3`.

[145] C. Mueller et al. 'The advanced LIGO input optics'. In: *Rev. Sci. Instrum.* 87 (2016). DOI: `https://doi.org/10.1063/1.4936974`.

[146] A. Muiznieks et al. *Handbook of crystal growth (second edition)*. Elsevier, 2015. ISBN: 978-0-444-56369-9.

[147] R. Nawrodt et al. 'High mechanical Q-factor measurements on silicon bulk samples'. In: *J. Phys.: Conf. Ser* 122 (2008). DOI: `https://doi.org/10.1088/0264-9381/31/22/224001`.

[148] R. Nawrodt et al. ‘Mirror thermal noise calculation for ET’. In: *ET-027-09* (2009).

[149] E. Oelker. Private Communication.

[150] M. Phelps. ‘Hydroxide catalysis and indium bonding research for the design of ground-based gravitational wave detectors.’ PhD thesis. University of Glasgow, 2018.

[151] M. Phelps et al. ‘Strength of hydroxide catalysis bonds between sapphire, silicon, and fused silica as a function of time’. In: *Phys. Rev. D* 98 (2018). DOI: `https://doi.org/10.1103/PhysRevD.98.122003`.

[152] Hamamatsu Photonics. *InGaAs PIN photodiode*. URL: `https://www.hamamatsu.com/eu/en/product/optical-sensors/infrared-detector/ingaas-photodiode/G12182-005K.html`. (accessed: 8.11.24).

[153] M. V. Plissi et al. ‘Aspects of the suspension system for GEO 600’. In: *Rev. Sci. Instrum.* 69 (1998). DOI: `https://doi.org/10.1063/1.1149054`.

[154] M. V. Plissi. ‘Cantilever blade analysis for Advanced LIGO’. In: *LIGO Document LIGO-T030107-00-D* (2003).

[155] M. V. Plissi et al. ‘GEO 600 triple pendulum suspension system: Seismic isolation and control’. In: *Rev. Sci. Instrum.* 71 (2000). DOI: `https://doi.org/10.1063/1.1150645`.

[156] Cosmic Explorer Project. *Cosmic Explorer Overview*. URL: `https://cosmicexplorer.org/`. (accessed: 27.08.2024).

[157] M. Punturo et al. ‘The Einstein Telescope: a third-generation gravitational wave observatory’. In: *Class. Quantum Grav.* 27 (2010). DOI: `http://doi.org/10.1088/0264-9381/27/19/194002`.

[158] D. Reitze et al. ‘Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO’. In: *arXiv preprint arXiv:1907.04833* 27 (2019). DOI: `https://doi.org/10.48550/arXiv.1907.04833`.

[159] L. Ricci et al. ‘A compact grating-stabilized diode laser system for atomic physics’. In: *Optics Communications* 117 (1995). DOI: `https://doi.org/10.1016/0030-4018(95)00146-Y`.

[160] N. A. Robertson et al. ‘Quadruple suspension design for Advanced LIGO’. In: *Class. Quantum Grav.* 19 (2002). DOI: `https://doi.org/10.1088/0264-9381/19/15/311`.

[161] S. Rowan et al. ‘Investigation of mechanical loss factors of some candidate materials for the test masses of gravitational wave detectors’. In: *Phys. Lett. A* 265 (2000). DOI: `https://doi.org/10.1016/S0375-9601(99)00874-9`.

[162] S. Rowan et al. ‘Thermal noise and material issues for gravitational wave detectors’. In: *Phys. Lett. A* 347 (2005).

[163] P. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, 2017. ISBN: 978-981-3143-07-4.

[164] P. Saulson et al. ‘Thermal noise in mechanical experiments’. In: *Phys. Rev. D* 42 (1990). DOI: https://doi.org/10.1103/PhysRevD.42.2437.

[165] R. Savage and P. Schwinberg. ‘Test Plan - Table-top Frequency Stabilization Servo (TTFSS)’. In: *LIGO Document LIGO-E040418-04-W* (2004).

[166] B. Schutz. *A First Course in General Relativity*. Cambridge University Press, 2009. ISBN: 978-0-521-88705-2.

[167] B. Schutz. ‘Determining the Hubble constant from gravitational wave observations’. In: *Nature* 323 (1986). DOI: https://doi.org/10.1038/323310a0.

[168] B. Shapiro. ‘Electronic Setup and Testing of Advanced LIGO Suspensions’. In: *LIGO Document LIGO-E1000078-v1* (2010).

[169] K. Shimamura et al. ‘Silicon Single Crystal Fiber Growth by Micro Pulling Down Method’. In: *Japenese J. Appl. Phys.* 35 (1996). DOI: https://doi.org/10.1143/JJAP.35.L793.

[170] S. Song et al. ‘Localised structuring of metal-semiconductor cores in silica clad fibres using laser-driven thermal gradients’. In: *Nature Communications* 13 (2022). DOI: https://doi.org/10.1038/s41467-022-29975-1.

[171] A. Spencer. ‘Advanced Techniques in Laser Interferometry for Current and Future Gravitational Wave Detectors’. PhD thesis. University of Glasgow, 2020.

[172] A. Spencer. ‘Instrument Science update from Glasgow Group: ET Symposium 2024 Maastricht’. In: *ET-0275A-24* (2024).

[173] A. Spencer et al. ‘The Glasgow Cryogenic Interferometry Facility’. In: *LIGO Document LIGO-G2201488-v3* (2022).

[174] J. Steinlechner. ‘Development of mirror coatings for gravitational-wave detectors’. In: *Phil. Trans. R. Soc. A.* 376 (2018). DOI: https://doi.org/10.1098/rsta.2017.0282.

[175] J. Steinlechner et al. ‘Silicon-Based Optical Mirror Coatings for Ultrahigh Precision Metrology and Sensing’. In: *Phys. Rev. Lett.* 120 (2018). DOI: https://doi.org/10.1103/PhysRevLett.120.263602.

[176] K. A. Strain and B. J. Meers. ‘Experimental demonstration of dual recycling for interferometric gravitational-wave detectors’. In: *Phys. Rev. Lett.* 66 (1991). DOI: https://doi.org/10.1103/PhysRevLett.66.1391.

[177] NSF MPS AC Subcommittee. ‘Next-Generation Gravitational-Wave Detector Concepts’. In: *National Science Foundation Report* (2024).

[178] W. D. Sylwestrowicz. ‘Mechanical properties of single crystals of silicon’. In: *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 7 (1962). DOI: https://doi.org/10.1080/14786436208213849.

[179] Stanford Research Systems. *Dynamic Signal Analyzer SR785*. URL: https://thinksrs.com/products/sr785.html. (accessed: 24.09.2024).

[180] N. Cornish T. Robson and C. Liu. 'The construction and use of LISA sensitivity curves'. In: *Class. Quantum Grav.* 36 (2019). DOI: `https://doi.org/10.1088/1361-6382/ab1101`.

[181] H. Tanaka et al. 'Improvement of mechanical loss measurement system of sapphire fibers for the cryogenic suspension system of KAGRA II'. In: *KAGRA Document JGW-G1504281-v1* (2015).

[182] ET Science Team. 'Einstein gravitational wave telescope conceptual design study'. In: *ET-0106C-10* 4 (2011).

[183] ET Science Team. 'Einstein Telescope: Science Case, Design Study and Feasibility Report'. In: *ET-0004A-20* (2020).

[184] The ETpathfinder Team. *ETpathfinder*. URL: `https://www.etpathfinder.eu/`. (accessed: 29.08.2024).

[185] The ETpathfinder Team. 'ETpathfinder Design Report'. In: *Public ET Document* (2020).

[186] Inc Technology Applications. *CRYOCOOLER SERIES (CS) CuTS CS-68B CuTS - Cryocooler Series Thermal Straps for Cryocoolers*. URL: `https://www.techapps.com/cryocooler-series-thermal-straps`. (accessed: 24.6.24).

[187] The LIGO Scientific Collaboration The KAGRA Collaboration and The Virgo Collaboration. 'LIGO-Virgo-KAGRA Cumulative Detection plot - O1-O4b'. In: *LIGO Document G2302098-v20* (2024).

[188] The LIGO Scientific Collaboration The KAGRA Collaboration and The Virgo Collaboration. *LIGO, VIRGO AND KAGRA OBSERVING RUN PLANS*. URL: `https://observing.docs.ligo.org/plan/`. (accessed: 02.09.2024).

[189] The LIGO Scientific Collaboration The KAGRA Collaboration and The Virgo Collaboration. *LIGO/Virgo/KAGRA Public Alerts*. URL: `https://gracedb.ligo.org/superevents/public/O4/`. (accessed: 03.09.2024).

[190] The LIGO Scientific Collaboration The KAGRA Collaboration and The Virgo Collaboration. 'Observing Scenario timeline graphic, post-O3'. In: *LIGO Document G2002127-v26* (2024).

[191] Marc T. Thompson. *Intuitive Analog Circuit Design*. Newnes, 2014. ISBN: 9780124058668.

[192] Thorlabs. *SAF1900S - Mounted SAF Gain Chip, Half Butterfly Pkg, CWL = 1900 nm, SM Fiber*. URL: `https://www.thorlabs.de/thorProduct.cfm?partnumber=SAF1900S`. (accessed: 23.10.24).

[193] C. Torrie. 'Development of Suspensions for the GEO 600 Gravitational Wave Detector'. PhD thesis. University of Glasgow, 2001.

[194] Y. S. Touloukian et al. *Thermophysical Properties of Matter: Volume 1. Thermal Conductivity, Metallic Elements and Alloys*. Plenum Publishing Corporation, 1970. ISBN: 146-159602-5.

[195] M. Tse et al. 'Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy'. In: *Phys. Rev. Lett.* 123 (2019). DOI: `https://doi.org/10.1103/PhysRevLett.123.231107`.

[196] T. Uchiyama et al. 'Mechanical quality factor of a cryogenic sapphire test mass for gravitational wave detectors'. In: *Phys. Lett. A* 261 (1999). DOI: `https://doi.org/10.1016/S0375-9601(99)00563-0`.

[197] T. Ushiba et al. 'Cryogenic suspension design for a kilometer-scale gravitational-wave detector'. In: *Class. Quantum Grav.* 38 (2021). DOI: `https://doi.org/10.1088/1361-6382/abe9f3`.

[198] S. P. Vyatchanin V. B. Braginsky. 'Corner reflectors and Quantum-Non-Demolition Measurements in gravitational wave antennae'. In: *Phys.Lett. A.* 324 (2004). DOI: `https://doi.org/10.1016/j.physleta.2004.02.066`.

[199] V. P. Mitrofanov V. B. Braginsky and V. I. Panov (Translated by Erast Gliner). *Systems with Small Dissipation*. University of Chicago Press, 1987. ISBN: 978-0226070735.

[200] M. van Veggel. 'Cryogenic suspensions for future gravitational wave detectors'. In: *Conference Presentation: Gravitational Waves Advanced Detector Workshop* (2018).

[201] M. van Veggel et al. 'Final Design Document ETM/ITM ears'. In: *LIGO Document LIGO- T0900447-v3* (2009).

[202] L. Vieira et al. 'Simulation of crucible-free growth of monocrystalline silicon fibres for mirror suspension in gravitational-wave detectors'. In: *Journal of Crystal Growth* 629 (2024). DOI: `https://doi.org/10.1016/j.jcrysgro.2023.127549`.

[203] F. F. Villa. *Silicon Sensors and Actuators: Silicon Properties and Crystal Growth*. Springer, Cham, 2022. ISBN: 978-3-030-80134-2.

[204] Virgo. *Virgo Image Gallery*. URL: `https://www.virgo-gw.eu/images/?_paged=3`. (accessed: 18.08.2024).

[205] Z. Wang et al. 'High-quality semiconductor fibres via mechanical design'. In: *Nature* 626 (2024). DOI: `https://doi.org/10.1038/s41586-023-06946-0`.

[206] J. Weber. 'Evidence for Discovery of Gravitational Radiation'. In: *Phys. Rev. Lett.* 22 (1969).

[207] J. Weber. 'Gravitational-Wave-Detector Events'. In: *Phys. Rev. Lett.* 20 (1968).

[208] University of Western Australia. *Gingin Gravity Precinct*. URL: `https://www.uwa.edu.au/Research-uwa/Centres/Gingin-Gravity-Precinct`. (accessed: 27.08.2024).

[209] G. K. White and M. L. Minges. 'Thermophysical Properties of Some Key Solids'. In: *International Journal of Thermophysics* 15 (1994). DOI: `https://doi.org/10.1007/BF01458841`.

[210] B. Willke. 'Stabilized lasers for advanced gravitational wave detectors'. In: *Laser Photonics Rev.* 4 (2010). DOI: `https://doi.org/10.1002/lpor.200900036`.

[211] G. Woan. *The Cambridge Handbook of Physics Formulas*. Cambridge University Press, 2000. ISBN: 0-521-57507-9.

[212] Wolfram. *WOLFRAM MATHEMATICA*. URL: `https://www.wolfram.com/mathematica/`. (accessed: 6.6.24).

[213] J. L. Wright. 'Optical springs to create macroscopic optical traps and negative inertia for gravitational wave detectors'. PhD thesis. University of Glasgow, 2022.

[214] O. Tabata Y. B. Gianchandani and H. Zappe. *Comprehensive microsystems, vol. 1*. Elsevier, 2008. ISBN: 978-0-444-52190-3.

[215] T. Yamada et al. 'High performance thermal link with small spring constant for cryogenic applications'. In: *Cryogenics* 116 (2021). DOI: `https://doi.org/10.1016/j.cryogenics.2021.103280`.

[216] T. Yamada. 'Low-Vibration Conductive Cooling of KAGRA Cryogenic Mirror Suspension.' PhD thesis. University of Tokyo, 2020.

[217] H. D. Young and R. A. Freedman. *University Physics with Modern Physics*. Pearson, 2020. ISBN: 9780136874331.

[218] A. Zimmer et al. 'Mechanical losses in low loss materials studied by Cryogenic Resonant Acoustic spectroscopy of bulk materials (CRA spectroscopy)'. In: *Journal of Physics: Conf. Series* 92 (2007). DOI: `https://doi.org/10.1088/1742-6596/92/1/012095`.