Simon C. Tait
University of Glasgow, 2020/2021

Studies of Mechanical and Optical Properties of Thin Film Coatings for Future Gravitational Wave Detectors

School of Physics and Astronomy

First Supervisor: Dr. Iain Martin
Second Supervisor: Professor Sheila Rowan

Contents

List of Tables

List of Figures

aLIGO	Advanced Laser Interferometric Gravitational Wave Detector
AdV	Advanced Virgo
BH	Black Hole
BBH	Binary Black Hole
BNS	Binary Neutron Star
CAD	Computer Aided Design
CMB	Cosmic Microwave Background
CBC	Compact Binary Coalescence
CP	Compensation Plate
CryoGeNS	Cryogenic Gentle Nodal Support
CTN	Coating Thermal Noise
cSi	crystalline Silicon
DARM	Differential Arm Length
ECR	Electron Cyclotron Resonance
EFI	Electric Field Intensity
EM	Electro-Magnetism/Electro-Magnetic
ERM	End Reaction Mass
ESA	European Space Agency
ESD	Electro Static Drive
ECR	Electron Cyclotron Resonance
ET	Einstein Telescope
ETM	End Test Mass
FEA	Finite Element Analysis
FWHM	Full With at Half Maximum
GeNS	Gentle Nodal Support
GRB	Gamma Ray Burst
GW	Gravitational Wave
GWD	Gravitational Wave Detector
HAM	Horizontal Access Module
HAR	Hughes Aircraft Research Laboratories
HeNe	Helium-Neon (laser)
HEPI	Hydraulic External Pre Isolator
HR	Highly Reflecting / Highly Reflective
HWS	Hartmann Wave-front Sensor
IBS	Ion Beam Sputtering
IFO	Interferometer

IMC	Input Mode Cleaner
INTEGRAL	International Gamma-Ray Astrophysics Laboratory
IR	Infra Red
ITM	Input Test Mass
KAGRA	Kamioka Gravitational Wave Detector
LF	Low Frequency
LHe	Liquid Helium
LHS	Left Hand Side
LIGO	Laser Interferometric Gravitational Wave Detector
LISA	Laser Interferometer Space Antenna
LHO	LIGO Hanford
LLO	LIGO Livingston
LMA	Laboratoire des Matériaux Avancés
LN2	Liquid Nitrogen
MIT	Massachusetts Institute of Technology
MMI	Multi-Mode Interaction
NB	Nota bene
NASA	National Aeronautics and Space Administration
NEMO	Neutron Star Extreme Matter Observatory
NGC	New General Catalogue
NS	Neutron Star
NSBH	Neutron Star Black Hole
OMC	Optical Mode Cleaner
PBS	Polarising Beam Splitter
PCI	Photothermal Common-path Interferometry
PEEK	Polyether-Ether-Ketone
PID	Proportional–Integral–Derivative
POC	Points of Contact
PRC	Power Recycling Cavity
PSD	Power Spectral Density
PTFE	Polytetrafluoroethylene
QWL	Quarter Wave Layer
RBS	Rutherford Back Scattering
RF-ICP	Radio Frequency Inductive Coupled Plasma

RH	Ring Heater
RHS	Right Hand Side
RLVIP	Reactive Low Voltage Ion Plating
ROI	Region of Interest
SIMS	Secondary Ion Mass Spectrometry
SNR	Signal to Noise Ratio
SRM	Signal Recycling Mirror
SSE	Sum of Squares Error
STD	Standard Deviation
TCS	Thermal Compensation System
TEM	Tunnelling Electron Microscopy
TN	Thermal Noise
US	United States of America
UTC	Coordinated Universal Time
UV	Ultra Violet
VCA	Voice Coil Actuator
XRD	X-Ray Diffraction
ZPK	Zero Pole gain

Acknowledgements

Firstly, I would like to express my deep and sincere gratitude to my first supervisor and mentor, Dr Iain Martin, for his invaluable guidance throughout my PhD. I would like to thank him for his unwavering support, for the late nights puzzling things out in the lab, and for putting up with the 3 am "can I just check this one thing" conversations. I would hope to do half as good a job if I were in his position.

I would like to thank Dr Peter Murray, for going over and above what was required during this process. I would also like to thank Dr Jessica Steinlechner, Prof. Stuart Reid. I have learned so much from each one of you all, and I would not be the scientist I am today without your help.

I am grateful for the aid of my second supervisor, Professor Sheila Rowan for her aid throughout my PhD and for allowing me to join the Institute for Gravitational Research (IGR), enabling me to spend the past four years doing what I love.

Spending four years in the IGR has been some of the most stressful and the most enjoyable times of my life, and I will always cherish the people that I have had the privilege to know during my time here. I would like to thank the (now) Dr Raymond Robie and (now) Dr Mark Fletcher for the terrible office pranks, for putting up with my terrible excel sheets and for taking me under their wings when I first started. I learned a lot from both of you. I also need to mention the key contributions from Dr Angus Bell, Dr Bryan Barr, Dr Alan Cumming, Dr Liam Cunningham, Mr Russell Jones, Dr Jamie Scott and Dr Borja Sorazu, and Mr Colin Craig.

I also need to thank my office mates Dr Joeseph Bayley, Dr Rebecca Douglas, Mr Hunter Gabbard, Dr Hafiza Isa, Maya M. Kinley-Hanlon, Mr Graeme McGhee and

Dr Richard Middlemiss for your advice and for keeping things exciting.

I also need to thank those who have kept me afloat personally throughout this process. In particular, my family, my mum and sister, gave me the confidence to pursue physics as a career and have always been there for me even when it wasn't physically possible. Lastly, I want to thank my partner and rock, Laurence Datrier, for her everlasting love and support. You were always there for me when times got hard. You were there to drag me away from my computer when I hadn't been outside for days and remind me what was important. The past year has been particularly hard, and I definitely would not have made it this far if it wasn't for you.

Preface

This thesis is an account of the research carried out between October 2016 and January 2021. The majority of this time was spent in the Institute for Gravitational Research (IGR), where the focus of this work was on optical and mechanical characterisation of dielectric thin film coatings, which could be applicable for future ground-based gravitational wave detectors. Four months between January and May 2018 were spent working at the Laser Interferometer Gravitational Wave Observatory (LIGO) in Livingston, Louisiana, as part of the LIGO Fellowship programme. Over this time, the focus of work was the development infrastructure to indirectly measure optical absorption of highly reflecting (HR) test mass coatings.

Chapter 1 describes the nature of gravitational radiation and its possible sources. It also contains a description of the basic principles of using laser interferometry to detect gravitational waves.

Chapter 2 contains an in-depth description of coating thermal noise, which is a lower limit on the sensitivity of current gravitational wave detectors. This chapter also contains a description of the status of highly reflecting mirror coatings in the gravitational wave field.

Chapter 3 details the work carried out between January and May 2018 while working at the LIGO Livingston site. This involved modelling the heating of LIGO test masses, whose experimental frequencies were provided by Dr Carl Blair (University of Western Australia). Modelling of heat-flow through and optical scattering from the front face of a LIGO test masses was aided by Haoyu Wang (University of Birmingham) and Dr Carl Blair. Experimental tracking of test mass eigenmode frequencies was carried out using a MATLAB script based on previous

functions developed by Dr Carl Blair, with the aid of Dr Arnaud Pele (Massachusetts Institute of Technology). Implementations of zero pole gain (ZPK) transfer functions between experimental frequency shifts was created by Dr Carl Blair.

Chapter 4 contains measurements of the optical absorption of thin-film coatings using photothermal common-path interferometry (PCI). Measurements of thin-film Ta_2O_5 were all carried out by the author; however, guidance was sometimes sought from Dr Angus Bell (formerly University of Glasgow), Dr Jessica Stienlechner (Maastricht University), Dr Iain Martin (University of Glasgow) and Zeno Tornasi (University of Glasgow). Rutherford Backscattering (RBS) measurements of Ta_2O_5 coatings were carried out buy by Alexandre Lussier (Université de Montréal) and reviewed by Prof. François Schiettekatte, (Université de Montréal). The stabilisation of polarisation fluctuations of the 1064 nm fibre coupled laser and installation of pickoff optical components were carried out with the aid of Dr Angus Bell. The development of 2D optical absorption mapping was aided by a Python script that can synchronise multiple time-series, created by Dr Joseph Bayley (University of Glasgow). This was then used as an input for a MATLAB script written by the author. Calculations of k, the imaginary component of refractive index, were carried out using MATLAB functions written by Zeno Tornasi, which were then adapted by the author such that analysis of 2D absorption maps could be carried out.

The optical absorption of direct heat-treated samples at 2000 nm were carried out by Dr Jessica Stienlechner. Optical absorption of multi-material samples that were incrementally heat treated was carried out by the author between 100°C and 400°C 1064 nm, 1550 nm and 2000 nm. Due to restrictions placed on laboratory work starting in March 2020, heat treatment and measurements of these samples at 500°C were carried out by Graeme McGhee (University of Glasgow), working under the direction of the author. Dr Jessica Stienlechner provided calculations of k for multimaterial coatings. Samples of Zr:Ta_2O_5 and $LaTiO_3$ were provided by Dr Svetoslava Angelova (University of Strathclyde). The author carried out all optical absorption measurements and analysis.

Chapter 5 details the development of a gentle nodal support which can measure

the mechanical loss of coated $\varnothing = 3$" disks. This technique was first published by Dr Elisabetta Cesarini (Università degli Studi di Roma), and this was used as the premise for the apparatus, which the author built. The SolidWorks drawings and design of the stage and the integral table of the vacuum system were carried out by Russell Jones (University of Glasgow) with input from the author. The author carried out finite element modelling of uncoated cantilevers disks with input from Dr Mark Fletcher (formerly the University of Glasgow). The author carried out measurement of all uncoated cantilever and disk substrates with occasional guidance given by Dr Raymond Robie (California Institute of Technology), Dr Iain Martin and Dr Peter Murray (University of Glasgow). The mathematical form describing the decay of two beating oscillations was carried out by the author with the aid of Dr Christopher Messenger (University of Glasgow). The analysis of the two modes' beating decays used a python package written by Dr Gabriele Vajente (California Institute of Technology), where the author made adaptions. All measurements of silica disks coated with the *Upper Stack*, *Lower Stack* and *Full Stack* were carried out by the author between the as-deposited state and 300°C. Above this heat treatment temperature, measurements of these samples were carried out with the aid of Maya M. Kinley-Hanlon (University of Glasgow). The author carried out the analysis of ringdown measurements with the aid of Maya M. Kinley-Hanlon and Dr Peter Murray. The author carried out all finite element modelling of multi-material disks. Calculations of the Brownian thermal noise of the *Full Stack* full-stack HR coating were carried out by the author.

Chapter 6 details the adaption of a gentle nodal support in Chapter 5 to operate at cryogenic temperatures. The Solidworks drawings for the Disk retention plate, PTFE translation rods and external Cryostat support structure by Russell Jones (University of Glasgow) with input from the author. Russell Jones carried out FEA modelling and optimisation of the CryoGeNS baseplate. Setting up of the cryostat, levelling and mounting of the cryostat was carried out by the author. The author also carried out the wiring of temperature sensors with the aid from Maya M. Kinley-Hanlon and Dr Iain Martin. The author carried out all finite element modelling detailed in this chapter. The author carried out measurements of the initial disk calibration with Miss Cassady Smith's support (formerly the University of Glasgow), Graeme McGhee, Dr Iain Martin, Dr Peter Murray and Maya M. Kinley-Hanlon. Conversions between the temperature of the sample and internal

cryostat sensors below 220 K were carried out by the author, with conversions at higher temperatures aided by Graeme McGhee. During mechanical loss measurements, Graeme McGhee also helped to measure mechanical modes throughout the temperature range.

Chapter 7 contains measurements of $\varnothing = 2$" crystalline silicon disks and cSi cantilevers coated with a set of prototype multimaterial coatings. Dr Peter Murray carried out measurements and ringdown analysis of all cantilever samples. In this chapter, mechanical loss measurements of the same coated cSi disks were carried out using two different measurement systems, the CryoGeNS liquid-cooled system developed by the author and a Pulse Tube cooled cryostat. Measurements of uncoated disks in the CryoGeNS system of an uncoated disk and disks coated with the *Lower Stack* and *Full Stack* were carried out by the author with the aid of Graeme McGhee. Measurements of the *Upper Stack*, *Lower Stack*, and *Full Stack* coatings using the Pulse Tube cooled cryostat were carried out by Dr Peter Murray and Maya Kinley-Hanlon. The author carried out thermoelastic and Brownian thermal noise calculations of these samples with input from Dr Iain Martin and Dr Jessica Stienlechner.

Summary

Gravitational radiation in the form of gravitational waves was the last prediction to be verified from Einstein's general theory of relativity. Einstein suggested that when a body or bodies with an asymmetric distribution of mass accelerated, energy from the motion would create distortions in space-time which would propagate in all directions at the speed of light.

Until the first confirmed observation of gravitational waves from the coalescence of two black holes, this theory had not been experimentally proven. The first gravitational wave (GW) event, 'GW150914' confirmed Einstein's predictions, with the event releasing $\approx 3\,M_{\odot}$ worth of energy as gravitational radiation during the collision. Since this event, more than 50 confident GW events have been detected, including the first observation of an extremely rare kilonova event after the collision of two neutron stars.

Gravitational waves exert fluctuating strains on space as they propagate, resulting in changes in length of objects that they pass through. Current gravitational wave detectors use laser interferometry to measure this effect using a single laser source and beamsplitter. Two perpendicular laser beams are created and used to monitor the positions of suspended mirrors at the ends of km-scale perpendicular arms. The laser beams reflected from the mirrors are recombined at the beam splitter, creating an interference pattern. Any changes in these mirrors' position then result in a differential change in the arm length inside the detector, altering the generated interference pattern. As the expected change in differential arm length produced by a gravitational wave event is $<1\times10^{-18}\,\mathrm{m}$, all other sources of motion, or noise, must be reduced to exceedingly low levels to maximise the sensitivity to such events.

Throughout the range of frequencies to which a gravitational wave detector is sensitive too, its highest sensitivity occurs between $\sim$50 Hz and 150 Hz. In this frequency band, thermal noise stemming from thermal vibrations in the materials used to create highly reflecting mirror coatings for each test mass limits the sensitivity of the detector. Each material's contributions to the level of thermal noise are proportional to its level of mechanical loss, its temperature and dimensions of the laser beam on its surface.

This thesis will focus on the development of coating materials with low mechanical loss and low optical absorption, which can be used to decrease levels of thermal noise inside and increase the stability of a gravitational wave detector. As the amount of laser light absorbed into the coating layer also dictates the test mass's thermal state, a large part of this research will also focus on this aspect of coating measurement. A large part of the work in this thesis involves the first experimental verification of the so-called 'multimaterial coating' principle, through testing the optical absorption and room-temperature and cryogenic mechanical loss, of an example of this type of novel coating design.

Chapter 1 describes the nature of gravitational radiation and its possible sources. An introduction to the experimental interferometry techniques used in a gravitational wave detector is considered, and notable sources of noise are summarised.

Chapter 2 provides a detailed summary of coating thermal noise in gravitational wave detectors. This chapter also introduces some of the recent advancements and current avenues of research in HR gravitational wave detector coatings.

Chapter 3 is an account of work carried out by the author at the LIGO Livingston Observatory to develop a technique for monitoring the absorption of the detector mirrors in situ. By studying the resonant frequencies of coated test masses in a gravitational wave detector, a relationship between frequency and the change in test mass temperature by laser heating can be produced. If the total laser power and the level of optical absorption of the coated optic are known, predictions of how its resonant frequencies will change can be modelled using finite element analysis

(FEA). If the optical absorption of the coating material at the time of deposition was known, the shift in resonant frequencies of the test mass in response to laser heating could be used to predict any changes in the absorption of the optic.

Chapter 4 discusses the experimental Photothermal Common-path Interferometry (PCI) technique used by the author to measure the optical absorption of thin-film coatings. This technique is used to study the peculiar changes in optical absorption of tantalum pentoxide (Ta_2O_5) coatings after heat treatment in a laboratory atmosphere and under low vacuum. Implementation of a polarisation stabilisation and power correction system into the PCI techniques has allowed for 2-dimensional absorption maps to produce these samples. In this chapter, measurements of a novel multimaterial coating designed to decrease coating thermal noise and optical absorption are studied. The material is subjected to two different heat treatment studies, and its optical absorption is characterised using 1064 nm, 1550 nm and 2000 nm laser light.

In Chapter 5, the methods and techniques used to measure the mechanical loss (internal friction) of coating materials are introduced. Throughout this chapter, the development of an automated measurement technique used to characterise the mechanical loss $\varnothing = 3$" (76.2 mm), $t = 2.6$ mm is discussed and compared to existing measurement techniques.

Chapter 6 describes the use of the technique developed in Chapter 5 to study the mechanical loss of the same novel multimaterial coating as a function of heat-treatment temperature. These measurements allow the level of coating thermal noise produced by each multimaterial coating to be calculated using the methods described in Chapter 2.

The development of the gentle nodal support described in Chapter 5 is continued in Chapter 7, upgrading the apparatus to function at cryogenic temperatures ($80\,\mathrm{K} < T < 293\,\mathrm{K}$). By creating an automated gentle nodal support that operates at cryogenic temperatures, the mechanical loss of coating materials can be characterised for third-generation gravitational wave detector applications. This chapter describes the development of the CryoGeNS system and characterisation of uncoated $\varnothing = 2$" (50.8 mm), $t = 360\,\mu$m crystalline silicon disks.

Chapter 8 details cryogenic mechanical loss measurements of the prototype multimaterial coating, carried out using the CryoGeNS nodal support. Coating loss calculated from each disk is compared to measurements of the same samples carried out in a second commercial cryostat and cryogenic measurements of cSi cantilevers coated in the same materials. This allows the level of thermal noise improvement of these coatings to be calculated as a function of temperature.

Chapter 1

Gravitational Wave Detection

1.0.1 Introduction

Gravitational waves are produced by the acceleration of asymmetric distributions of mass. As the waves propagate, they exert fluctuating tidal strains on space and change the length of objects that they pass through. Due to the weak nature of the gravitational force, very large masses and accelerations are required to produce detectable gravitational waves. Such sources include the coalescence of a compact binary system (consisting of neutron stars, black holes or a neutron star and a black hole) and supernovae.

Gravitational waves were directly detected for the first time on September 14$^{\text{th}}$ 2015 by the Advanced Laser Interferometric Gravitational Wave Observatory (aLIGO) [8]. The signal, known as GW150914 (see section 1.3.1), resulted from the inspiral and merger of two black holes and was measured using two 4 km long Michelson-based interferometers one located in Hanford, Washington State and the other in Livingston, Louisiana in the United States of America. Since this event, the Advanced Virgo (AdV) detector [9] in Cascina, Italy and the KAGRA detector [10] near Hida in the Gifu Prefecture of Japan have joined the global network.

Since the first confirmed observation of a gravitational wave signal, the sensitivity of each detector has been upgraded over multiple observation runs, with the number of confirmed events increasing with detector sensitivity. To date, there have been over 40 inspiral signals measured from binary black hole pairs (BBH), or

Binary Neutron Stars (BNS) [11, 12]. Following the detection of gravitational waves from a BNS inspiral (GW170817, discussed in section 1.3.2), the first electromagnetic counterparts to a GW event were detected. Analysis of these measurements confirmed the source as a kilonova event.

In this chapter, the fundamental details of gravitational wave radiation and its sources will be explored. The working principles of current interferometric gravitational wave detectors, and the factors limiting their sensitivity, will be described, and plans for future detectors will be discussed.

1.1 Gravitational Radiation

Just as electromagnetic waves are the product of an acceleration of charge, gravitational waves are produced by an acceleration of mass. This results in a transverse wave that propagates at the speed of light c, imparting strains on space-time. In the electromagnetic case, radiation is di-polar in nature due to balancing positive and negative charges. However, due to the conservation of momentum, mass has no such parallel. The lowest moment of mass distribution which can produce gravitational waves is the quadrupole moment, meaning that GW's can only be produced by an acceleration of an asymmetric distribution of mass. There are two polarisation's of GW, $\hat{h}_{\times}$ and $\hat{h}_{+}$, where h denotes the strain exerted by the wave on space. The effect of a passing gravitational wave can be visualised by considering a ring of freely-falling particles, of diameter L as shown in Fig 1.1.

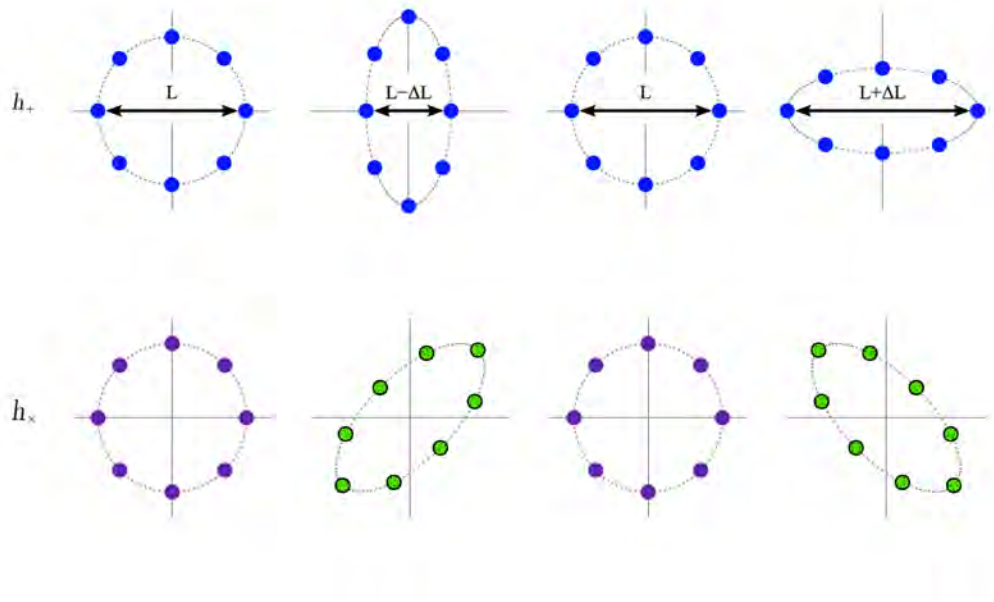

Figure 1.1: Resultant displacement of a ring of freely falling particles when subjected to (a) $\hat{h}_+$ or (b) $\hat{h}_\times$.

A gravitational wave propagating into the page distorts the ring's shape, causing the distances between particles to contract ΔL along the x-axis and expand ΔL along the y-axis. The effect that different polarisations of GW have on the ring of particles is then rotated through 45 ° for the $\hat{h}_\times$ compared to the $\hat{h}_+$ polarised quadrupole. The amplitude of the gravitational wave h, is then given by

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

Typical values of a GW signal passing through a detector would result in a measured shift in the differential arm length of $4{\times}10^{-18}\,\mathrm{m\,Hz^{-\frac{1}{2}}}$ [13], .

1.2 Sources of Gravitational Waves

Sources of gravitational waves are typically classified according to the type of signal they produce. Different strategies are used to search for different types of signal.

1.2.1 Continuous Sources

Continuous sources emit GW's at a constant, or nearly constant, frequency over a long period of time. Examples of these include certain categories of rotating neutron star. Neutron stars are produced by the core collapse of a massive star which exceeds the Chandrasekhar limit($1.4M_{\odot}$ [14]), typically $m{\sim}\,10M_{\odot}$ - $25M_{\odot}$ [15]. The pure nuclear matter at the core of the star cools, the density of the neutron matter increases at its centre. Under certain conditions, the crust's cooling results in a mass distribution that is anti-symmetric, and the neutron star fulfils all requirements to produce gravitational waves.

During the formation of a rotating neutron star, the neutron matter can flow much like a fluid producing instabilities in the stars rotation [16]. This flow of matter produces axial inertial modes (r-modes) in the star, which are believed to dissipate energy through gravitational waves. The accretion of matter from a second nearby star can increase the neutron star's speed of rotation, which is expected to increase the emission of GW's [17]. The exchange of mass in the accretion process may also lead to axial asymmetry of the star, producing gravitational wave radiation [18].

It is expected that gravitational wave sources such as rapidly rotating asymmetric neutron stars produce sinusoidal strain signals, which are considerably below the noise floor in current gravitational wave detectors [19]. Certain classes of highly magnetised rotating neutron stars (pulsars) produce jets of electromagnetic radiation at their poles which can be observed in the electromagnetic spectrum allowing their rotational frequency to be identified. By sampling detector strain data for extensive periods (e.g. longer than 400 days), coherent searches for the expected signals from pulsars can be achieved [20]. This sampling method requires monitoring the rotational period ($f = 2\nu$) of the pulsar.

By monitoring the pulsar, estimates of its principal moment of inertia I_{zz} can be made, allowing the expected strain amplitude h_0 to be calculated using the following expression:

$$h_0 = \frac{4\pi^2 G}{c^4}\frac{I_{zz},\, f^2}{d}\boldsymbol{\epsilon} \tag{1.2}$$

where G is the gravitational constant, c is the speed of light in vacuum, $\boldsymbol{\epsilon}$ is the

ellipticity of the star at a distance d from the observer [21]. This rotation coupling is often referred to as 'wobbling' and is thought to be a source of continuous gravitational waves. Calculations suggest that typical amplitude strain values for a rotating pulsar $h \sim \times 10^{-18} - \times 10^{-19}$ would be expected [22].

1.2.2 Burst Sources

Burst sources produce short-lived signals from a rapid acceleration of mass, such as the explosion of a star or the final merger of two inspiralling compact objects [23]. For a star of $< 5 M_{\odot}$, usually a white dwarf, this typically involves gaining mass from another companion star to produce a type 1 a supernova. For other types of supernovae, the mass of the star is rapidly accelerated, with larger stars emitting more energy in GWs.

Two coalescing objects, such as neutron stars, or a collapse of a massive star, can also produce concentrated beams of high energy radiation in the form of a gamma-ray burst (GRB). Long GRBs (duration $\approx$2 s) can be produced through the collapse of massive stars [24]. Short GRBs (duration $\leq$0.5 s) have been studied in less detail than their long counterparts as they occur less frequently. The first strong evidence for the origins of short GRBs came through observations of the coalescence of two neutron stars [24] (see section 1.3.2).

1.2.3 Stochastic Sources

Stochastic gravitational waves are a combination of GW signals with a random probability distribution or pattern that can be analysed statistically but cannot be precisely predicted. Signals from sources such as BBH, BNS, supernovae, kilonovae and the resultant waves from asymmetrically rotating neutron stars have created a stochastic background which can be observed today. The stochastic background also encompasses the residual gravitational wave remnants, which were produced between approximately 10^{-36} s to 10^{-32} s after the Big Bang [25].

As the study of the cosmic microwave background (CMB) produced new information about the early universe's formation, observations of the stochastic gravitational wave background are also of great interest to cosmology. Further

studies of primordial gravitational waves, which were emitted much sooner after the Big Bang than sources of the CMB, may reveal information about the unsystematic nature of the Big Bang - and by extension, reveal new information about the nature of the universe.

1.2.4 Inspiral Sources

Through optical observations, it is thought that more than 80% of stellar systems exist as binary systems. If the system contains one of the more dense astronomical objects such as a white dwarf star, neutron star, or a black hole, the system is classified as a compact binary. The total orbital kinetic energy of each body in a compact binary system is slowly lost in the form of gravitational waves. As a result, the separation between the bodies decreases.

During this process, the orbital period of the two objects decreases, and energy is lost via the emission of gravitational waves. In order to estimate observed strain produced from an inspiral event, a factor known as the chirp mass $\mathcal{M}$ is used:

$$\mathcal{M} = \mu^{3/5} M^{2/5}, \tag{1.3}$$

where $\mu = \frac{m_1 m_2}{m_1 + m_2}$ is the reduced mass ratio, the total mass of the system is denoted by M, and the mass ratio of the two objects, m_1, m_2. Viewing the orbital frequency as a function of time produces what is known as a 'chirp signal'. As the separation between both objects decreases, the orbital frequency *f(t)* increases until the two objects coalesce. The evolution of frequency over time and $\mathcal{M}$ can then be used to calculate the expected gravitational wave strain [26].

1.3 Gravitational Wave Detections

1.3.1 Binary Black Hole Detections

The first confirmed detection of gravitational waves (GW150914) occurred on the 14^{th} of September 2015 [27]. Analysis showed the signal arose from the coalescence of a BBH system, with component masses $36^{+5}_{-4} M_\odot$ and $29^{+4}_{-4} M_\odot$. The signal was measured in both aLIGO facilities, with 0.2 s delay between the detections. This

was the first direct measurement of a BBH, the first experimental evidence that black holes can exist in this mass range, the first direct measurement of a black hole and proved the final 'postulate' of Einstein's General Theory of Relativity.

The signal which was measured in both detectors swept up in frequency from 35 Hz to 250 Hz over a period of 0.2 s during the BBH inspiral, and reached a peak strain of 1×10^{-21} with a signal to noise ratio (SNR) of 24 [13]. The sky localisation of the event was inferred from the aLIGO and the time delay between the signal arriving at each detector: its origins could be localised to a region of sky in the southern hemisphere [8]. Its luminosity-distance of 410^{+160}_{-180} Mpc, (1.4 ± 0.6 billion light-years) from Earth was estimated from the amplitude of the signal.

Over the next two years, during the first and second observational cycles of the detectors, seven more detections of similar black hole systems were made by aLIGO [28–31] including the first event detected by both aLIGO and AdV [31] on the 14$^{\text{th}}$ of August 2017. Detecting the same event in three detectors with a wide geographical separation across the Earth's surface significantly improves the resolution with which the sky-location of the source can be calculated.

After several upgrades to both aLIGO and AdV [32–42], the third observational period started in 2019. In the first half of this observational run more than 35 observations from inspiral sources were found with more than 70% of detections registering in both the aLIGO and AdV detectors [12].

1.3.2 Binary Neutron Star Detection

1.3.2.1 Gravitational Wave Observations

On the 17$^{\text{th}}$ of August 2017 (12:41:04 UTC) aLIGO detected a signal sweeping up in frequency lasting $\sim$100 s in the detectors highest sensitivity band (10 Hz - 100 Hz) with the signal being detected first at LIGO Livingston, followed by LIGO Hanford after a three ms delay. The amplitude evolution of the signal over the last 30 s of the event is shown in Fig 1.2. The duration of the measured inspiral and the final measured frequency identified the signal as being produced from the coalescence of a low-mass two-body system.

Analysis of the component masses of the system predicted that one body involved in the inspiral was 90% lower than the smallest black holes previous measured, with the mass of both bodies between $2.73\,M_{\odot}$ and $3.29\,M_{\odot}$, [30]. The total mass of the system before the inspiral was calculated to be below the upper mass limit of neutron stars calculated by Tolman, Oppenheimer and Volkoff in 1939 of $\sim 5\,M_{\odot}$ [43], suggesting that this could be an explanation for the calculated mass of the system. However, this calculation did not provide enough evidence to unequivocally confirm the astrophysical objects' nature of the binary objects.

Plotting the recorded strain data from each detector as a spectrogram allows the signal strength as a function of frequency and time, the evolution of the inspiral, to be viewed [44] (see Fig 1.2). It is shown that the measured frequency of the oscillation passes through the 200 Hz range of detector sensitivity, suggesting a separation of ∼94 km between the two objects. If the lower mass object in the binary pair were to be a white dwarf star, it would be expected to have a radius greater than 1100 km with a lower limit of $\sim 1.4\,M_{\odot}$ on its mass [45]. Should the radius of a white dwarf decrease below this limit, the star would become dynamically unstable. [46].

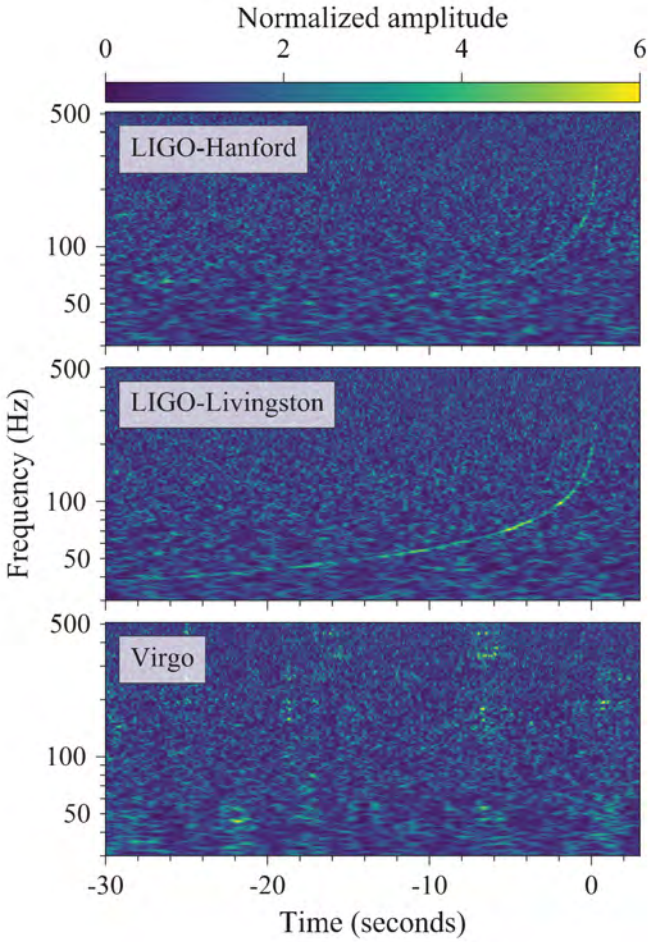

Figure 1.2: Time-Frequency representations of the gravitational-wave event GW170817 observed by the LIGO Hanford(top), LIGO-Livingston (middle), and Virgo (bottom) detectors. Times are shown relative to August 17, 2017 12:41:04 UTC. Figure and headings reproduced from [44].

Approximately two seconds after the gravitational wave signal was detected (12:41:06 UTC), the Fermi Gamma-Ray Burst monitor [47] aboard the Fermi Observatory was triggered, with an event significance 4.6σ [48]. This information then took 14 s to be relayed to the International Gamma-Ray Astrophysics Laboratory (INTEGRAL) [49]. Using the associated times that the signal was detected, the signal's origins were constrained to a region of $28\,\mathrm{deg}^2$ with 90% credibility. The overlap of the confidence regions calculated with gravitational wave observations and gamma-ray observations is shown in Fig 1.3.

1.3.2.2 Electromagnetic Counterpart Observations

After ground-based GWDs and the Fermi-GMB detector had localised the event, these coordinates were announced to the broader astronomical community. Observatories including space and ground-based telescopes ranging in size from 40 cm to 10 m across the entirety of the electromagnetic spectrum, including optical, ultraviolet, near-infrared and long-baseline radio telescopes, were trained on one spot [44]. The One-Meter, Two-Hemisphere team was the first to observe an optical counterpart designated SSS17a in the region outlined by GW170817. This observation produced a list of multiple potential galaxies from which the GW source could have originated. This led to a potential supernova event in Galaxy NGC 4993 being identified/discovered within the predicted bounds of GW170817. The region was then photometrically observed in the infra-red and ultraviolet by the Gemini-South [50, 51] and Hubble observatories [52] producing an estimate of the sources effective temperature over the next two weeks.

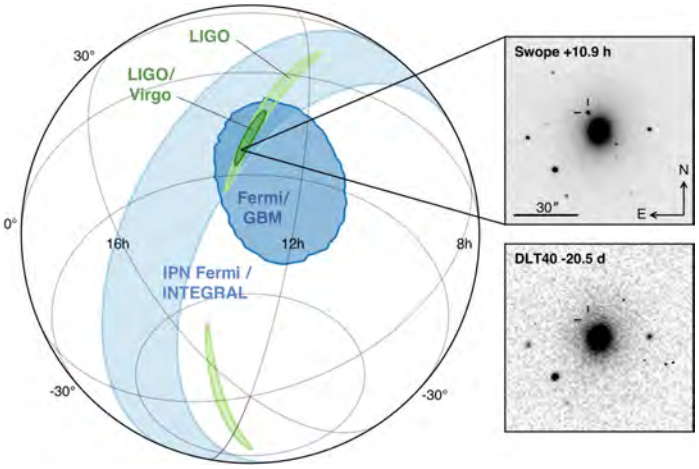

Figure 1.3: Localisation of the gravitational-wave, gamma-ray, and optical signals. The left panel shows an orthographic projection of the 90% credible regions from LIGO (190 deg^2; light green), the initial LIGO-Virgo Localisation (31 deg^2; dark green), IPN triangulation from the time delay between Fermi and INTEGRAL (light blue), and Fermi-GBM (dark blue). The inset shows the location of the apparent host galaxy NGC 4993 in the Swope optical discovery image at 10.9 hours after the merger (top right) and the DLT40 pre-discovery image from 20.5 days before the merger (bottom right). Figure and headings reproduced from [44].

Two days into the measurements, multiple groups [53–56] produced coincident measurements of a rapid decline in the emitted UV-blue spectrum brightness, consistent with rapid cooling, with a seemingly conflicting increase in the object's infrared output. Further inspection of the spectra revealed that none of the

markers, such as traces of ions in the ejecta, usually consistent with a supernova event, could be found. Instead, absorption lines related to the emission of lanthanides were observed [56] which can only be created with higher mass radioactive elements which are not produced in younger/low mass stars [57], ruling out any of the known young supernova remnants in NGC 4993. The observation of the rapid increase in IR and a decrease in UV did not match previous measurements of observed astronomical objects, suggesting that the observed event was similar to computational models of a possible kilonova/macronova. These are only thought to occur after the binary coalescence of two compact stellar objects, such as a black hole-neutron star (BH-NS) or neutron star-neutron star (NS-NS), which were theorised to emit an s-GRB after their coalescence. Considering all of the data produced from all EM and GW sources used to localise and monitor the emissions across the electromagnetic spectrum, the most likely source that could have produced such an astounding astronomical narrative is merging two neutron stars resulting in a kilonova explosion.

1.4 A Brief History of Gravitational Wave Detectors

1.4.0.1 Resonant Bar Detectors

In 1960 Joseph Weber built the first apparatus to measure gravitational wave radiation from a 'double star' or binary star system [58]. Weber produced a long cylindrical mass with multiple piezoelectric transducers placed around each end's circumference. These 'bar detectors' would be excited by a passing gravitational wave, where the excited resonance would persist far longer than the wave that induced it. The bar's vibrational motion was then converted into an electrical signal that could be monitored [59].

Due to the range of resonant frequencies of the bar, the sensitivity range of these detectors was very small. Weber's first bar detector had resonant frequencies which would overlap with GW radiation produced by the Crab Pulsar [60]. Weber

produced two of these detectors, separating them by $\sim 1000\,\mathrm{km}$ in order to reduce the chance that a single terrestrial event could trigger both detectors. Weber reported coincident events in both detectors and claimed that they were produced by gravitational waves [61]. However, the results could not be reproduced by other groups and were deemed invalid [62]. There was growing concern that the sensitivity of these detectors was not high enough to detect gravitational-wave events, with some calculations suggesting that Weber's sensitivity was limited by the thermal noise of the bars [59], reducing their sensitivity by ten orders of magnitude below the calculated requirements [63].

Multiple iterations of bar detector were produced to increase the probability of GW detection. One of the most notable upgrades was implementing cryogenic cooling of the resonant bar [64]. With five cryogenic bar detectors commissioned, the combined strain sensitivity and capability of source localisation were increased. However, no confirmed measurements of gravitational waves were ever produced. Due to the lack of results from these detectors, all resonant bar detectors have since been decommissioned.

1.4.0.2 Interferometric Gravitational Wave Detectors

Gertsenshtein and Putovoit proposed a much more sensitive type of detector [63]. The proposed device was based on a Michelson-Morley interferometer [65] which uses two perpendicular laser beams created from a single optical source (see Fig 1.4). Changes in the interference pattern created by combining the two laser beams on one photo-detector could then monitor the relative position of each end mirror.

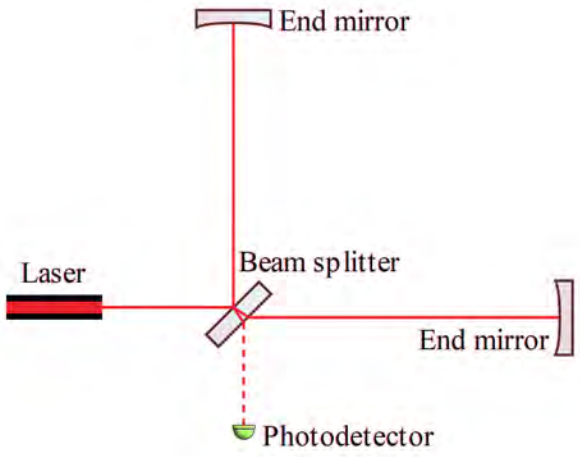

Figure 1.4: Representation of a Michelson-Morely Interferometer. A Single laser source produces monochromatic laser beam equally split by the beam splitter and reflected off each perpendicular end mirror. The interference signal measured at photo-detector.

By holding the resulting interference pattern close to a dark fringe state, any changes in the position of each mirror with respect to the beam splitter produces a change in the phase of the light in one arm with respect to the other, causing the interference pattern to shift from its dark fringe state, producing measurable photons at the output photo-detector[1].

R. Forward of Hughes Aircraft Research (HAR) Laboratories produced the first functioning large scale prototype interferometer in the early 1970s, using a low power ($\sim$45 mW) helium-neon (HeNe) laser and a total arm length of 8.5 m. Around the same time, a second interferometer based on the works of Pirani [66] was also under construction at the Massachusetts Institute of Technology (MIT) [67]. Both devices demonstrated Gertsenshtein's working principle, but their sensitivity was limited by photon shot-noise (see 1.5.3). Several other prototype detectors were developed with the Max Planck Institute for Astrophysics (Garching, Germany) and the now Institute for Gravitational Research (Glasgow, Scotland) both embracing the technology. The Garching design was 30 m in arm length and used a $P = 10$ W, $\lambda = 514$ nm laser [68], while the detector in Glasgow had a total arm length of 10 m [69]. Both detectors had also utilised Fabry-Pérot cavities in the detector arms, effectively increasing each arm's length and the sensitivity of each

[1]The interferometer is locked close to the edge of a dark fringe. A passing GW produces a relative change in the mirrors' position, which is then reversed using a feedback control system.

detector [68, 69].

Through a coalition of the Scottish and German groups, funding was secured to develop and build a 600 m long interferometric detector, GEO-600, in Hannover, Germany. This detector was designed to reach strain sensitivities down to $10^{-22}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at frequencies above 50 Hz [70]. One of the key developments implemented by GEO-600 was to suspend the test mass mirrors, at the end of each detector arm, by silica fibres [71, 72]. Simultaneously, the initial LIGO (iLIGO) development was also underway, where steel cables suspended each test mass. Further funding was secured through the National Science Foundation to build two LIGO sites, one in Hanford, Washington and the other in Livingston, Louisiana, beginning operation in the early 2000s. Fibres made of silica would later be adopted by Advanced LIGO (aLIGO), detectors that began operation in 2015. These detectors, along with GEO600 and a second European detector, Virgo [9], and a cryogenic detector in the Gifu Prefecture of Japan (KAGRA) [73] would become the basis of a now four detectors strong gravitational detector network.

Through each iteration of detector, upgrades and fine-tuning of apparatus, [32–42] has lead to detectors which have broadband sensitivity between 10 Hz - 1000 Hz, with a peak sensitivity several orders of magnitude higher than the first prototype developed at HAR. However, sensitivity is limited by several noise sources related to fundamental physics, each of which has its own set of challenges to be overcome.

1.5 Detector Operation and Limiting Noise Sources

Fig 1.5 shows the optical layout of an Advanced LIGO gravitational wave detector. Following the principle of Michelson's design, a single 125 W laser (1064 nm) is used as the primary means of detecting a change in position of the end test mass optics (ETMs). This differential change in arm length is then monitored as a function of time at the detector output.

Laser light passing into the detector first passes through the input mode cleaner

(IMC) optics which reduce the intensity of stray polarised laser light and filters the laser light's optical modes, transmitting only the Gaussian TM_{00} mode.

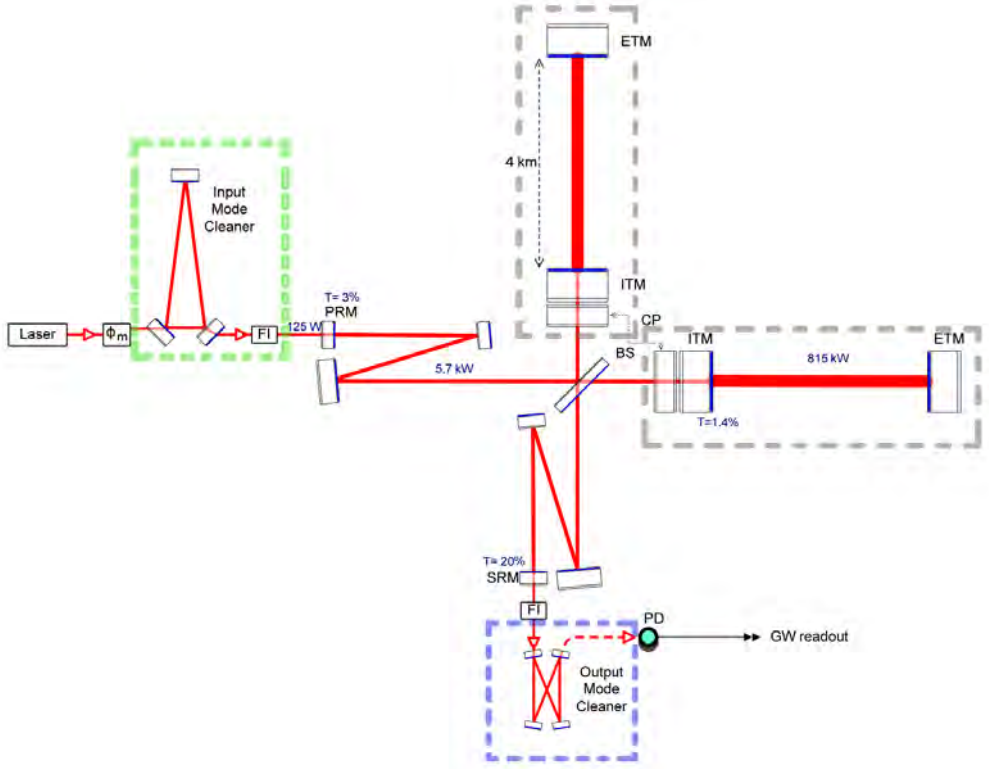

Figure 1.5: Optical layout of a second generation advanced LIGO gravitational wave detector [3]. Fabry-Pérot cavities (grey), input (green) and output mode cleaning optics (blue) are highlighted.

The laser is then incident onto a 50/50 beam splitter mirror and passes into each arm of the detector. Each arm contains a Fabry-Pérot cavity, which effectively lengthens the arms of the detector without the need for a physical length increase of the system. The laser light is continuously built up inside these cavities, reflecting back and forth between the partially transmissive input test mass (ITM) mirrors and end test masses (ETM). In this state, the optical cavity is 'on resonance', and this allows a significant amount of laser power to build up through each consecutive 8 km round trip. A small portion of the laser light is transmitted through each ITM and recombined at the beam splitter. The recombined beams then pass through a final set of output mode cleaner filters, removing any higher-order optical modes . This greatly reduces the power that must be detected at the output [74].

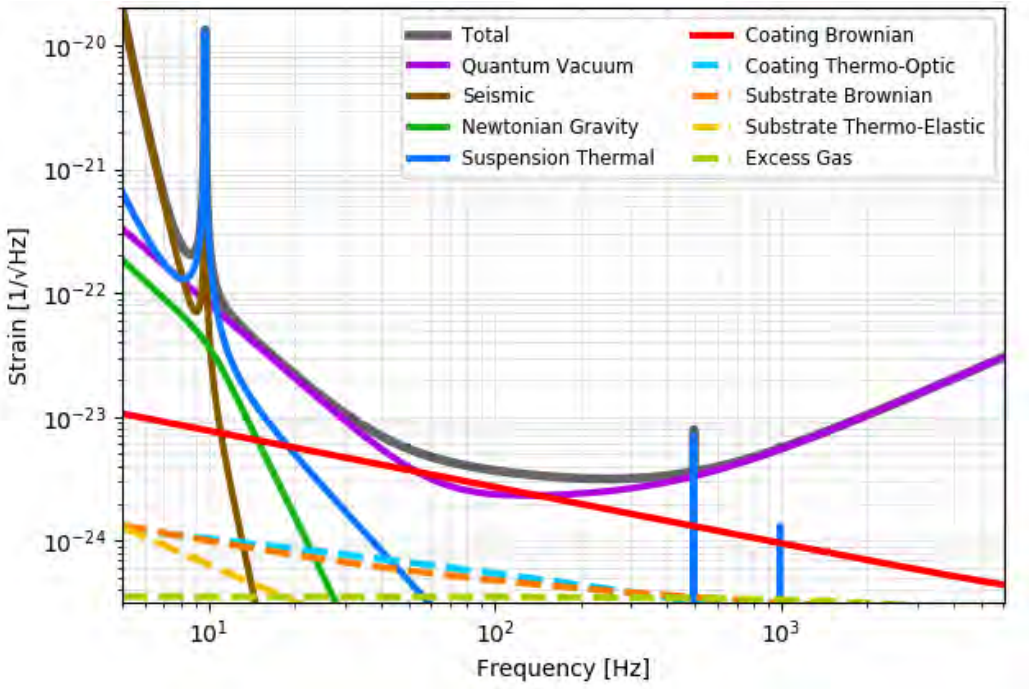

Figure 1.6: Calculated contributions of multiple noise sources in an aLIGO GWD. The total noise budget from each noise source is calculated using GWINC [75]

Each component of the detector must be optimised such that its sensitivity to astronomical events is as high as possible. Fig 1.6 shows the estimated level of noise produced by different processes or systems in the aLIGO detector. In the following subsections, noise sources that place the greatest restrictions on strain sensitivity will be discussed.

1.5.1 Gravity Gradient Noise

Gravity gradient noise refers to changes in the local gravitational field around the detector, coupling directly into the position of the detectors' mirrors through the force of gravity. This effect produces significant noise at $\leqslant$ 20 Hz [76] inhibiting GW detection below $\approx$1 Hz. The largest contribution comes from seismic waves travelling across the surface of the Earth. These seismic waves produce density fluctuations in the Earth around the detector, exerting a fluctuating gravitational force on the test mass. Placing multiple seismometers at different points around the detector can be used to measure the effects of surface waves, allowing some of their effects to be counteracted/removed from the interferometer [77].

As the energy of a surface wave will decrease exponentially with the amount of material which the wave has travelled through, building future detectors 100 - 200 m underground [10, 78] can reduce the effect of these surface seismic waves.

1.5.2 Seismic Noise

Seismic noise is produced from vibrations on the Earths' surface, with notable contributions from human activity around the detector. To mitigate vibrational coupling from these sources, multiple levels of active and passive isolation are employed. In aLIGO, this isolation's lowermost level is known as Hydraulic External Pre Isolator (HEPI) stages [79]. These consist of four corner stations on each vacuum tank locked to the ground's motion. Using hydraulic actuators to compensate for ground acceleration in 6 degrees of freedom provides low-frequency isolation between 0.1 Hz and 10 Hz. The second stage is then locked to the motion of the first, using a feed-forward loop to sense the motion of HEPI and reduce the motion of the second stage to essentially zero.

Inside the ETM and ITM vacuum chambers any residual motion is passively damped through four stages of blade springs and pendulums. Figure 1.7 shows a model of this suspension system. At frequencies above each pendulum's resonant frequency, the transfer function drops as $1/f^2$, isolating the mirror from the motion of its suspension point.

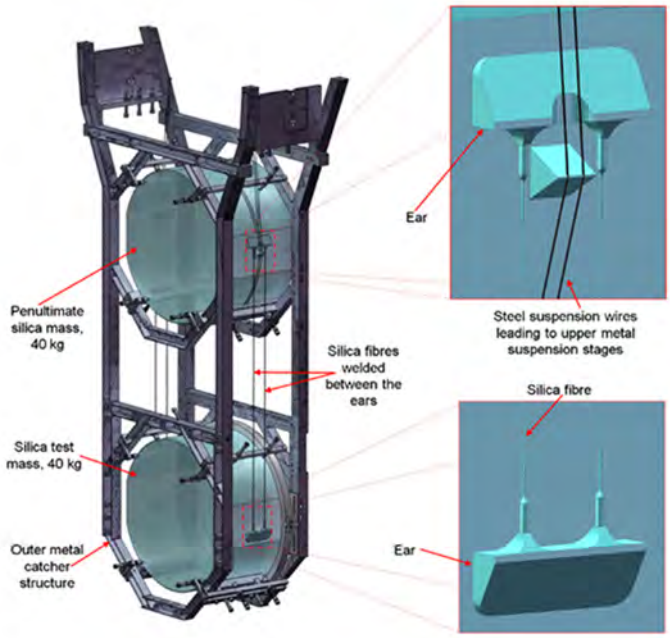

Figure 1.7: Rendering of monolithic silica suspensions used in initial LIGO and aLIGO detectors - not to scale [80].

1.5.3 Quantum Noise

Quantum noise is associated with laser light in a GW detector and consists of; shot noise, limiting sensitivity at high frequencies, and radiation pressure noise, limiting sensitivity at low frequencies. This section gives a classical description of shot noise and radiation pressure noise.

1.5.3.1 Shot Noise

Shot noise arises from the statistical counting of photons incident on the photo-detector at the interferometer output. As with any statistical counting measurement, there is an associated error proportional to the square root of the number of counts. To reduce this uncertainty, higher laser power (resulting in more photo-electrons being produced) should be used. The sensitivity limit from shot noise is given by [81]:

$$h = \frac{1}{L}\left(\frac{\mu c \lambda}{4\pi^2 P}\right)^{\frac{1}{2}}, \tag{1.4}$$

where h is the measured strain, L is the interferometer arm length, μ is Plank's constant and P, and λ are the power and wavelength of the laser light propagating at the speed of light in a vacuum, c. In an aLIGO style detector with 125 W of input laser power, $h = 2.95 \times 10^{-24}$ at 100 Hz would be expected. As the effects of shot noise scales with increasing frequency, the laser power inside the detector is increased, reducing the counting error and increasing sensitivity to higher frequency gravitational wave signals. However, increasing the power inside the detector increases another source of quantum noise, known as radiation pressure noise [82].

1.5.3.2 Radiation Pressure Noise

Radiation pressure noise is caused by photons in the detector arm imparting momentum ($2\hbar$k) to the mirrors, where k is the photon wave number. The differential displacements created by this effect place a lower limit on the detector's achievable strain sensitivity. By increasing the mass, m, of each reflecting optic, the effect of photon momentum transfer can be reduced, as can be seen in the following formula:-

$$h = \frac{N}{Lmf^2}\left(\frac{\mu P}{4\pi^4 \lambda c}\right)^{\frac{1}{2}}, \tag{1.5}$$

where f is frequency, N is the total number of reflections on each optic [82]. All other terms are defined in Eq. 1.4.

As shot noise is inversely proportional to the square root of the incident laser power, and radiation pressure noise is proportional to the square root of the incident laser power, a situation can be achieved for each frequency where the shot noise and radiation pressure noise curves intersect at their lowest possible values. Under these conditions, the detector's level of quantum noise is reduced and reaches the standard quantum limit (SQL) [83]. This limit stems from the Heisenberg Uncertainty Principle concerning momentum and position. Optical techniques have been developed which can surpass this limit, such as optical squeezing [84] (see section 1.6.0.2) and signal recycling cavities which allow the effects of shot noise and

radiation pressure noise to be correlated.

It should be noted that a formal quantum mechanical treatment of a GW detector shows that the source of quantum noise is zero-point fluctuations in the vacuum electromagnetic field entering the interferometer at the output port [59, 82, 85]. Fluctuations in this field's amplitude result in radiation pressure noise, while phase fluctuations result in shot noise. This is discussed in more detail in section 1.6.0.2.

1.5.4 Thermal Noise

Thermal noise arises from thermally-induced vibrations of the atoms and molecules in the detector mirrors, and their suspensions. These vibrations can excite each test mass's mechanical eigenmodes, which interacts with the primary sensing laser. The thermal noise contributions of each material are linked to the rate at which vibrational energy can be dissipated, the material's mechanical loss.

In the current generation of GWD's, fused silica was chosen as the bulk material for room temperature test masses and suspension fibres [86]. It has been shown to have low internal friction (mechanical loss) and low optical absorption. To increase the mirrors' reflectivity, alternating high and low refractive index coating layers are deposited onto the optic's front (laser incident) face. The number of layers and their material composition determines the reflectivity of the mirror. The thermal noise produced by these coating materials, coating thermal noise (CTN), limits the peak sensitivity of current detectors.

1.6 Interferometric Techniques

A gravitational wave passing through a GWD changes the relative length of each of the detector arms. As each arm contains a Fabry-Pérot cavity, the sensitivity to changes in the phase of the light-induced by a passing GW is dependent on the length of the cavity and the storage time τ_s of the laser light. The maximum sensitivity to such a phase change is achieved when τ_s is equal to half the period of the gravitational wave [87]. Thus, the sensitivity of a GWD to GW sources at different frequencies can be changed by changing the lengths of these cavities, i.e. a

much longer cavity, increases τ_s and lowers the sensitive frequency range of the detector. To produce the infrastructure to make a GWD with detector arms spanning over much longer distances is impractical and would ultimately be limited by the curvature of the Earth over long distances [88]. A more practical approach is to create a system where the laser light's propagation distance is far greater than the physical arm cavity length.

One approach to achieve artificial lengthening is to use a delay-line configuration where the ETM mirror is placed at a small angle to the beam path creating what is known as a 'folded cavity' [89] (see Fig 1.8(a)). This technique is employed at the GEO-600 detector [90].

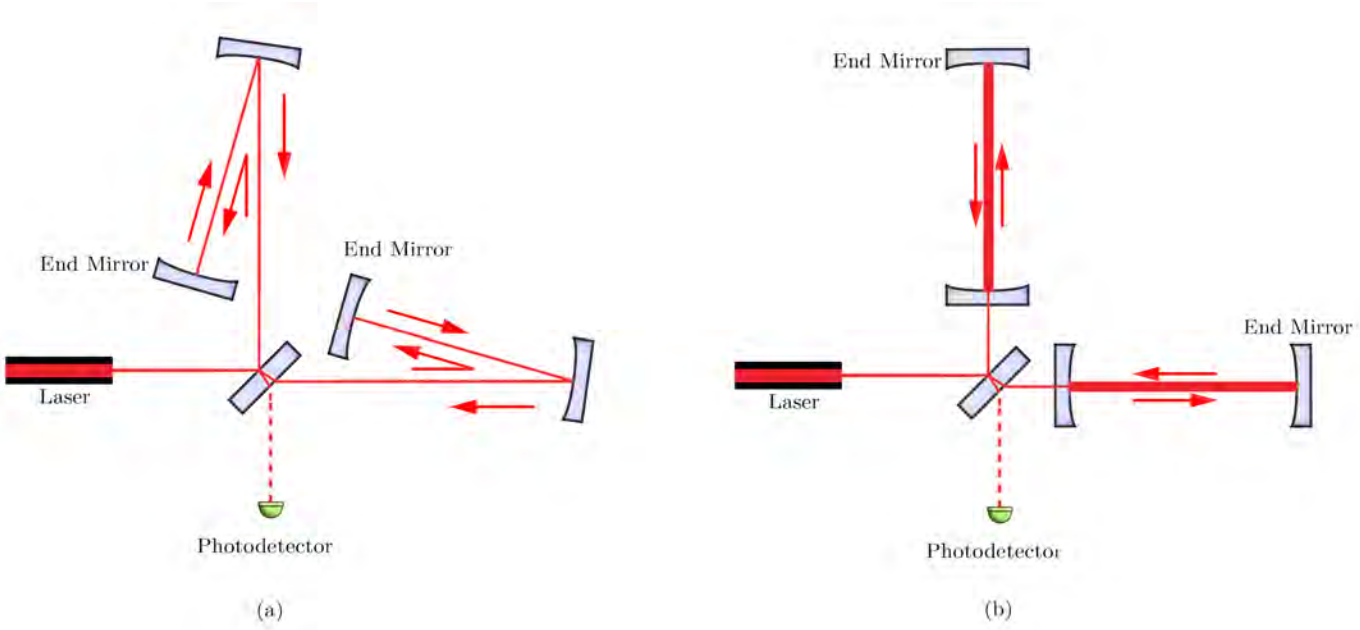

Figure 1.8: Test mass mirror configuration in (a) 'folded cavity' interferometer and Fabry-Pérot cavity (b) configurations [89]. The red arrows denote the laser propagation direction.

A second approach to increasing arm length while constraining the mirror angle is to use Fabry-Pérot cavities (see Fig 1.8(b)). The use of cavities in this manner allows laser light to travel over the same distance over multiple reflections without increasing the physical length of the detector arms. Over a specific time, known

as the storage time of the cavity[2], a small portion of the resonating laser power is transmitted back out of the cavity and can be measured. Over each bounce, the phase difference between the light injected into the cavity and the light returned from it is enhanced, producing a greater sensitivity to the effects of a passing gravitational wave.

1.6.0.1 Power and Signal Recycling Cavities

At the detector output, the recombined laser light from both Fabry-Pérot cavities are held close to a dark fringe. In doing so this increases sensitivity to passing gravitational waves, but also squanders the majority of laser light produced by each cavity. This loss of light can be remedied by creating another cavity within the detector. Placing a mirror before the 50/50 beam splitter (see Fig 1.9), known as the power recycling mirror (PRM) can reflect light back into the Fabry-Pérot cavities, increasing the total incident power which linearly scales the interferomter output signal.

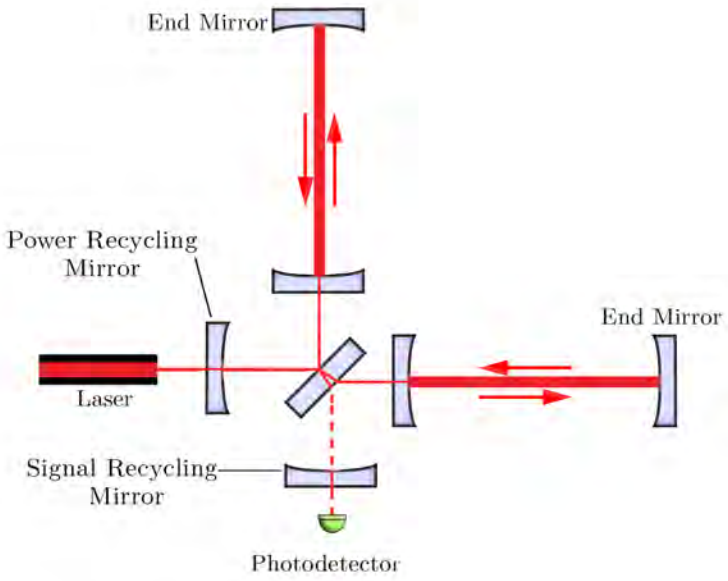

Figure 1.9: Positions of signal and power recycling mirrors in a Fabry-Pérot cavity based interferometer.

When the interferometer is in operation, the light incident on the photo-detector

[2]The storage time τ_s is determined by the separation and properties of each mirror [89]

is reduced. The carrier frequency of the laser light and its side-bands are contained within the cavities. An additional signal recycling mirror (SRM) between the interferometer output and the photo-detector forms an additional cavity. In this 'dual recycling' configuration, the interferometer can be thought of as two separate cavities: one which contains both arms and the PRM, and the second using the SRM and detector arms. The main carrier light of the laser is recycled using the PRM, while its side-bands (the imaginary components of the main carrier) can resonate between the arms and SRM [87]. A passing gravitational wave changes the relative length of each detector arm, increasing the amplitude of the side-band. By placing the SRM at different distances from the beam splitter, the frequency of the resonating side bands can be changed, determining the frequency of light that reaches the photo-detector. Thus, the interferometer's frequency sensitivity can be tuned [87, 89]. By changing the distance between the SRM and the rest of the interferometer, the detector's sensitivity can be shifted to different parts of the frequency spectrum, i.e. moving the detector's sensitivity to lower values of f to mitigate the effects of radiation pressure noise.

1.6.0.2 Squeezed Light

In Eq. 1.4, and 1.5, a classical description of shot noise and radiation pressure are given, where momentum transferred from photons described as particles. In this regime, the standard quantum limit, the points at which laser power is optimised to reduce shot noise and radiation pressure, is an unsurpassable boundary. While a classical coherent electromagnetic wave has a well-defined amplitude and phase, this is not true for a quantised coherent field. There is then uncertainty in the amplitude and the phase, independent of each other [91]. By manipulating the light field to lower the uncertainty on the amplitude, this inadvertently increases the fluctuations in the field's phase.

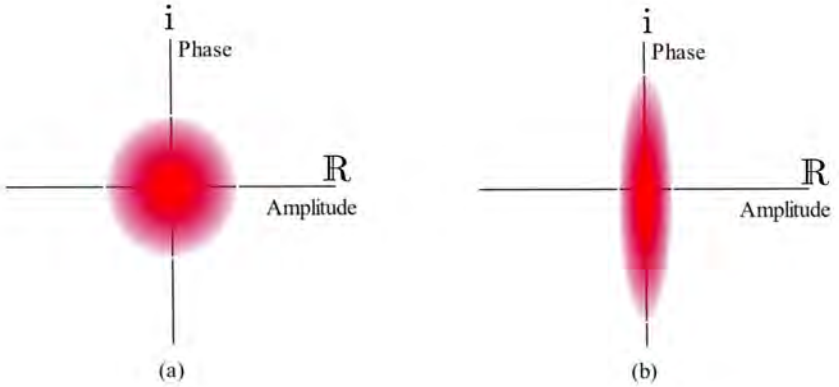

Figure 1.10: Quadrature representation of the probability and uncertainty in a classical coherent light field (a) compared to a coherent quantised field after amplitude squeezing (b). Higher transparency indicates a lower probability. Larger magnitude on a given axis shows a higher uncertainty.

By manipulating the light field, uncertainties in the amplitude or the phase of the quantised field can be reduced below the level of uncertainty produced by vacuum fluctuations. Fig 1.10(a) shows a comparison between the probability and uncertainty of an un-squeezed light field, analogous to a classical quantised coherent field. Fig 1.10(b) shows the same light field after manipulation of squeezed states to reduce its amplitude uncertainty.

Injecting squeezed light states into a steady-state interferometer at the 'dark port', the detector's output, the sensitivity of the detector to the amplitude of the gravitational wave strain at low frequencies can be improved. Manipulating the field so that the phase uncertainty is lowered increases the detector's sensitivity at higher frequencies.

1.7 Current Status of Gravitational Wave detectors

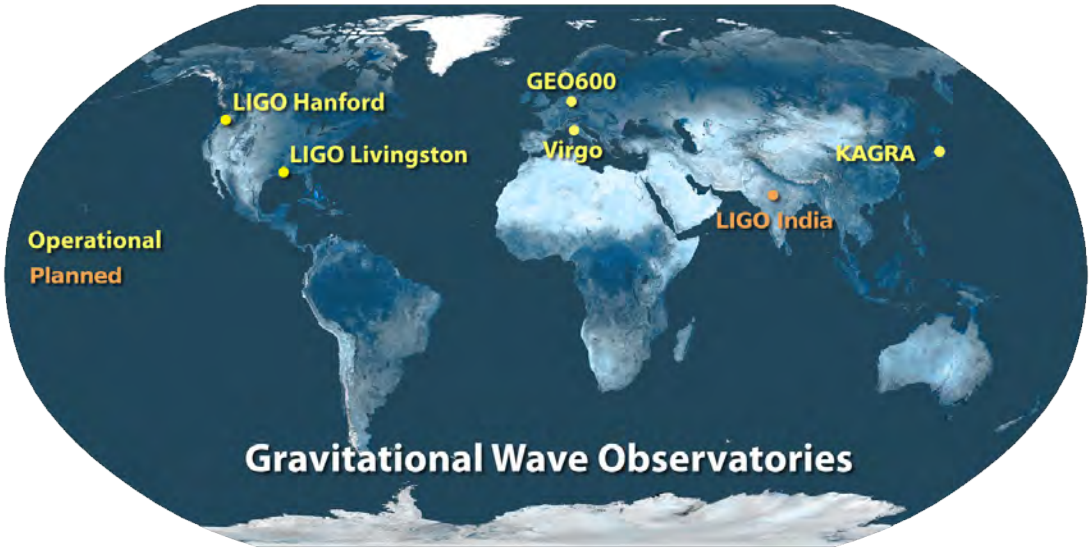

Figure 1.11: Gravitational-wave observatories across the globe. In addition to the twin LIGO detectors in Hanford, Washington and Livingston, Louisiana, the Virgo detector located in Italy, current (as of 2021) operating facilities in the global network include GEO600 in Germany and the Kamioka Gravitational-wave Detector (KAGRA) in Japan. A sixth observatory is being planned in India.

Since the developments of the German/Scottish collaborative venture of GEO600 and Caltech based iLIGO detectors, the number of operating gravitational wave observatories has grown. There are currently five long-baseline GW detectors in the current detector network, aLIGO Hanford, aLIGO Livingston, AdV, KAGRA and GEO600. Fig 1.11 shows the locations of each detector currently in operation, with the proposed LIGO India detector also highlighted. With each new active detector and upgrades to current installations the sensitivity and source localisation precision of the gravitational wave detector network increases [92].

Fig 1.14 is the calculated strain sensitivity of current and planned gravitational wave detector as a function of frequency. The differences between each detector will be discussed in the following sections.

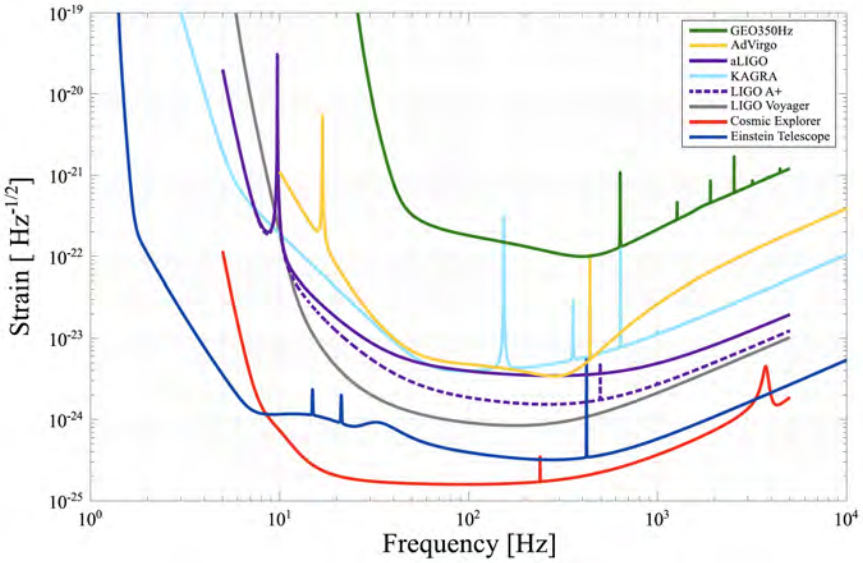

Figure 1.12: Calculated Strain sensitivities of current and planned ground based gravitational wave detectors [93].

1.7.1 Advanced LIGO (aLIGO)

Advanced LIGO, run by MIT and Caltech, are two gravitational wave detectors in the United States of America. LIGO Livingston (Louisiana) and LIGO Hanford (Washington) are designed such that their optics and components are as similar as possible. Doing so reduces the number of possible differences in each detector, making noise reduction a joint effort. Since their initial construction in the early 1990s, the iLIGO detectors have undergone multiple upgrades.

The aLIGO detectors are made up of two 4 km-long Fabry-Pérot cavities in a perpendicular configuration. The current maximum input laser power to each cavity is 125 W for each detector, which translates to a total circulating cavity power of 750 kW [94]. Each test mass mirror (40 kg) is produced from high grade fused silica, 34 cm in diameter, and monolithically suspended from fused silica suspension fibres bonded to 'ears' on the side of each test mass [95–97]. To increase the reflectivity and lower the effects of laser heating, each test mass is coated in a highly reflecting mirror stack with low optical absorption, composed from layers of titania doped tantala ($Ti:Ta_2O_5$) and silica glass. While significant progress has been made in improving each detector's stability and sensitivity (a factor 10 improvement over the inital LIGO

detectors), the aLIGO sites have not yet reached their design sensitivity.

(a) (b)

Figure 1.13: Images of advanced LIGO detectors in Hanford, Washington [98](a) and Livingston, Louisiana [99](b).

In the most recent observational run, both detectors were operational for more than 150 days, with LIGO Hanford and LIGO Livingston maintaining observational status for 71.2% and 75.8% of the run. An average inspiral detection range of 113 Mpc (Hanford) and 134 Mpc (Livingston), was achieved, contributing to measurements of over forty gravitational wave events [12].

1.7.2 Advanced Virgo (AdV)

Figure 1.14: Image of advanced Virgo detector Cascina, Pisa.

Advanced Virgo is the third iteration of 3 km long gravitational wave detector in Cascina, Italy. Like aLIGO, the Virgo detector, run in collaboration with institutions from France, Germany, Hungary, Poland and The Netherlands, has been through multiple upgrades to reach its current sensitivity. These upgrades include fused silica suspension fibres in test mass suspensions, a thermal compensation system for test mass optics (see chapter 3), and vacuum system improvements to mitigate the effects of residual gas noise.

Comparing the calculated sensitivity of AdV to aLIGO, there are clear differences in the functional form of the noise. One of the key contributors to the detector's sensitivity at low frequencies is AdV's unique seismic isolation system [100]. The vibration isolation system, known as the Virgo *Superattenuator* is composed of multiple stages of passive isolation, which can suppress seismic effects down to frequencies of 4 Hz [100]. The detector's sensitivity in this frequency range surpasses the active isolation system used in aLIGO detectors [72, 79].

AdV has contributed to the measurements of gravitational wave events since it joined the second observational run (O2) in 2017, contributing to the detection and localisation of GW170817. Since then, the detector has been heavily involved in the first half of the third observing run (O3a). It maintained observational status for several days more than either aLIGO detector, contributing to more than 70% of gravitational wave events.

1.7.3 GEO600

GEO600, the first gravitational wave detector to employ monolithic silica suspended test masses, is a collaborative effort between the Max Planck Institute for Astrophysics (Garching, Germany) and the Institute for Gravitational Research (Glasgow, Scotland). Construction of the detector and its 600 m long (1200 m folded cavity) arm started in the 1990s, with its completion shortly after the millennium.

Due to its folded arm design, the strain sensitivity of GEO600 is no longer competitive with advanced detectors such as aLIGO or AdV. However, many of the technologies now employed at these detectors have stemmed from techniques first developed in GEO600. This detector was the first to develop and implement monolithic silica fibre suspensions and injection of squeezed light into the interferometer.

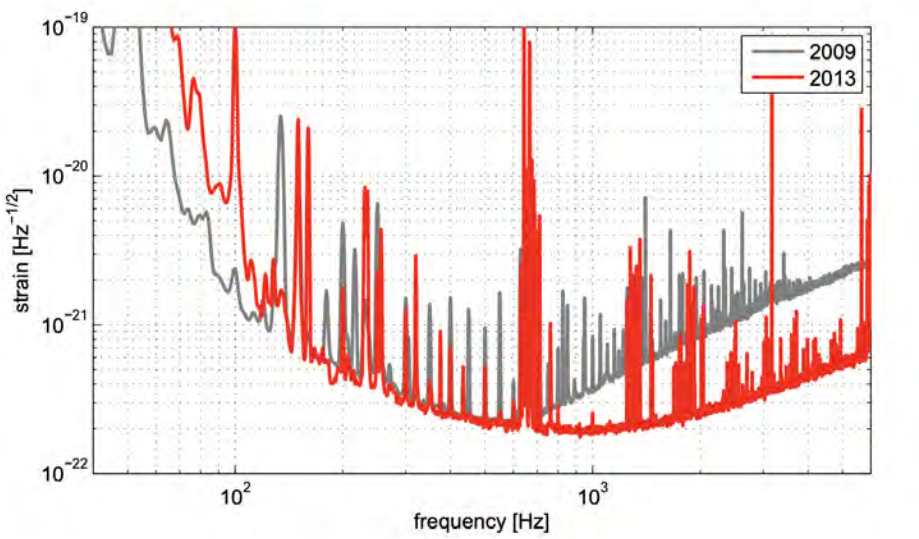

Figure 1.15: GEO600 Strain sensitivity between 40 Hz - 1.5 kHz (grey) compared to upgraded strain sensitivity of GEO600-HF (red). Figure reproduced from [101].

As one of the first detectors to implement squeezed light [102], the focus of GEO600 has pivoted to increasing its strain sensitivity at frequencies where quantum shot noise currently dominates ($f > 600$ Hz) [101]. Fig 1.15 is a comparison the strain sensitivity of GEO600 in 2009 before several focused upgrades to its high frequency capabilities, re-branding as GEO600-HF [103].

1.7.4 Kamioka Gravitational Wave Detector (KAGRA)

The Kamioka Gravitational Wave Detector (KAGRA) is a 3 km long sapphire based interferometer that started construction in early 2012. Situated near Hida, Japan, the GW detector is the first detector of its kind, utilising active cryogenic cooling of its optical components. The detector itself is under construction 300 m below surface level, inside the Kamioka Mine to reduce the effects of seismic and gravity gradient on its low-frequency sensitivity [73].

By operating the detector at cryogenic temperatures (∼20 K), thermal noise effects from coating and substrate materials are reduced. However, to produce comparable levels of strain sensitivity to the advanced detectors, (aLIGO, AdV), all major optical components, including test masses and suspension fibres, are composed of crystalline sapphire, due to its better thermal noise performance at these temperatures compared to fused silica. As KAGRA is the first detector to attempt cryogenic operation, the techniques and methods developed here will be of great interest to third-generation gravitational wave detectors, which also aim to operate in the cryogenic regime [78, 104].

KAGRA is currently able to reach a sensitivity distance of 594 kpc, a fraction of current advanced detectors' sensitivity. During O3a, KAGRA could not maintain stability for prolonged periods and could not contribute to any GW detections. As the detector's full design sensitivity has not yet been reached, it is expected that KAGRA will make meaningful contributions to GW detections in the coming years.

1.7.5 LIGO India

Through collaboration between LIGO Labs at Caltech and the Indian Department of Atomic Energy (DAE), and the Department of Science and Technology (DST), plans to construct a gravitational wave detection in India have been produced. Using schematics, optical and electronic components previously used in LIGO Hanford, an identical 4 km long detector has been designed. Currently, multiple sites for the construction of LIGO India are still being assessed. Through the use of another GWD in the network of functioning detectors, localisation of GW sources then becomes much more accurate. [92].

1.8 Third Generation Detectors

Current gravitational wave detectors such as aLIGO are classed as second-generation detectors. As multiple upgrades have been implemented at each detector, the noise sources which limit their capabilities are becoming limited by the fundamental physical laws which allow them to function. To improve their sensitivities beyond their current sensitivity, techniques must be employed to circumvent these limits. In this section, there will be descriptions of multiple proposed third-generation gravitational wave detectors, detailing the noise lowering techniques which are to be employed in each case.

1.8.0.1 LIGO Voyager

Based on the current aLIGO detectors design, LIGO voyager is a project that aims to improve upon current detector strain sensitivity by more than a factor of 4 [104]. Currently, there are plans to operate the detector at cryogenic temperatures ($\sim 123\,\mathrm{K}$) using some of the techniques first tested in KAGRA. It will use crystalline silicon as an optical mirror substrate with a primary interferometer laser wavelength of 2000 nm. Doing so reduces thermal noise from the 200 kg optical substrates [104], with ongoing research into low optical absorption, HR mirror coatings which can operate successfully under these conditions. By injecting squeezed light states into the detector, using techniques first implemented by GEO600-HF, the detector's high-frequency sensitivity will also be significantly improved compared to current second-generation detectors [104].

1.8.0.2 Einstein Telescope

The Einstein Telescope (ET) is a proposed European detector, comprised of six independent 10 km interferometers planning to be built between 100-300 m below the ground. Half of the interferometers will be optimised for a different bandwidth of frequency space, and orientated in a xylophone configuration, shown in Figure 1.16. Each of the nested interferometers aiming to increase on current second-generation sensitivities by a factor of 100. Such a configuration of interferometers, previously suggested for aLIGO, allows for two separate detectors whose sensitivities are optimised for high frequency (HF) 100 Hz to 10 kHz and low frequency (LF) between 2 Hz and 100 Hz.

To gain such a high resolution over a large frequency range, three detectors will focus on low-frequency sensitivity and the others, high frequency.

- As a low power main interferometer beam is being used in the LF interferometer to avoid excess heating in a cryogenic system, it also has the advantage of reducing any incident radiation pressure on the sensing optics. Silicon test will be implemented in a cryogenically cooled system (<20 K) with a 1550 nm, 18 kW main interferometer beam. Silicon is used in place of traditional silica optics in this case as when cooled, silica exhibits unavoidably large mechanical loss peaks at $\sim$ 20 K and $\sim$ 250 K [105] which is not present in silicon optics at these temperatures, reducing substrate thermal noise.

- High-frequency sensitivity will be gained through a room temperature system, where the effects of the cryogenic loss peak in silica is no longer a limit, and its low mechanical and optical properties can be taken advantage of. Each test mass optic's size and mass increased to 200 kg fused silica, 60 cm in diameter. Each Fabry-Pérot cavities will be designed to produce 3 MW of circulating optical power.

Plans of complimenting this with frequency-dependent light squeezing methods currently being used at GEO-600 [90] reduce the associated increase in radiation pressure shot noise indicative of such an increase in laser power. With such a detector it is projected that an overall strain-sensitivity of 2×10^{-24} m $(\sqrt{\text{Hz}})^{-1}$ at 100 Hz will be achieved [78].

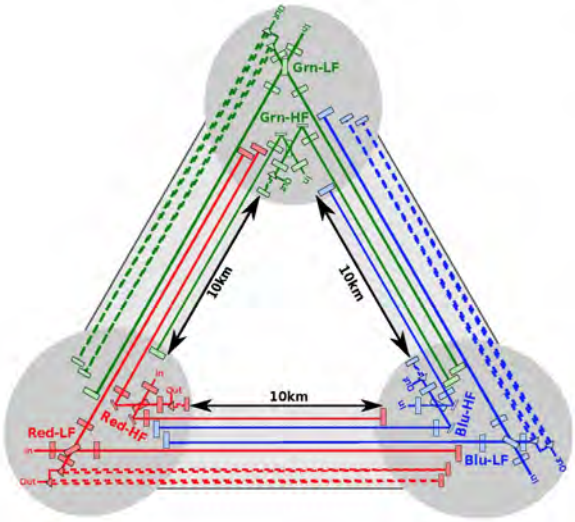

Figure 1.16: A schematic layout of the proposed Einstein Telescope GWD comprised of six independent interferometers (3× LF, 3× HF) optimised for low frequency and high frequency gravitational wave detection. [78].

1.8.0.3 Cosmic Explorer

Cosmic Explorer is the name given to the next iteration of US ground-based gravitational-wave detectors, focusing on sensitivity towards high mass Compact Binary Coalescences (CBC). These events are thought to produce a measurable signal in the LF band of a current detector sensitivity (2 Hz - 100 Hz). To increase detector sensitivity in this frequency range, the proposed length of the detector arms is 40 km, with other shorter iterations being considered [106]. Such a detector would use technology similar to what is currently implemented in aLIGO, but with added frequency dependant squeezing. The increased detector length is calculated to improve sensitivity by over an order of magnitude compared to current 2nd generation detectors [107]. After this phase of its operation, there are also plans to operate Cosmic Explorer at cryogenic temperatures.

1.8.0.4 Neutron Star Extreme Matter Observatory (NEMO)

The Neutron Star Extreme Matter Observatory (NEMO) concept design has been proposed to probe the underlying nuclear physics that occurs during two neutron stars' coalescence. In order to produce useful observations of such events, the sensitivity of the detector is directed towards low noise at high frequencies

(f>1000 Hz) [108]. By increasing the sensitive frequency band of NEMO, more information on the intra and post-merger of a BNS can be produced.

Like other planned third-generation detectors such as ET-LF and LIGO Voyager, NEMO plans to take advantage of cryogenically-cooled silicon test masses to reduce thermal noise effects. A 500 W, 2000 nm seed laser will be used with a typical Fabry-Pérot cavity arm configuration to reach 4.5MW of stored laser light power.

1.9 Space Based Detectors

1.9.0.1 Laser Interferometer Space Antenna (LISA)

The Laser Interferometer Space Antenna (LISA) is a joint venture by the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA) to launch three spacecraft capable to the third Lagrange point (L3) in the Earths' orbit around the Sun. The LISA interferometer consists of three spacecraft placed in a heliocentric orbit. These craft will be separated by 2.5×10^6 km[3]. These three spacecraft will be capable of sending and receiving 2 W, 1064 nm optical signals from the optical benches and test masses aboard each spacecraft, creating a three-point interferometer. With each interferometer achieving arm lengths inconceivable for ground-based detectors, the spacecraft strain sensitivity will be unparalleled at low frequencies (f <1 Hz).

Space-based gravitational wave detectors such as LISA and others [110, 111], are not subject to any gravity gradient noise created by seismic motion, which is the limiting factor on low-frequency gravitational wave measurements for detectors on Earth, gaining sensitivity to sources in the audio and sub-audio bands. This low-frequency sensitivity could allow inspiral sources to be measured in ground-based detectors to be detected much earlier in their coalescing orbits. This would allow LISA to be used as an early warning trigger for ground-based detectors. The exceptional low-frequency performance would also allow for measurements of continuous gravitational waves, which have remained unmeasured [19, 20], ultra-compact binaries in our galaxy, super-massive black hole mergers, and extreme mass ratio inspirals [109].

[3]Current Launch date of LISA is estimated early 2030s [109]

1.10 Conclusions

Since the initial theory of the existence of gravitational waves, multiple scientific instruments have been produced allowing more than 40 confident gravitational wave signals to be detected from inspiral sources. The sensitivities required for these distortions in space-time require advanced optical techniques, which have been implemented into a Michelson based design. The chambers which contain the test mass optics often referred to as the quietest rooms in the world, have been well isolated from surrounding environmental effects allowing their motion to be detected to an accuracy of $<2\times10^{-20}$ m. The sources of noise that now limit these detectors' sensitivity now stem from the underlying physical laws that facilitate their function. However, multiple upgrades and planned gravitational wave detectors continue to push the boundaries on sensitivity.

Since the development of the first gravitational waves detector, a global collaboration containing more than 100 academic institutions have created a network of advanced gravitational wave detectors. We are now truly in the era of gravitational-wave astronomy. Upgrades to GWD installations have been central in the groundbreaking measurements of gravitational waves leading to new scientific insights into black holes, gamma-ray bursts and kilonovae, fundamentally changing our understanding of the universe. To further our understanding of different types of GW sources that have not yet been measured, there must be improvements to multiple limiting sources of noise that limit the sensitivity of current gravitational wave detectors.

Chapter 2

Coating Thermal Noise

On a molecular level, any material above absolute zero has $\frac{1}{2}k_BT$ of thermal motion associated with every vibrational degree of freedom. In a gravitational wave detector, thermal energy can result in fluctuations of the measured position of the test mass, limiting sensitivity. This sensitivity limit is known a 'thermal noise'. Thermal noise can arise from thermal energy inducing vibrations (Brownian thermal noise) and from temperature fluctuations changing the optical path length in the detector via the thermal expansion coefficient or refractive index of the material.

2.1 Brownian noise

If one were to imagine a macroscopic material composed of molecular scale strings, then plucking of one of these strings would cause it to vibrate. The more energy available to the string (analogous to a higher temperature (T) in the system), the higher amplitude the motion would be. These small molecular vibrations then couple into the macroscopic mechanical resonances of the material. This analogy likens the plucked movement of a string to the random motion of fundamental particles in a solid at a temperature $\mathrm{T} > 0\,\mathrm{K}$. This random motion of particles was first observed by R. Brown in 1828 [112] when he observed the motion of pollen particles suspended in water. Upon first observation, the movement of the pollen particles was described as chaotic, and no explanation for the motion of the particles could be produced. This would not be mathematically explained until 1901 when A. Einstein published his understanding of the phenomena [113]. It was noted that the motion of the particles

was due to their interactions with the surrounding water molecules. Einstein noted that the motion of these pollen particles depended on the total thermal energy of the system. This theoretical prediction was then experimentally verified by J.Perrin in 1908 [114], who received the Nobel Prize in Physics for his work.

2.1.1 The Fluctuation-Dissipation Theorem

The pollen particles moving through their suspension liquid is an example of a random, thermally driven dissipative process. Callen published a series of papers [115] in which the thermodynamics of an irreversible dissipative process was explained using the Fluctuation Dissipation Theorem. This theorem states that any system in thermodynamic equilibrium will exhibit random thermally-induced fluctuations. The magnitude of the fluctuations is related to the real component of the impedance of the system (which measures the dissipation of energy in the system) $\mathbb{R}[Y(\omega)]$. The mobility of a given system can be expressed as the ratio of the force applied during the reaction *F(ω)* and the induced velocity $\mathrm{v}(\omega)$ where ω is the angular frequency

$$Y(\omega) \equiv \left(\frac{F(\omega)}{\mathrm{v}(\omega)}\right)^{-1}. \tag{2.1}$$

Callen showed that in order to predict the total thermal noise which would be produced by a given system as a function of frequency, Eq. 2.1 the mechanical admittance of a given system is expressed by scaling Eq. 2.1 by $k_{\mathrm{B}}T$, where k_{B} is Boltzmann's constant, and T is the temperature of the system and by $\frac{1}{\omega^2}$, creating a power spectral density $S_x(\omega)$ of thermal noise

$$S_x(\omega) = \frac{4k_{\mathrm{B}}T}{\omega^2}\mathbb{R}[Y(\omega)]. \tag{2.2}$$

In a gravitational wave detector, thermal energy can excite the pendulum mode of the optic and resonances of its suspensions resulting in 'suspension thermal noise'. Internal modes of the mirror result in 'substrate thermal noise'. However, this chapter will be mainly concerned with thermal noise arising from the mirror coatings.

2.1.2 Mechanical Dissipation

Thermal noise in a gravitational wave detector is associated with the magnitude of mechanical damping in the suspended mirror system as discussed in Eq 2.1. Mechanical dissipation (damping) can arise from internal sources and external phenomena. All sources of external damping must be suitably minimised. For example, to avoid damping of the pendulum via interactions with gas molecules, the detector is operated at pressures $< 2.6 \times 10^{-8}$ mbar [116]. When all sources of external damping are minimised, the remaining dissipation arises from internal processes with the material itself.

2.1.2.1 Internal Dissipation

A purely elastic body, acted on by external stress σ will experience a linear restoring force opposing the acting force. The total stress and strain are related by Hooke's Law, which describes the restoring force produced by a coiled spring, stretched over a given distance [117]. The total strain ε on such a body, can then be defined by the Young's Modulus Y and the total stress acting on the body:

$$\varepsilon = \frac{\sigma}{Y}. \tag{2.3}$$

An ideal elastic body responds instantly to an applied stress. Considering an anelastic body under the same force, the response has some delay [118]. Assuming that the applied stress σ is periodic, it can be defined in terms of an amplitude σ_0

$$\sigma = \sigma_0 e^{i\omega t}, \tag{2.4}$$

with a periodic component defined by the angular frequency of the oscillation ω over a given time t. The resulting strain on the same body ε, maintains the same form as Eq. 2.4 offset by a phase term Φ:

$$\varepsilon = \varepsilon_0 e^{i\omega t - i\Phi}. \tag{2.5}$$

The phase lag Φ between ε and σ is also known as the mechanical loss $\phi(\omega)_{\text{mech}}$, or internal friction, of the material. For anelastic materials, the internal friction $\Phi \equiv \phi_{\text{mech}}$, can be directly related to the magnitude of thermal noise exhibited by the

material. It has been shown that any changes in the internal material structure such as grain boundaries [119], dangling bonds [7] and imperfections in the glass structure can all contribute to a higher internal friction level. There is also experimental evidence that the internal friction of a glass-like material can be influenced by temperature change, heat treatment and doping [120–122].

2.1.3 Thermal Noise of a Single Resonant Mode

To understand the complexities of thermal noise in a system with multiple resonances, first, the simplified case of a single-resonance system will be investigated. In this case the thermal noise can be described by the behaviour of a damped simple harmonic oscillator. $\phi(\omega)_{\text{mech}}$ is given by the total energy dissipated over one oscillation cycle, divided by the stored energy. The dimensional properties such as the spring constant k and mass m must be known

$$\phi(\omega)_{\text{mech}} = \frac{E_{lost\,per\,cycle}}{2\pi E_{\text{stored}}}. \tag{2.6}$$

These properties in turn with the total of the body displacement x, and the force F imparted by the thermal motion can be used in Hooke's Law. Starting with Hooke's law and substituting the spring constant k for the complex spring constant $\overline{\mathbf{k}} = k(1 + i\phi(\omega))x$ [118], gives

$$ma = k(1 + i\phi(\omega))x - F, \tag{2.7}$$

where $\phi(\omega)$ can be expressed as:-

$$\phi(\omega) = \frac{-ima + ikx + iF}{kx}. \tag{2.8}$$

In order to calculate the thermal noise contributions of the system, Eq. 2.7 must be expressed in terms of the velocity of the mass m. The acting force on the system then takes the form

$$F(\omega) = \frac{v\,(k[\phi(\omega) - i] + im\omega^2)}{\omega} \tag{2.9}$$

where velocity v of the oscillating mass is a function of position and angular frequency, $v = x\,i\,\omega$. The total mechanical admittance of the sample $Y(\omega)$ can then be calculated by substituting Eq. 2.9 into Eq. 2.1

$$\mathbb{R}[Y(\omega)] = \frac{k\phi(\omega)\,\omega}{k^2\phi(\omega)^2 + (m\omega^2 - k)^2}. \tag{2.10}$$

Knowledge of this relationship allows the thermal noise of the oscillator to be calculated by using equations 2.10 and 2.2 producing the power spectral density $S_x(\omega)$:-

$$S_x(\omega) = \frac{4k_\mathrm{B}T}{\omega}\frac{k\phi(\omega)}{k^2\phi(\omega)^2 + (m\omega^2 - k)^2}. \tag{2.11}$$

This function gives the power spectrum of the thermal motion due to a single resonance as a function of ω. To calculate the frequency dependent power spectral density $S_x(f)$, one must convert from angular frequency and take the equivalence of $k = m\omega_0^2$ [123] where ω_0 is the angular resonant frequency

$$S_x(f) = \frac{k_\mathrm{B}T\phi(f)f_0^2}{2\pi^3 f m\left(\left(f_0^2 - f^2\right)^2 + f_0^4\phi^2(f)\right)}. \tag{2.12}$$

$S_x(f)$ is related to the mechanical loss $\phi(f)$ of the resonance. It is interesting to note the behaviour of the thermal noise of the resonance at frequencies some distance away. In this consideration, it will be assumed that the mechanical loss of the resonator is much lower than unity ($\phi(f) << 1$). For $f << f_0$, the f term in the subtraction can be considered virtually zero; in this case, the predicted thermal noise decreases as a function of frequency

$$S_x(f) = \frac{k_\mathrm{B}T\phi(f)}{2\pi^3 m f_0^2 f}. \tag{2.13}$$

Conversely, at frequencies higher than the resonance $f >> f_0$ the power spectral density continues to decrease as a function of

$$S_x(f) = \frac{k_\mathrm{B}T f_0^2\phi(f)}{2\pi^3 m f^5} \tag{2.14}$$

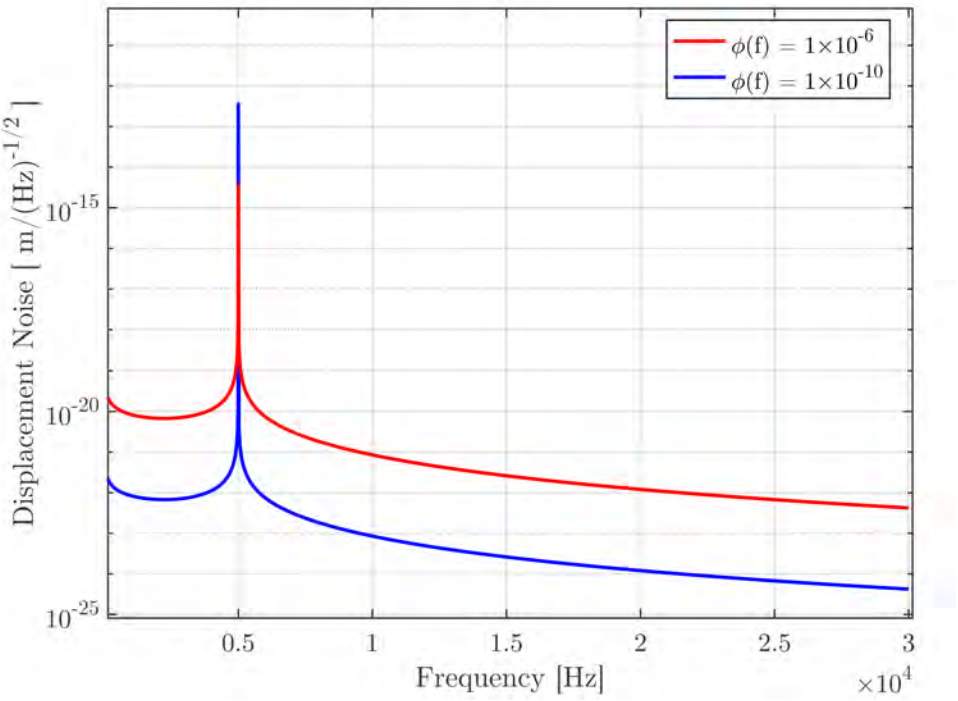

Figure 2.1: The effect of mechanical loss on thermal noise spectrum of two otherwise identical oscillators (mass 40 kg, resonant mode at 5 kHz, temperature 293 K). Lower loss (shown in blue) can be seen to concentrate more of the thermal motion close to the resonant frequency, reducing the off-resonance thermal noise.

On resonance (where $f = f_0$) the maximum value of thermal noise is produced and $S_x(f)$ takes the following form:

$$S_x(f) = \frac{k_B T}{2f^3 m\phi(f_0)\pi^3}. \tag{2.15}$$

Figure 2.1 shows an example of this behaviour where the total displacement noise of the oscillator is given as $(S_x(f))^{-\frac{1}{2}}$. For a lower loss (1×10^{-10}) and higher loss (1×10^{-6}) oscillator of 5 kHz calculated using Eq. 2.12 are shown. This clearly illustrates that for an oscillator with lower mechanical loss, the thermal noise is higher at $f = f_0$, and the resonance peak is narrower, lowering the thermal motion far from the resonance. This is a crucial concept for gravitational wave detector applications where the use of low loss materials concentrates the majority of the thermal motion in narrow bands close to the resonances. The pendulum/bounce resonances of the suspensions occur at lower frequencies (<20 Hz) compared to the internal frequencies of the test mass optics ($f_0 > 5$ kHz). High-quality factor

suspension 'violin modes' however, are present throughout the detection band, but can be effectively filtered out from detector data [124].

2.1.4 Substrate Brownian Thermal Noise

In reality, a gravitational wave detector mirror has many vibrational modes which can be excited by thermal energy. Raab *et al.* [125] estimated the thermal noise of a mirror as the sum of contributions from n mechanical modes:

$$S_x(f) = \sum_n \frac{4k_{\mathrm{B}}T\phi_n(f)}{2\pi^3\alpha_n f_0^{n\,2} m}, \tag{2.16}$$

where ϕ_n is the mechanical loss of the n^{th} mode and α_n describes the spatial overlap between the intensity profile of a Gaussian beam and the motion of the front surface of the test mass at a resonance f_0^n. In Eq. 2.16 the mechanical loss of each mechanical mode of a test mass is assumed to be spatially homogeneous. However, the presence of inhomogeneous losses - for example, from the mirror coating applied to the front surface of the mass - can result in interactions between different modes. In this scenario, the total thermal noise is no longer simply given as the summation over many modes shown in Eq. 2.16. A second point which is not considered in this approach is the possibility that the mechanical loss of the substrate and coating layer can be inhomogeneous across the surface and depth of each material.

To account for inhomegeneous loss, Levin [126] suggested the following thought experiment: A test mass of radius R, with thickness t, is subject to an oscillating force F described by a Gaussian profile identical to that of the interferometer laser beam

$$F(\boldsymbol{r}) = \frac{2}{\pi r_0^2} \exp\left(-\frac{2r^2}{r_0^2}\right), \tag{2.17}$$

radius r_0 whose intensity drops to $\frac{1}{e^2}$ at a position r. In the case where the loss of the substrate is inhomogeneous the thermal noise of the system cannot be calculated using Eq. 2.16. Levin's thought experiment involves calculating the deformation of the test mass due to the Gaussian pressure at every point inside the test mass. The strain energy at every point, normalised to the total strain energy in the mirror, can be multiplied with the loss of every point. The summation of this product over the

whole test mass gives the total dissipated power, $W_{\rm diss}$. This can then be calculated using the general theory of elasticity [126] and expressed in the following form [127]:

$$W_{\rm diss} = 2\pi f \int_{\rm vol} \varepsilon(x,y,z)\phi(x,y,z,f)dV, \tag{2.18}$$

where ε is the energy density of the elastic deformation when the test mass is at a point of maximum deformation. The total deformation of the test mass sensed by the interferometer laser is then accounted for in Eq. 2.19 by weighting the power spectral density by the spatial profile of the Gaussian beam. If the resulting strain and the mechanical loss of every point in the test mass are known, the thermal noise sensed by $F(\boldsymbol{r})$ is then related to the total dissipated power $W_{\rm diss}$, the sum of the mechanical loss and strain in the optic and F_0 is the peak magnitude of the oscillating force

$$S_x(f) = \frac{2k_{\rm B}T}{\pi^2 f^2}\frac{W_{\rm diss}}{F_0^2}. \tag{2.19}$$

As a result of Levin's work, it was realised that a test mass with a concentration of lossy material at the front face of the optic, close to the sensing laser beam, will produce a higher level of thermal noise than an identical layer of material embedded in the test mass. In the real world case, an Advanced LIGO (aLIGO) test mass is coated with 18 bi-layers of silica and titania doped tantala, whose loss $\phi_{\rm coating} = 2.4 \times 10^{-4}$ [128], is orders of magnitude higher than that of the fused silica substrate $\phi_{\rm silica} \approx 1 \times 10^{-9}$ [86]. To achieve the required reflectivity, the high loss coating material must be deposited on the front face of each optic, at the point of interaction with the sensing laser, producing a significant source of thermal noise. Therefore Levin's approach shows the significant role that a lossy HR mirror coating has on thermal noise in a gravitational wave detector.

2.1.5 Coating Brownian Thermal Noise

Nakagawa *et al.* [129] made the first attempt to theoretically calculate the coating thermal noise in a mirror. As the radius r_0 ($I \equiv \frac{1}{e^2}$) of the interferometer beam in Eq. 2.17 is much smaller than the radius of the test mass $\boldsymbol{r}$, the front surface of the optic is considered to be semi-infinite while the thickness of the optic is finite. The coating was approximated as a body of thickness t and loss $\phi_{\rm coat}$ on the front surface of the test mass, which, following Levin's approach [126], has uniform loss $\phi_{\rm sub}$. For simplicity the HR coating stack was modelled as a single layer with the

average material properties the component materials. The power spectral density of the thermal noise is described by the dissipated energy $W_{\rm diss}$ and the total mechanical loss of the test mass $\phi_{\rm total}$

$$S_x(f) = \frac{2k_{\rm B}TW_{\rm diss}}{\pi f F_0^2}\phi_{\rm total}(f). \tag{2.20}$$

In this paper a derivation by Liu and Thorne [130] who calculated the total loss of a coated test mass based on elastic theory was used. The energy dissipation $W_{\rm diss}$ has also been reformulated to take into account the presence of the additional coating layer

$$W_{\rm diss} = \frac{F_0^2\,(1-\sigma_{\rm sub}^2)}{2\sqrt{\pi}r_0 E_{\rm sub}}, \tag{2.21}$$

where $E_{\rm sub}$ and $\sigma_{\rm sub}$ are the Young's Modulus and Poisson Ratio of the substrate material. Nakagawa showed that by combining equations 2.20 and 2.21 the power spectral density of the thermal noise of a coated test mass could be stated as follows:-

$$S_x(f) = \frac{k_{\rm B}T\,(1-\sigma_{\rm sub}^2)}{\pi^{3/2} f r_0 E_{\rm sub}}\phi_{\rm sub}\left\{1 + \frac{2}{\sqrt{\pi}}\frac{(1-2\sigma_{\rm sub})}{(1-\sigma_{\rm sub})}\frac{\phi_{\rm coat}}{\phi_{\rm sub}}\left(\frac{t}{r_0}\right)\right\}. \tag{2.22}$$

2.1.6 Coating Thermal Noise in a Multi-layer Material

The simplification taken by Nakagawa fails to fully represent the distribution of loss inside the coating layers, and as such introduces a source of error on their calculated values. Illustrated by Levin, any inhomogeneities in the distribution of loss can affect the mechanical admittance of the material, and thus the level of thermal noise produced by the material.

Harry *et al.* [131] accounted for the distribution of material layers with different mechanical losses. As the mechanical admittance of a material is dependent on the loss at a given point, the elastic strain energy density, formally shown as $E_{\rm max}$ in Eq, 2.18 is resolved into two components. Each component describes the imparted stress $\sigma_{\rm ij}$ and the resultant strain $\varepsilon_{\rm ij}$ normal and perpendicular to the coating surface. Harry showed that the integral of elastic strain energy in the coating layer with respect to a given plane could be calculated. Expressions for the elastic strain energy density parallel $\delta E_{\parallel}$ and perpendicular $\delta E_{\perp}$ to the coating surface were produced

$$\delta E_{\parallel}/E_{\mathrm{sub}} = \frac{1}{\sqrt{\pi}r_0}\frac{Y_{\mathrm{coat}}(1+\sigma_{\mathrm{sub}})(1-2\sigma_{\mathrm{sub}})^2 + Y_{\mathrm{sub}}\sigma_{\mathrm{coat}}\left(1+\sigma_{\mathrm{coat}}\right)\left(1-2\sigma_{\mathrm{sub}}\right)}{Y_{\mathrm{sub}}\left(1+\sigma_{\mathrm{coat}}\right)\left(1-\sigma_{\mathrm{coat}}\right)\left(1-\sigma_{\mathrm{sub}}\right)}, \tag{2.23}$$

by taking into account the Young's Moduli and Poisson ratios of the coating and substrate materials. When normalised to the total energy contained within the substrate E_{sub}

$$\delta E_{\perp}/E_{\mathrm{sub}} = \frac{1}{\sqrt{\pi}r_0}\frac{Y_{\mathrm{sub}}\left(1+\sigma_{\mathrm{coat}}\right)\left(1-2\sigma_{\mathrm{coat}}\right) - Y_{\mathrm{coat}}\sigma_{\mathrm{coat}}(1+\sigma_{\mathrm{sub}})(1-2\sigma_{\mathrm{sub}})}{Y_{\mathrm{coat}}\left(1-\sigma_{\mathrm{coat}}\right)\left(1+\sigma_{\mathrm{sub}}\right)(1-\sigma_{\mathrm{sub}})}. \tag{2.24}$$

This formalism gives rise to two separate dissipation angles for a given coating layer, each associated with equations 2.23 and 2.24 using the same nomenclature. It was stated that the sum of these components scaled by the respective energy density in each plane and coating thickness t_{coat} would be equal to the total loss of a coated substrate

$$\phi_{\mathrm{coated}} = \phi_{\mathrm{substrate}} + \frac{\delta E_{\parallel} t_{\mathrm{coat}}}{E_{\mathrm{sub}}}\phi_{\parallel} + \frac{\delta E_{\perp} t_{\mathrm{coat}}}{E_{\mathrm{sub}}}\phi_{\perp}. \tag{2.25}$$

Factoring this into the Fluctuation Dissipation Theorem, the power spectral density coating thermal noise becomes:

$$\begin{aligned} S_x(f) =& \frac{2k_{\mathrm{B}}T}{\pi^{3/2}f}\frac{1-\sigma_{\mathrm{sub}}^2}{w_0 Y_{\mathrm{sub}}}\left\{\frac{1}{\sqrt{\pi}}\frac{d}{w_0}\frac{1}{Y_{\mathrm{sub}}Y_{\mathrm{coat}}\left(1-\sigma_{\mathrm{coat}}2\right)\left(1-\sigma_{\mathrm{sub}}^2\right)}\right. \\ &\times\left[Y_{\mathrm{coat}}^2(1+\sigma_{\mathrm{sub}})^2(1-2\sigma_{\mathrm{sub}})^2\phi_{\parallel}\right. \\ &+ Y_{\mathrm{sub}}Y_{\mathrm{coat}}\sigma_{\mathrm{coat}}(1+\sigma_{\mathrm{sub}})\left(1+\sigma_{\mathrm{coat}}\right)\left(1-2\sigma_{\mathrm{sub}}\right)\left(\phi_{\parallel}-\phi_{\perp}\right) \\ &\left.\left.+Y_{\mathrm{sub}}^2\left(1+\sigma_{\mathrm{coat}}\right)^2\left(1-2\sigma_{\mathrm{coat}}\right)^2\phi_{\perp}\right]\right\}. \end{aligned} \tag{2.26}$$

When calculating the $S_x(f)$ using the method shown in Eq. 2.26, Harry *et al.* used coated silica substrates, where $\phi_{\parallel}$ was determined using the ringdown method (this is discussed in more detail in Chapter 5). However, this method is not capable of measuring $\phi_{\perp}$. Harry *et al.* therefore made the assumption that $\phi_{\perp} = \phi_{\parallel}$ when calculating the thermal noise levels of a multilayer coating. They also state that the accuracy of this method for determining coating thermal noise depends on the accuracy to which $\phi_{\perp}$ is known.

2.1.7 Bulk and Shear Losses

Hong *et al.* shows that if the assumption of $\phi_{\perp} = \phi_{\parallel}$ is not used in the method derived by Harry *et al.* the model can produce un-physical, negative values for the dissipated energy [132]. Hong *et al.* developed an alternative approach, based on the assumption that a coating material has two separate loss mechanisms, but are instead associated with different types of deformation - ϕ_{bulk}, associated with bulk motion, and ϕ_{shear}, associated with shear motion. The deformation associated with any resonant mode can be decomposed into a sum of bulk (volumetric deformation) and shear (shape deformation), [133], and the total loss of this mode will depend on the proportion of elastic energy stored in bulk and in shear deformation

$$E_{\text{T}} = E_{\text{bulk}} + E_{\text{shear}}. \tag{2.27}$$

The energy stored in bulk motion (E_{bulk}), shown in Eq. 5.4 is expressed in the form of an integral with respect to the volume of the geometry:

$$E_{\text{bulk}} = \int \frac{1}{2} K \theta^2 dV. \tag{2.28}$$

This shows that the amount of energy stored in bulk motion is defined by the materials bulk modulus K, scaled by the first component of a 3×3 strain matrix S_{ii} denoted by θ^2. The energy dissipated through shear motion (E_{shear}) can also be calculated similarly:

$$E_{\text{shear}} = \int \mu \xi_{ij} \xi_{ij} dV. \tag{2.29}$$

Here μ denotes two components of a 3×3 compliance matrix which denotes the real components of the materials shear modulus. Terms denoted by ξ_{ij} refer to the respective components of the strain matrix S_{ijk} scaled by a Kroneker delta function, g_{ij}

$$\xi_{i,j} = \frac{1}{2} \left(S_{ij} + S_{ji} \right) - \frac{1}{3} g_{ij} S_{kk}. \tag{2.30}$$

The first approach taken by Hong is to show how the thermal noise of a single layer coating can be calculated where E_{shear} and E_{bulk} are taken into consideration. In this case, the effects of light penetration into the coating are ignored. In the physical case where the $t_{\text{coat}} << t_{\text{sub}}$ and $\boldsymbol{r} >> r_0$ the coated test mass can be considered

semi-infinite, similar to the approach taken by Nakagawa [129]. The dissipation of energy $W_{\rm diss}$ in reaction to an oscillating force F($\mathbf{r}$) can then be expressed as the sum of losses associated with the coating layer and the substrate, respectively

$$W_{\rm diss} = \phi_{\rm sub} E_{\rm sub} + \phi_{\rm bulk} E_{\rm bulk} + \phi_{\rm shear} E_{\rm shear}. \tag{2.31}$$

To correctly account for fluctuations in the thickness of the coating due to thermal noise, the implications of the substrate/coating interface also had to be included. The expanded forms derived by Hong *et al.* to calculate the normalised energy stored in bulk $E_{\rm bulk}$ and shear $E_{\rm shear}$ are shown below in equations 2.32 and 2.33. Here it was stated that the total dissipation of energy inside the single coating layer is inversely proportional to the term defining the effective laser beam area $\mathcal{A}_{\rm eff}$.

In these formulations, the effects that the coating and substrate have on the total energy dissipation, and indeed each other, is brought to light. During bulk deformation of the coating layer, the volume of the coating expands in all directions increasing the total thickness of the coating layer. This expansion also shifts the position of the coating/substrate interface. Contrary to this effect, during shear expansion of the coating layer there is no volume expansion, only a shape change in the material. This shape change causes the coating to contract with respect to the coating/substrate interface. In their paper, Hong *et al.* calculated the relationships between the bulk and shear energy dissipation in a multi-layer test mass. Normalising each term to F_0 produces the following formulations:-

$$\begin{aligned}\frac{E_{\rm bulk}}{F_0^2} =& \frac{(1-2\sigma_{\rm coat})\,l}{3}\left[\frac{Y_{\rm coat}}{Y_{\rm sub}^2}\frac{(1-2\sigma_{\rm sub})^2\,(1+\sigma_{\rm sub})^2}{(1-\sigma_{\rm coat})^2}\right. \\ &\left. + \frac{1}{Y_{\rm sub}}\frac{2\,(1-2\sigma_{\rm sub})\,(1+\sigma_{\rm sub})\,(1+\sigma_{\rm coat})}{(1-\sigma_{\rm coat})^2} + \frac{1}{Y_{\rm coat}}\frac{(1+\sigma_{\rm coat})^2}{(1-\sigma_{\rm coat})^2}\right] \\ &\times \frac{1}{\mathcal{A}_{\rm eff}}\end{aligned} \tag{2.32}$$

$$\frac{E_{\rm shear}}{F_0^2} = \frac{2l}{3}\left[\frac{Y_{\rm coat}}{Y_{\rm sub}^2}\frac{\left(1-\sigma_{\rm coat}+\sigma_{\rm coat}^2\right)\left(1+\sigma_{\rm sub}\right)^2\left(1-2\sigma_{\rm sub}\right)^2}{\left(1-\sigma_{\rm coat}\right)^2\left(1+\sigma_{\rm coat}\right)}\right.$$
$$-\frac{\left(1+\sigma_{\rm coat}\right)\left(1-2\sigma_{\rm coat}\right)\left(1-2\sigma_{\rm sub}\right)\left(1+\sigma_{\rm sub}\right)}{Y_{\rm sub}\left(1-\sigma_{\rm coat}\right)^2} \tag{2.33}$$
$$\left.+\frac{\left(1-2\sigma_{\rm coat}\right)^2\left(1+\sigma_{\rm coat}\right)}{Y_{\rm coat}\left(1-\sigma_{\rm coat}\right)^2}\right]\frac{1}{\mathcal{A}_{\rm eff}}.$$

Using the Fluctuation Dissipation Theorem and the dissipation formulated by Hong for a single layer coating in Eq. 2.31 the thermal noise of a coated test mass can be formulated as

$$S_{\bar{\xi}} = \frac{4k_{\rm B}T}{\pi f}\left[\phi_{\rm B}\frac{E_{\rm bulk}}{F_0^2}+\phi_{\rm S}\frac{E_{\rm shear}}{F_0^2}\right]. \tag{2.34}$$

While ignoring light penetration into the coating layer, Hong showed that the treatment of Harry *et al.* underestimated the total thermal noise in an aLIGO mirror. Hong *et al.* also goes on to produce a complete formation of thermal noise in a multilayer coating which requires knowledge of the light field passing through each coating layer [132]. This work was then simplified by Yam *et al.* [134] under the assumption that $\phi_{\rm bulk} \equiv \phi_{\rm shear}$. In this case, the power spectral density of coating Brownian thermal noise is given by

$$S_z^{\rm Br} = \frac{4k_{\rm B}T}{\pi r_G^2 r_0}\frac{1-\sigma_{\rm sub}-2\sigma_{\rm sub}^2}{Y_{\rm sub}}\sum_j b_j d_j \phi_{Mj}, \tag{2.35}$$

where $k_{\rm B}$ is the Boltzmann constant, T the mirror temperature, f the frequency and r_0 the radius of the interferometer laser beam on the coating. $Y_{\rm sub}$ and $\sigma_{\rm sub}$ are Young's modulus and the Poisson ratio of the substrate. d is the coating thickness and ϕ the coating mechanical loss. The index j refers to the material parameters of the $j^{\rm th}$ layer in the coating (starting from the outermost layer). b_j is a weighting factor described by:

$$b_j = \left[\frac{(1-2\sigma_j)(1+\sigma_j)}{(1-2\sigma_{\rm sub})(1+\sigma_{\rm sub})}\right]\frac{1}{1-\sigma_j}$$
$$\times\left[\left(1-n_j\frac{\partial\theta_c}{\partial\theta_j}\right)^2\frac{Y_{\rm sub}}{Y_j}+\frac{(1-\sigma_{\rm sub}-2\sigma_{\rm sub}^2)^2}{(1+\sigma_j)^2(1-2\sigma_j)}\frac{Y_j}{Y_{\rm sub}}\right], \tag{2.36}$$

where σ_j and $\sigma_{\rm sub}$ are the Poisson's ratios of each coating layer and the substrate and n_j is the refractive index of the $j^{\rm th}$ layer is multiplied by $\partial\theta_c/\partial\theta_j$. Together, these terms describe the phase noise contributions of each layer θ_c to fluctuations in the round-trip phase in each layer, θ_j. The weighting factor b_j takes into account the contributions of several effects inside the coating layer. As was described by [132], the expansion and contraction of each layer affects the position of the substrate. This is taken into account by scaling the phase noise sensitivity by the ratio of Young's modulus of the substrate and Young's modulus of that layer respectively. Conversely, Hong also stated that the coating/substrate interface could also influence the expansion and final position of each coating layer. This effect is also considered by the function of Poisson ratio's for the substrate and j$^{\rm th}$ layers, which are scaled by $\frac{Y_j}{Y_{\rm sub}}$. In Eq. 2.35 b_j is then scaled by d_j the thickness of each layer, accounting for the changes in the thickness of the j$^{\rm th}$ layer and the total loss of material ϕ_j. This leads to the interesting result that the level of thermal noise produced by a coating layer is also a function of its distance from the coating/substrate interface discussed by [132] in the derivation of Eq. 2.36.

However, it has also been noted by [134], the author, and others [135–137] that the ability to calculate the total thermal noise using Hong's approach is dependent on the knowledge of $\phi_{\rm bulk}$ and $\phi_{\rm shear}$. As the evaluation of $\phi_{\rm bulk}$ and $\phi_{\rm shear}$ from ringdown measurements is still in its infancy methods used to evaluate these terms developed in part by the author are described in Chapter 5.

2.2 Thermoelastic Noise

A body in thermal equilibrium will undergo spontaneous temperature fluctuations. Due to the coupling of the temperature and strain fields in the body, via the coefficient of thermal expansion, this will result in deformation of the mirror and motion of its front face, resulting in so-called thermoelastic noise. The thermal noise of a multi-layer coating due to pure thermoelastic noise was derived by M. Fejer *et al.* [138]

$$S_x(f) = \frac{8\sqrt{2}k_{\rm B}\,\alpha_{\rm S}^2 C_{\rm S}^2\,l^2\,(v+1)^2\,T^2\mathbf{C}^2}{\pi w^2\sqrt{\omega}\sqrt{C_{\rm S}k(s)}}\tilde{\Delta}^2. \tag{2.37}$$

In order to simplify the complex calculations of thermoelasticity for a multilayer

coating, the heat capacity $\bar{C}$ is used to denote the thickness averaged heat capacity of the coating stack, where other average properties are shown with the same notation. Terms which are denoted with a subscript 's', such as the thermal expansion α_{S} and heat capacity C_{S} refer to the material properties of the test mass substrate. The term $\tilde{\Delta}^2$ is a dimensionless constant which accounts for the differences in Young's moduli, Poisson ratio ν between the average coating layer and the substrate

$$\tilde{\Delta}^2 \equiv \left\{ \frac{C_s}{2\alpha_s \bar{C}} \left(\frac{\bar{\alpha}}{1-\bar{\nu}} \left[\frac{1+\bar{\nu}}{1+\nu_s} + (1-2\nu_s) \frac{\bar{E}}{E_s} \right] \right) - 1 \right\}^2 . \tag{2.38}$$

From Eq. 2.38 it is shown that the thermoelastic noise produced by a coated substrate reduces in situations where the properties of both materials are similar. With equal properties of coating and substrate, levels of thermoelastic noise reduces to zero.

2.3 Thermo-Refractive Noise

Statistical fluctuations in the temperature T of the coating layers can result in changes of its optical properties. The high reflectivity of these coatings is produced through Bragg reflection, where the optical path length and phase of light reflected at each interface determines the phase of the reflected laser light. As the refractive index n is temperature dependent, the phase of the laser light which penetrates into or is reflected by the coating layers can be altered by these fluctuations in temperature. The susceptibility of each material to changes in temperature is defined by its thermo-optic coefficient of $\beta = \frac{dn}{dT}$. The phase shift $\delta\phi$ produced from temperature fluctuations in a multi-layer coating is given in [139]

$$\delta\phi = 4\pi\, \beta_{\mathrm{eff}} \overline{u}. \tag{2.39}$$

The average fluctuation in temperature is given by $\overline{u}$ and β_{eff} and takes into account the change in the refractive index of the high and low index layers in each Bragg reflecting pair

$$\beta_{\mathrm{eff}} = \frac{n_2^2 \beta_H + n_H^2 \beta_L}{4\,(n_H^2 - n_L^2)}. \tag{2.40}$$

Expressing the resultant phase noise as a power spectral density such that it can be easily compared to other detector noise sources yields the following:-

$$S_x(\omega) = \frac{\sqrt{2}\beta_{\text{eff}}^2 \lambda^2 \kappa T^2}{\pi r_0^2 \sqrt{\omega \rho C \lambda^*}}. \tag{2.41}$$

The total contributions of this source of phase noise from an HR coating can then be quantified as a function of angular frequency ω: the wavelength of incident laser light λ; the materials thermal conductivity, κ; density ρ; specific heat capacity C; and the effective wavelength of light inside the coating layers λ^* [139].

2.4 Thermo-Optic Noise

Thermo-elastic noise and thermo-refractive coating noise are related, as the same temperature fluctuations drive the local expansion/contraction of the coating and the local change in refractive index. Evans *et al.* developed a coherent treatment of both noise sources, referring to the total effect as thermo-optic noise [140]. He found that the level of phase noise produced by this effect could be derived using the same formalism developed by Zener to calculate thermoelastic noise. The power spectrum observed by an incident Gaussian beam is given by

$$S_{\text{TO}}^{\Delta T}(f) = \frac{2 k_{\text{B}} T^2}{\pi^{3/2} w_r^2 \sqrt{\kappa_{\text{S}} C_{\text{s}} \rho_{\text{s}} f}}. \tag{2.42}$$

As the optic expands and contracts due to temperature fluctuations, the optical path length inside the material also changes. Thus the total noise due to temperature fluctuations in the coating is given by

$$S_{\text{TO}}^{\Delta z}(f) \simeq S_{\text{TO}}^{\Delta T}(f) \left(\tilde{\alpha}_{\text{c}} d - \tilde{\beta}_{\text{c}} \lambda - \tilde{\alpha}_{\text{s}} d \frac{\bar{C}_{\text{c}}}{C_{\text{s}}} \right). \tag{2.43}$$

Using the same formalism produced by Evans *et al.* [140] the total power spectral density produced in the Δz direction (normal to the test mass optic) is a function of the laser wavelength λ. The heat capacity of the coating and substrate C_c, C_s and the thickness of the HR coating d. It should be noted that to calculate the noise produced from a multilayer HR coating, the thermal expansion α_c and β_c, the description of phase reflection are both values which are averaged with respect to the thickness of each coating layer. α_c is described as a function of the thermal expansion of a coating material at position X and its respective Poisson ratio ν

$$\tilde{\alpha}_{\mathrm{c}} \sim 2\bar{\alpha}_{\mathrm{X}}\left(1+\bar{v}_{\mathrm{X}}\right). \tag{2.44}$$

This assumption only holds if the thermal expansion of the substrate and coating materials are similar. The term β_c takes into account the differences in refractive index between high (H) and low (L) index layers in the HR coating

$$\tilde{\beta}_{\mathrm{c}} \simeq \frac{B_{\mathrm{H}} + B_{\mathrm{L}}\left(2\left(n_{\mathrm{H}}/n_{\mathrm{L}}\right)^2 - 1\right)}{4\left(n_{\mathrm{H}}^2 - n_{\mathrm{L}}^2\right)}, \tag{2.45}$$

where B_{L} and B_{R} are terms which account for the small changes in optical path length in each material. Eq. 2.43 takes both the thermoelastic and thermo-refractive noise into consideration where $\alpha_{\mathrm{s}}d$ contains information of thermoelastic noise, and with $\tilde{\beta}_{\mathrm{c}}\lambda - \tilde{\alpha}_{\mathrm{s}}d\frac{\bar{C}_{\mathrm{c}}}{C_{\mathrm{s}}}$ detailing the contribution of thermo-refractive effects. It is interesting to note that the effects of thermo-refractive noise are negative (with respect to thermo-elastic noise) as the coating material expands as a function of temperature and the optical path length that the reflected laser light follows increases in turn. The subtraction of both of these terms shows that certain materials (depending on the sign of β), can produce partial cancellation of the thermoelastic and thermo-refractive noise terms, resulting in a smaller overall effect.

2.5 Status of Current HR Coatings

The operation of a gravitational wave detector relies on highly reflecting (HR) mirror coatings (R>99.999%) which absorb $<1\%$ of the light from the main interferometer laser. The mechanical dissipation and the Brownian thermal noise produced by these materials currently limit detector sensitivity between 50 Hz and 150 Hz [141]. Reduction of coating thermal noise and maintaining low optical absorption is then essential for improving detector sensitivity to gravitational wave sources and maintaining system stability over prolonged periods.

Current gravitational wave detector end test mass (ETM) coatings consist of $\approx$18 bi-layers of silica and titania doped tantala (TiO_2:Ta_2O_5) deposited by ion beam sputtering (IBS) (see section 4.3.2). The mechanical loss of this coating stack is influenced by the layers of Ti:Ta_2O_5 which have ϕ_{mech} $\sim 4\times10^{-4}$ compared to ϕ_{mech} $\sim 5\times10^{-5}$ produced by the silica layers [131, 142].

By producing an alternating stack of high refractive index $n_{\text{high}} = 2.12$ and low $n_{\text{low}} = 1.44$, this creates multiple Bragg reflecting layers. The thickness of each coating layer t_{layer} is such that $t_{\text{layer}} = \lambda_{\text{IFO}}/4$ where $\lambda_{\text{IFO}} = 1064\,\text{nm}$ is the wavelength of the main interferometer laser.[1]

For reference, values of thermal noise sources for an aLIGO detector are shown in Fig 2.2 where coating thermal noise has been calculated using Hong's approach.

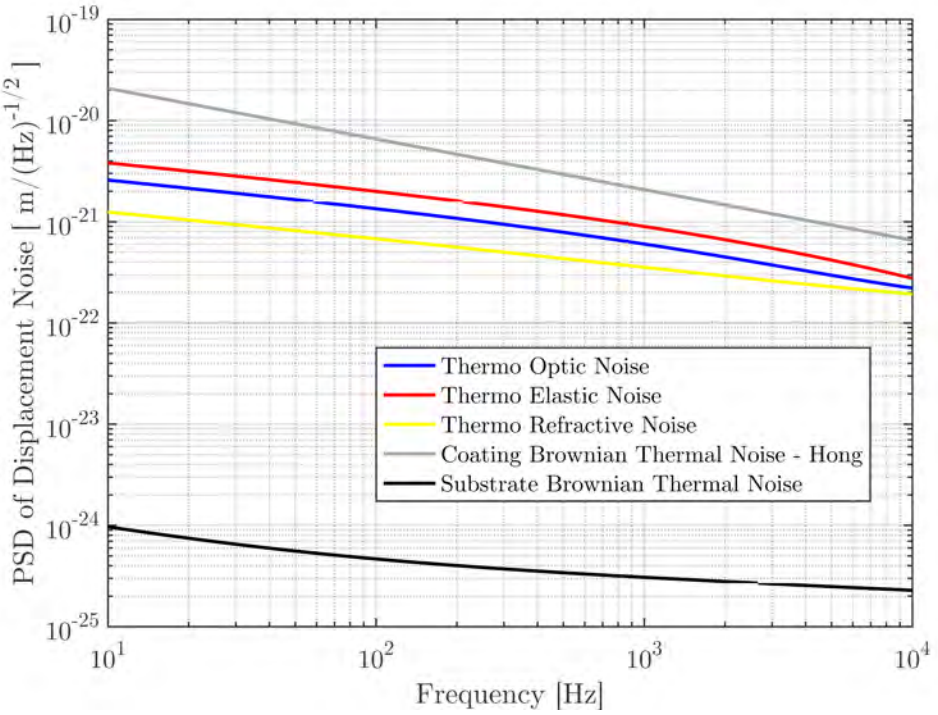

Figure 2.2: Comparison of thermally driven noise sources in an aLIGO test mass calculated using GWINC. Contribution of substrate thermal noise has been shown for comparison.

The internal atomic structure of the coating materials influences the coatings mechanical loss and optical absorption. Differences in the internal structure can be represented as a system of potential energy states with different stable positions. Increasing the total thermal energy available to the material by heat treating the material after deposition can change the glass' structure, producing a material with

[1] In practice the coating thickness also takes into account a secondary laser $\lambda = 533\,\text{nm}$ which is used for lock acquisition.

lower mechanical losses [122, 143].

Recent works [143, 144] have shown that depositing some coatings at an elevated temperature commonly known as 'elevated temperature deposition', can increase the mobility of atoms in the coating layer. This can lead to an 'ideal glass' state [143] with a lower density of two-level tunnelling states, connected with a reduction in mechanical loss. However, higher heat treatment temperatures can produce more 'brittle' glasses, which are then likely to crack or crystallise; due to increased optical scattering, it is impractical for use in a gravitational wave detector.

Research into alternative high index coating materials is in progress. Recent works have shown that by doping the high refractive index layers of Ta_2O_5 with TiO_2 or ZrO_2 increases the materials crystallisation temperature and can provide $>40\%$ reduction its mechanical loss [137, 145]. As such current aLIGO and AdV detectors use layers of Ti:Ta_2O_5 which produce $\phi_{\text{coating}} \approx 2{\times}10^{-4}$.

Alternative methods to prevent material crystallisation are also being developed. This involves using progressively thinner layers of SiO_2-TiO_2 coating material ($<$40 nm) which produces a composite material whose average refractive index is comparable to TiO_2. It has been shown that by stacking these 'nano-layers', this can also suppress the crystallisation temperature of the TiO_2 layers [146, 147]. The use of materials such as silicon nitride (Si_3N_4) in this technique is also of interest to detectors operating at lower temperatures as this method can suppress increases in loss which occur at these temperatures [147, 148].

Alternative high index amorphous materials such as amorphous silicon (aSi) or amorphous germanium (aGe) are also of interest to third-generation gravitational wave detectors which plan to operate at cryogenic temperatures. Due to their high optical absorption ($>$1000 ppm) at 1064 nm these materials cannot be a direct high index replacement material. However, by using longer wavelength laser light and operating at lower temperatures, the optical and mechanical properties of these materials become of great interest to third-generation gravitational wave detectors [78, 104].

The main focus of this thesis will be the experimental verification of a novel

multi-material coating design. This coating stack uses SiO_2-Ta_2O_5 layers to reduce the transmission of laser light to lower layers to a point where the high absorption of SiO_2-aSi layers can be tolerated. This allows the thermal noise benefits of aSi to be partially exploited, without significantly increasing the total coating absorption. J. Stienlechner and I. Martin *et al.* [149] first proposed this coating design which details an HR stack configuration which requires less coating layers to produce the required $R > 99.999\,\%$, which meets the requirements in reflectivity for a test mass optic, and reduces the total level of coating thermal noise inside a cryogenic detector.

2.6 Conclusions

Thermal noise from coated test masses limits the sensitivity of gravitational wave detectors with thermally driven processes in the HR coating layers dominating over effects from the fused silica substrate. At the most sensitive frequency band ($\sim 50\,\mathrm{Hz}$ - $150\,\mathrm{Hz}$) a gravitational wave detectors sensitivity is limited by coating Brownian noise, stemming from thermally driven vibrations in the coating layers. The level of coating Brownian noise arising from an HR coating is ultimately determined by the temperature (i.e. the amount of thermal energy available to the system), the level of mechanical dissipation of each material, the thickness of the coating and the radius of the interferometer laser beam on the mirror. In addition to Brownian thermal noise, thermal-optic noise associated with temperature fluctuations in the test mass can also occur due to the temperature dependence of the thermal expansion coefficient and the refractive index.

With careful consideration of the materials used in the coating and substrate, however, this can be significantly reduced. In the current detector's like aLIGO, the contribution of thermo-optic noise is an order of magnitude lower than coating Brownian thermal noise (see Fig 2.2).

Third-generation detectors are already planned to operate in the cryogenic regime, at temperatures below $200\,\mathrm{K}$. Silica test masses, which are currently used in room temperature detectors will not be suitable for use in cryogenic detectors. Due to the to an increase in mechanical loss (by more than 5 orders of magnitude) at these temperatures [150]. This increase is associated thermally activated transitions of

oxygen atoms between two energy states inside the material [151]. Ongoing work on molecular modelling of oxide materials suggests that the bonding structure of oxide-based glasses could be correlated to the level of mechanical dissipation as a function of temperature [152]. The use of silicon as a cryogenic test mass material has been proposed [153] as its low temperature mechanical loss is significantly lower [83, 154]. At these operating temperatures the level of noise produced from the thermo-optic effects in silicon substrates is comparable to coating Brownian noise and must be carefully considered.

Chapter 3

An In Situ Absorption Estimate of aLIGO Test Masses

3.1 aLIGO Thermal Compensation System (TCS)

The Advanced LIGO (aLIGO) detectors are made up of two 4 km-long Fabry-Pérot cavities in a perpendicular configuration. The current maximum input laser power is 125 W for each detector, which translates to a total circulating cavity power of 750 kW [94]. Some of the laser light which is transmitted through the input test mass (ITM) optics is absorbed, introducing a change of its bulk refractive index ($\frac{dn_{\mathrm{Bulk}}}{dT}$) along with a shape change in proportion to the absorbed laser power. Inside the Fabry-Pérot cavities, a small portion of the total laser light is also absorbed by the highly reflecting coatings (HR) deposited on the test mass optics producing a thermal lens on the surface of the optic, changing the optics' radius of curvature. The combination of bulk and coating absorption effects will hereafter be referred to as self-heating, which must be taken into consideration and controlled by a thermal compensation system (TCS) (see Fig 3.1) in order to maintain a high detector duty cycle. As discussed in more detail in Chapter 4, the power absorbed into a coated optic by laser heating is dominated by the coating layers and is proportional to the input laser power, P_{Input}, and α_{coat}, the absorption of its surface material.

3.1.0.1 Background

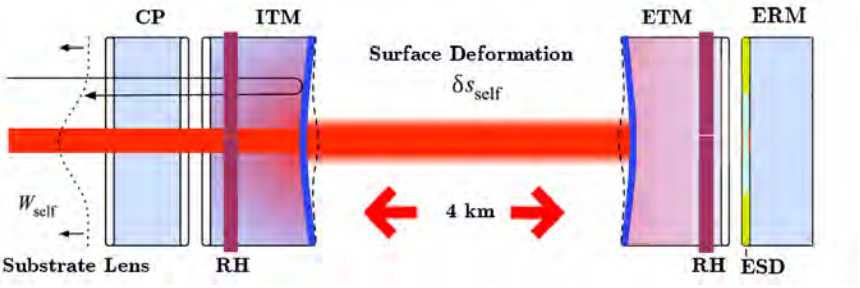

Figure 3.1: An illustration of the thermoelastic surface deformation Δ_{self} on and ITM and ETM optics on cavity resonance. Compensation plate (CP), Electro-Static Drive Plate (ESD) bonded to the End Reaction Mass optic (ERM) and Ring Heaters (RH) also pictured.

Fig 3.1 shows an aLIGO Fabry-Pérot cavity alongside the TCS apparatus. The Compensation Plate (CP), which is situated behind the Input Test Mass (ITM) [94], uses an external CO_2 laser (not pictured) to produce a variable bulk thermal lens which compensates for any beam distortion produced by the signal recycling cavities. The End Reaction Mass (ERM), situated behind the End Test Mass (ETM), allows the orientation of the ETM to be altered through the bonded Electro-Static Drive Plate (ESD).

An optical cavity can be 'locked' on resonance using the Pound-Drever Hall technique [155] by controlling the frequency of the input laser such that the phase-modulated laser light injected into the cavity is resonant in the cavity.

The aLIGO design specifications for ETM coatings require 0.5 ppm optical absorption and 99.999% reflectivity at 1064 nm. The total heat introduced into the test masses is dominated by the effects of coating absorption, as optical absorption of bulk SupraSil 3001 silica (0.25 ppm) at 1064 nm [156] is less than the absorption of an HR mirror coating. This means that 375 mW of equivalent heating power is dissipated into the test mass by the coating [94] compared to < 35 mW by both the ITM and CP substrates combined [157]. The localised laser heating creates a 2D radial gradient in temperature and refractive index across the test mass surface, changing its radius of curvature (RoC). This causes a distortion of the fundamental TEM_{00} Gaussian cavity mode (which has the spatial profile of a simple Gaussian distribution), causing the wave-fronts inside the cavity to shift or *defocus*, by δs,

where s is the initial height of the deformation (sagitta) with respect to the un-deformed mirror surface [158].

The normalised change in sagitta can be expressed in terms of the thermal expansion coefficient α, thermal conductivity κ, laser wavelength λ and absorbed laser power P_{α}:

$$\frac{\delta s}{s} \approx \frac{\alpha}{2\kappa\lambda}\mathrm{P}_{a}. \tag{3.1}$$

Since the refractive index of the coating, $\frac{dn(\mathrm{T})}{d\mathrm{T}}$, and Youngs Modulus $Y(\mathrm{T})$ of the optic both have a temperature dependence, any heating of the optics will also result in measurable changes in its mechanical properties such as mechanical mode frequency ω_{m}. This dependence can result in as much as a $\sim 0.2\,\mathrm{Hz}$ shift in measured a frequency of $\approx 5\,\mathrm{kHz}$ as laser power inside the cavity increases. In a case where the frequency spacing between the fundamental mechanical mode of the test mass and the frequency of the TEM modes becomes zero, the radiation pressure exerted on the test mass is now at the same frequency as the mechanical resonance. This effect can transfer energy into the mechanical resonance of the optic, producing 'parametric instabilities' in the detector [5]. If these are not actively damped, the mechanical resonances of the optics will become overly excited, causing the cavity to lose lock reducing detector stability and duty cycle (the total time that the interferometer is online). Further generations of gravitational wave detectors such as the Einstein Telescope - High-Frequency detector plan to increase the total circulating cavity power to a maximum laser power to 3 MW [78] to increase detector sensitivity. The effects of parametric instabilities are projected to increase by a factor 1.5 with this increase in laser power [5].

To combat effects of self-heating and parametric instabilities in aLIGO, an opposing thermal lens is created in the bulk of the optic with an external ring heater (RH). During detector operation the total heating required to combat self heating can be calculated if the total absorption of the HR coating layer is known. Currently methods used to combat test mass deformation rely on the measured absorption of the coating at the time of installation, and sporadic measurements taken by a Hartmann Wave-front sensor, discussed more in Section 3.2.

However, there is no way to directly produce a long term measurement of coating absorption α_{coat} after the optic has been installed and the detector is in observational mode. Any changes to the coating absorption during installation or operation will therefore make it difficult to mitigate the thermal lensing produced by self-heating. It has been shown upon inspection of the input test mass from LIGO Hanford's X-arm (ITM-X), that regions of anomalously high absorption ($\approx$ 1000$\times$ greater) can occur in an HR coating during detector operation [159]. Direct detection methods of such a defect during detector operation do not currently exist and as such this effect could worsen without the knowledge of its existence, affecting the duty cycle of the detector. To be able to produce a continuous estimate of the coating's absorption after installation would allow for more accurate correction of ring heater power during detector operation, thus increasing system stability. If upper and lower limits on the expected change in the mechanical mode frequencies of each optic with heating can be calculated, this will allow for frequency windowing of broadband frequency data produced during detector operation allowing measurements of how each mechanical mode shifts with temperature to be provided. As the absorption of the coating is the dominant source of heating in an ETM optic, if the measured change in mechanical mode frequency differs from predictions over long time scales (several months) it follows that changes in coating absorption are likely to be responsible for the measured changes.

In the next section, the creation of a finite element model of an ETM optic to model its response to thermal loads will be discussed to predict the behaviour of its mechanical mode frequencies. This model will be compared to the measured changes in the frequencies of test mass mechanical modes with laser heating over short time scales ($\approx$ 30 min) when the laser is powered on/off. Comparing the model to data over several months, any changes in the coating absorption can be estimated.

3.2 Finite Element Modelling

The test mass mechanical modes can be directly observed in the differential arm length signal (DARM) of the interferometer: the final output signal from the detector. During interferometer operation, the interferometer optics are heated due

to absorption of the laser light circulating in the arm cavities. Due to the temperature dependence of silica's Young's modulus $Y(T)_{\text{bulk}}$, the mechanical mode frequencies of the material will shift with heating. It can be shown that for the rotationally symmetric mechanical modes, frequencies of the optic can be calculated using:

$$\omega_m = \beta_m \sqrt{\frac{Y(T)_{\text{bulk}}}{\rho(1+\nu)}}, \tag{3.2}$$

where ρ is the material density, ν is its Poisson ratio and β_{m} is a constant which is dependent on the dimensions of the mirror [160].

A more accurate prediction of the room temperature resonant frequencies of the ETM optics, was calculated using COMSOL Multiphysics by C. Blair [5] (between 5 kHz and 40 kHz). A realistic geometry, shown in Figure 3.2, was used which takes into account the flat cut-away on either side of the optic's barrel used for suspensions bonding. The computational demands of running a model based on a realistic geometry could be reduced by using a simplified 2-dimensional asymmetric representation (2DAS) of the system. Fig 3.3 shows an example of the model produced by the author to reduce computational load and increase solution speed, using the 3D geometry as a reference for resonant mode frequencies.

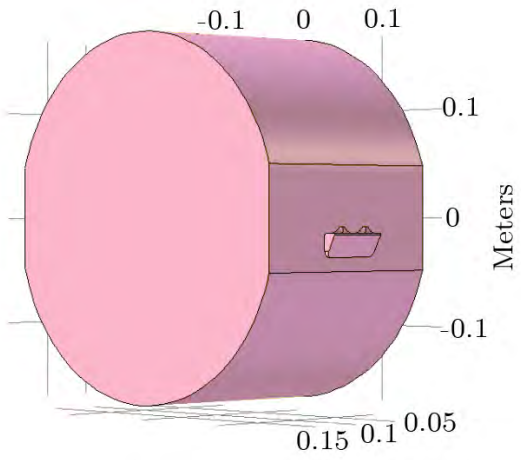

Figure 3.2: Realistic geometry from COMSOL Multiphysics used for mechanical resonance estimation of ETMs [5].

The geometry was approximated as a simple 2D slice of the 3D system, which can then later be visualised in '3D' by rotating the slice through 360° producing

a 3D representation of that system. This allows for multiple heating loads such as the effects from the main interferometer beam on the optical absorption of the coating, as well as scattering from the surrounding suspensions to be taken into consideration, without exponentially increasing the solution time. A comparison of realistic and 2DAS geometries was performed [161], which showed that this geometry approximation could represent the resonant frequencies of the geometry to within 0.4%. However, due to the confines of modelling resonances of the 3D system in a 2DAS regime, only rotationally symmetric mechanical modes can be modelled in this manner. As the frequencies of a realistic geometry are known, relative changes in the 2DAS modelled resonant frequencies due to heating can be compared with the predictions of the realistic geometry.

Material Property	Symbol	Fused Silica	Aluminium	Unit
Density	ρ	2203	2700	kg/m^3
Poisson's Ratio	ν	0.17	0.33	-
Youngs Modulus	E	7.3×10^{10}	7×10^{10}	Pa
Thermal Conductivitiy	κ	1.38	160	W/(mK)
Thermal Expansion Coefficient	α	$5.5\text{x}10^{-7}$	2.31×10^{-5}	1/K
Heat Capacity	C_p	740	900	J/(kg K)
Relative Permittivity	ϵ	3.8	1	-
Electrical Conductivity	σ	1×10^{-14}	3.77×10^{7}	S/m
Surface emissivity	ϵ_e	0.93	0.1	-

Table 3.1: Material Properties used in the 2DAS model of a coated LIGO test mass. [1].

To predict the frequency shift of the resonances due to self-heating, it is assumed that there is no change in α_coat =1.44 ppm from the time of installation. The incident laser beam is approximated as a *boundary heating term* with a Gaussian spatial profile in COMSOL applied to the front surface of the optic. The first order drumhead mode (8126 Hz) shown in Figure 3.4, is the first mechanical mode predicted by COMSOL, this frequency was consequently used to estimate the frequency response with heating.

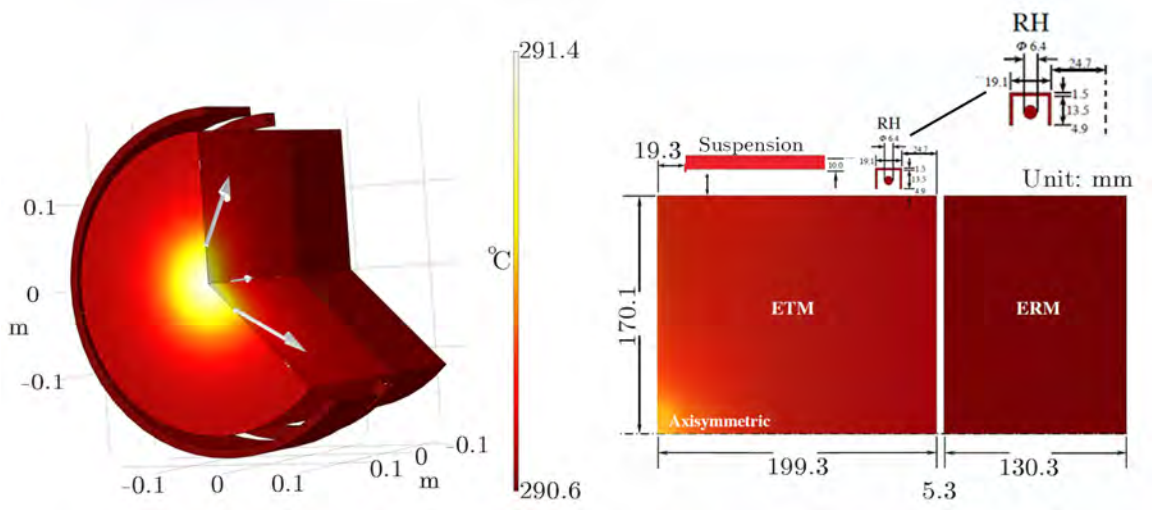

Figure 3.3: LHS shows an image of the 2DAS model rotated through 225° during a 30-hour heating cycle the total temperature shift in degrees Celsius. White arrows denote the direction of the radial temperature gradient through the optic. RHS shows the full dimensions of the 2DAS model taking into account the scattering contribution from the optics suspensions as an extra heating term [1].

To accurately account for the total heating power incident on the ETM coating surface, the total absorbed power P_α must have a radial profile related to the profile of the laser beam on the mirror. This can be achieved by taking into account the intensity profile of a Gaussian beam, assumed by the *boundary heating term* and scaling for the gain of the power recycling cavity $\mathrm{k_{PRC}}$ [162]

$$P_\alpha = P_{\mathrm{laser}} \times \mathrm{k_{PRC}}. \tag{3.3}$$

Taking the value of $\mathrm{k_{PRC}}$ from stable detector operation at 40 W input yields a value of $\mathrm{k_{PRC}}$ =270 [5]. Using the Gaussian intensity profile, the incident laser power absorbed by optic can be described by:-

$$P_{\mathrm{laser}} = 2\,\frac{P_{\mathrm{in}}}{\pi\omega^2}\exp\left(\frac{2r}{\omega}\right)^2\frac{1}{\alpha_{\mathrm{coat}}} \tag{3.4}$$

where ω is the beam radius and r, is the distance from the central axis of the beam.

Using an incident power of 100 kW with a coating absorption of 1.4 ppm equates to a total absorbed power of 0.14 W into the coating. Due to the coating's high reflectivity, effects of any transmitted light into the bulk of the optic can be

considered negligible. Measurements from the Hartmann Wavefront Sensor (HWS) were used to corroborate this value. HWS consists of an array of apertures which are used to measure the total thermal lens inside the test mass substrate [94]. A low power probe beam is transmitted through an aperture array from the HWS, through the back surface of the ETM and ERM/CP and reflected to the sensor. The diffraction of the reflected beams are measured with respect to an internal reference beam which allows the tilt of each reflected beam to be calculated to its output aperture. Integrating over the calculated tilt values for each aperture allows for a reconstruction of the reflected wavefront [163] from the test mass. This allows the radius of curvature of the test mass to be inferred. The change in optical path length from the thermoelastic effect can also be used to produce an estimate of coating absorption.

Using the HWS on ETM-Y at LIGO Livingston (LLO), the coating absorption was estimated to be between 1.3 to 2.2 ppm [5], much higher than the 0.23 ppm predicted by Pinard *et al.* [164].As this system is highly sensitive to external factors such as air currents, temperature fluctuations and the drifting of the optics over prolonged periods [94], the most accurate readings are produced over the first ∼3000 s after calibration. The method used by the author to predict coating absorption over timescales longer than one hour were tested against the HWS. The best estimate of coating absorption produced by the HWS of 1.4 ppm will be used as a reference basis in this chapter. As the most significant changes in test mass temperature occur over the first three hours after the interferometer gains lock, the HWS can only be used as a guide for coating absorption outwith these timescales.

Additional heating terms from point and figure scattering from the surrounding suspensions and ITM's are also taken into consideration [1]. From this it was shown by Wang *et al.* [161] that a 0.16 ppm absorption due to heating from figure scatter and 0.37 ppm from point scattering, equated to a total heating contribution of 0.5 ppm from the ring heater (RH) onto the ETM. These calculations allowed the effects of scattering to be taken into consideration by the author. The shift in ETM mechanical mode frequency was calculated for 30 hrs of self-heating, shown in Figure 3.4(a), giving significant overlap with HWS data.

Building on the work carried out by Wang [1] the 2DAS geometry of the ETM

was altered, taking into account the End Reaction Mass (ERM), Ring Heaters (RH) and suspensions to produce a more accurate representation of the physical system. By subjecting the front surface of the optic to 100 kW of laser power, the power absorbed by the coated face of the model is given by Eq. 3.4. The addition of each body to the model increases re-radiation effects from their respective surfaces onto the optic whereas the cooler surrounding components absorb photons which would have been radiated from the ETM. These bodies then re-radiate the lower energy photons back to the ETM, raising its temperature compared to the isolated state. The larger temperature increase of the test mass then causes a greater shift in mechanical mode frequency.

For simplicity the suspensions were approximated as a large block above the ETM, allowing the effects of scattering to be included. This modified model shows an increase in the calculated Δf. Figure 3.4(a) shows the calculated change in frequency with the addition of the surrounding bodies. The full geometry results in the frequency change being 0.04 Hz higher than with the optic alone.

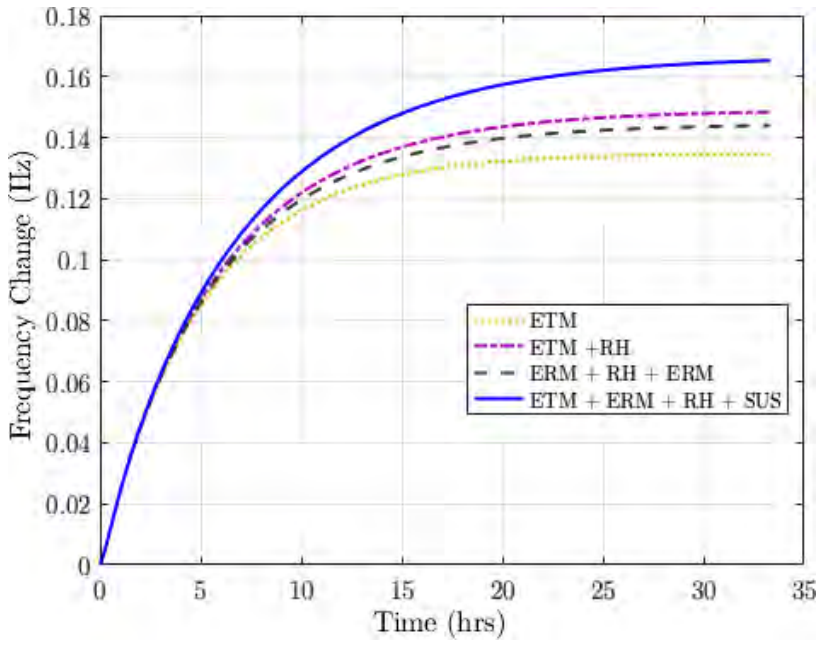

Figure 3.4: Modelled frequency shift over 33 hours of continuous laser heating of the ETM optic in COMSOL, and its response to adding more bodies to the solution.

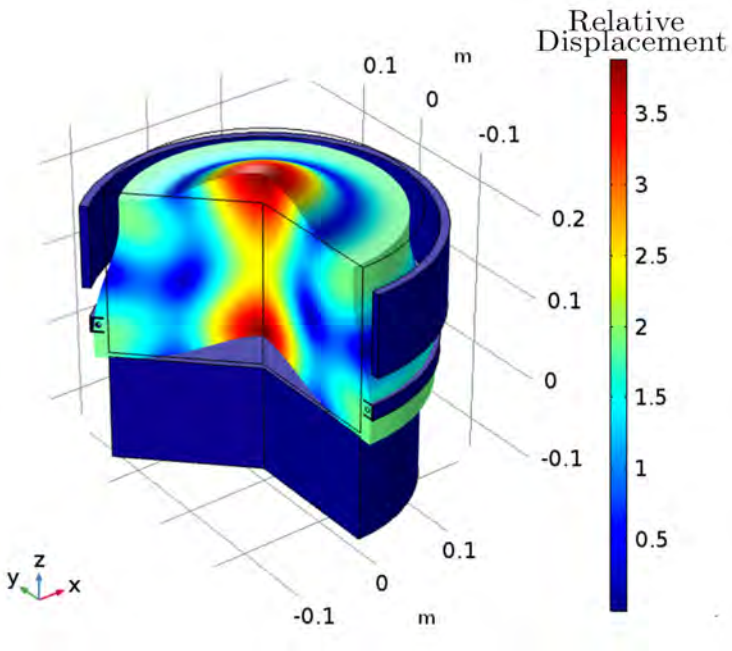

Figure 3.5: Modelled surface displacement of the 8129 Hz mechanical mode shape of the 2DAS ETM and its surroundings including the ERM, Ring Heater and approximated suspension rotated through 190°. Units of displacement are relative.

In each case, the modelled heating of the system increases exponentailly over the first 10 hrs, with the largest $\frac{dT}{dt}$ occurring over the first three hours. After 20 hrs, the model suggests that the ETM should then begin to reach thermal equilibrium. To take the results of this model and apply its predictions to the experimental temperature of the test mass, a function which describes the modelled temperature changes must be produced. Fitting the modelled curve which includes all geometry terms from COMSOL with a single term exponent with respect to time, allows the time constant and gain of the model to be obtained which can later be used to produce an adequate transfer function between temperature and frequency:

$$\frac{\Delta f}{\alpha_{\text{coat}}}(1 - \exp^{-\tau(t)}) + f, \tag{3.5}$$

where Δf is the total change in frequency in Hz, scaled to the coating absorption ($\alpha = 1.4$ ppm), τ is the exponential time constant, and f is the mechanical mode frequency at room temperature. To solve for the expected change in temperature of the 8126 Hz mechanical mode, taking into account the pre-installation coating absorption α_{coat} of 1.4 ppm Eq. 3.5 takes the following form:

$$\frac{0.1136}{0.14}(1 - e^{-0.159(t)}) + 8126. \tag{3.6}$$

From this, the modelled response for a test mass with $\alpha_{\text{coat}} = 1.4\,\text{ppm}$ is produced. This model then allows the change in resonant frequency of the coated test mass to be calculated as its temperature is increased. This 'modelled estimate' with constant optical absorption can then be compared to experimental measurements of frequency shift when the interferometer is in operation. Any differences between measured and modelled frequency shifts suggests that the heating of the experimental test mass differs from the model, implying the coating layer is absorbing more/less power than predicted.

3.3 Experimental Frequency Tracking

To accurately measure ETM resonances inside the interferometer, the detector must be operating in a 'low noise' state. Under these conditions, the automated state machine known as GUARDIAN (which is used for management and control of interferometer systems) has fully 'locked' all sections of the interferometer [165]. From here, the laser power inside the Fabry-Pérot cavities will increase from $<40\,\text{mW}$ to $100\,\text{kW}$ $\pm 5\%$ within $\sim 100\,\text{seconds}$. Using the time constant predicted using Eq. 3.6, it was shown that it takes $\sim 30\,\text{hours}$ for an ETM to reach thermal equilibrium from a 'cold state' with the largest changes in mechanical mode frequency occurring in the first 30 mins. If the system reaches thermal stability, Δf becomes dominated by environmental temperature effects from the surrounding vacuum system which are not taken into consideration in this model.

Due to a multitude of external effects such as seismic coupling or inadequate alignment, GUARDIAN does not always succeed in moving from a state of 'unlock' to a long-term stable lock. As the most significant frequency change will be observed in the first 30 mins when GUARDIAN transitions from a prolonged unlocked state to a fully locked state, logic must be defined and be used to identify regions of time over which data can be requested.

Figure 3.6 shows the GUARDIAN lock state of LIGO Livingston over 22 hours, starting at the 11^{th} of July 2017. This program returns the current detector state as

a value between 0 and 2000, allowing the lock state of the interferometer to be tracked. In the case where the GUARDIAN state >1800, a full 45 W beam is incident on the Fabry-Pérot cavities, and the detector is considered to be in 'low noise' state where the system is stable and gravitational wave measurements can be taken. As was shown previously with COMSOL in Figure 3.4, when the laser is activated, the resonant modes of the optic will begin to shift following the α_{coat}. Regions of GUARDIAN data which meet this condition are then used as timestamps which can be used to request data from other interferometer channels. The output mode cleaner (OMC) is the final optical system through which the primary IFO sensing laser passes before reaching the output photodetector. Using pick-off photo-diodes the circulating laser power inside the OMC can be monitored after passing through the interferometer, allowing the resonant frequencies to be tracked. This data is available to all operators of the interferometer and this method of analysis has been attempted previously. However, due to inadequate frequency tracking this has yielded mixed results. As discussed in the next setion, the improved frequency tracking methods implemented by the author, which take into account COMSOL predictions should improve this process.

So far the process used to estimate the 'locked absorption' of a coated test mass has been described. However, by trying to predict how the test mass temperature changes when the interferometer loses lock, greater insight into coating absorption over long timescales can be gained. When the interferometer then loses lock there is no laser light circulating inside the interferometer, and thus no data to track at the OMC. During these times results from the COMSOL model suggest show that the optic and its surroundings will begin to dissipate heat, at nominally the same rate which the heat was applied following Eq.3.5. Modelling the heating and cooling of the test mass in COMSOL showed that both the heating and cooling states are inverse but equal exponential trends, with the modelled frequency decreasing back to its initial value when the system reaches thermal equilibrium. This then produces two expressions which describe the thermal state of the interferometer during operation which can be used to provide a modelled estimate of its thermal state. This prediction of frequency shift and absorption estimation as the detector gains and loses lock has not been attempted before and should allow any changes in the absorption of coating absorption to be estimated over prolonged periods of detector operation.

As this method of absorption analysis requires historical interferometer data and cannot currently be run in real-time, 30 day stretches of GUARDIAN data were analysed. Fig 3.6 shows an example of the changes in GUARDIAN state starting on July 11th 2017. If the IFO has moved from an unlocked state (GUARDIAN = 0) to a 'stable lock' (GUARDIAN>1800) and has remained stable for more than 2 hours the beginning of this region is marked with a red diamond, and an increase in test mass frequency is expected. In the periods where the interferometer is locked, there is a continuous signal from the output mode cleaner (OMC) (see Chapter 1).

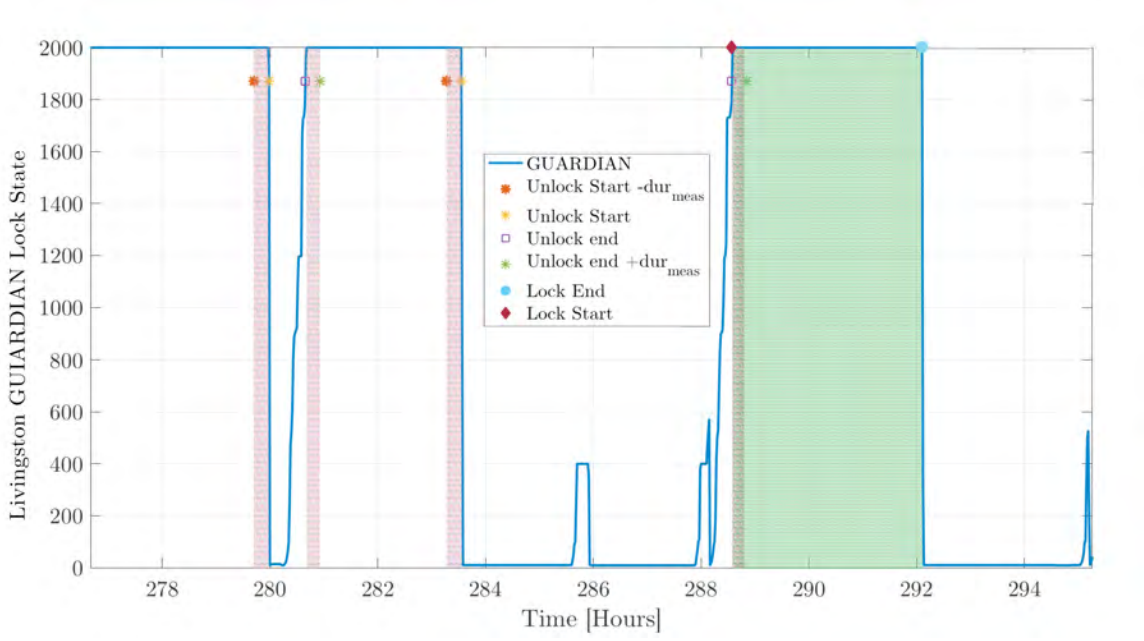

Figure 3.6: Plot of the L1 GUARDIAN-State starting from the 11^{th} of July 2017 and ending 22 hours later. Green denotes regions where 'stable lock' data has been identified, Pink identifies regions 1000 s before and 1000 s after the detector has lost lock. Dark green identifies a region where both 'stable lock' and 'stable unlock' logic occurs at the same time.

If the system has remained locked for more than 2 hours, then loses lock and does not regain lock for more than 1 hour, the cooling trend of the system can also be predicted as the cooling curve of the test mass is expected to follow the inverse of the trend denoted by Eq.,3.5. These regions are denoted as blue circles in Figure 3.6. If the interferometer loses lock there is no longer laser light circulating in the interferometer and data from the OMC ceases. Thus, the cooling trend of the optic must be calculated from the previous known state of the optic. Regions of data for which both the 'good lock' and 'good un-lock' conditions are satisfied are

highlighted in green and the entire region of data will be used.

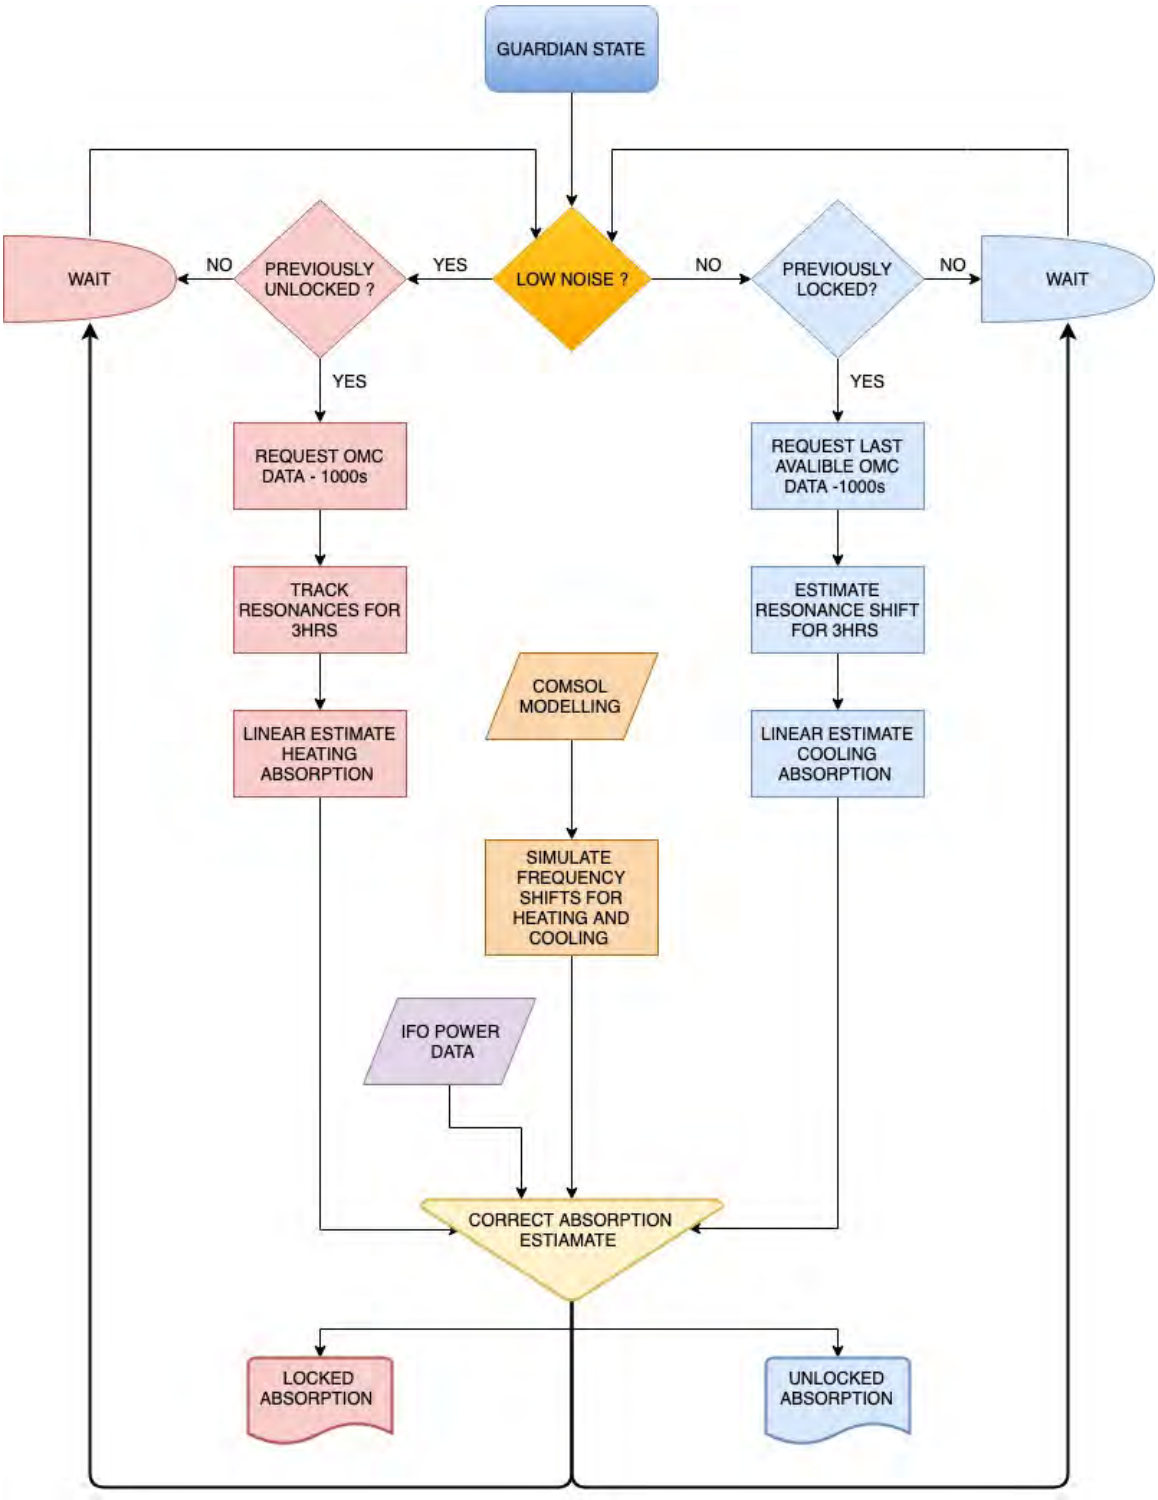

Figure 3.7: Flow chart describing the processes from which an absorption estimate can be produced.

On average the total time taken to go from an unlocked state to 'low-noise' state is ≈100 seconds, where bounds of ±1000 s are taken on all qualifying sets of data (highlighted in pink). This removes regions of data where the interferometer has not reached a 'low noise' state such as the region around 288 hours from the final data-set. These attempts to achieve the 'low noise' state can contain sporadic heating and cooling of the test mass as different systems are powered off and on by GUARDIAN. These periods can cause the temperature of the test mass to deviate in a manner difficult to predict with the COMSOL modelled system.

Two different FEA models were used in order to match measured frequencies recorded in OMC data to its parent optic. As the 2DAS model shown in Fig 3.4 is only able to represent mechanical modes which are rotationally symmetric, this model could only be used to predict the fundamental drum-head (0,1) mode [166] of each optic at 8.2 kHz. A 3D representation of the test mass geometry was produced allowing mechanical frequencies between 5 kHz and 6.1 kHz to be identified in the output of the OMC. The results of this model were benchmarked against a similar model produced by C. Blair [5], with the predicted frequencies the measured values. The frequencies for each optic are listed below in Table 3.2. No mechanical modes above 8.2 kHz could be identified as this is bordering on the edge of several digital filters placed on the OMC channel used to reduce other noise sources in the output signal.

Optic	Mode 1 (Hz)	Mode 2 (Hz)	Mode 3 (Hz)
ITM-X	5937.2	6039.0	8160.0
ITM-Y	5939.1	6043.1	8161.4
ETM-X	5948.4	6053.2	8158.3
ETM-Y	5946.3	6050.4	8155.6

Table 3.2: Measured eigenfrequencies of end test mass optics in OMC output. Relative surface deformations of each mode shape are shown in Fig 3.12

3.3.0.1 Stable-Lock Absorption

For regions of data which meet 'stable lock' conditions, the signal from the OMC is analysed, tracking the 12 frequencies listed in Table 3.2 for the time that the detector

remains locked. As the shifts in these frequencies are small (maximum 0.17 Hz for 33 hours of a continuous stable lock) the power spectral density of the output data can be windowed, placing $\pm$ 0.5 Hz bounds on the region for the estimated frequency. The resonant frequency then appears as a peak above the mean noise in these regions. An example of each mode frequency measured in the OMC is shown below in Fig 3.8.

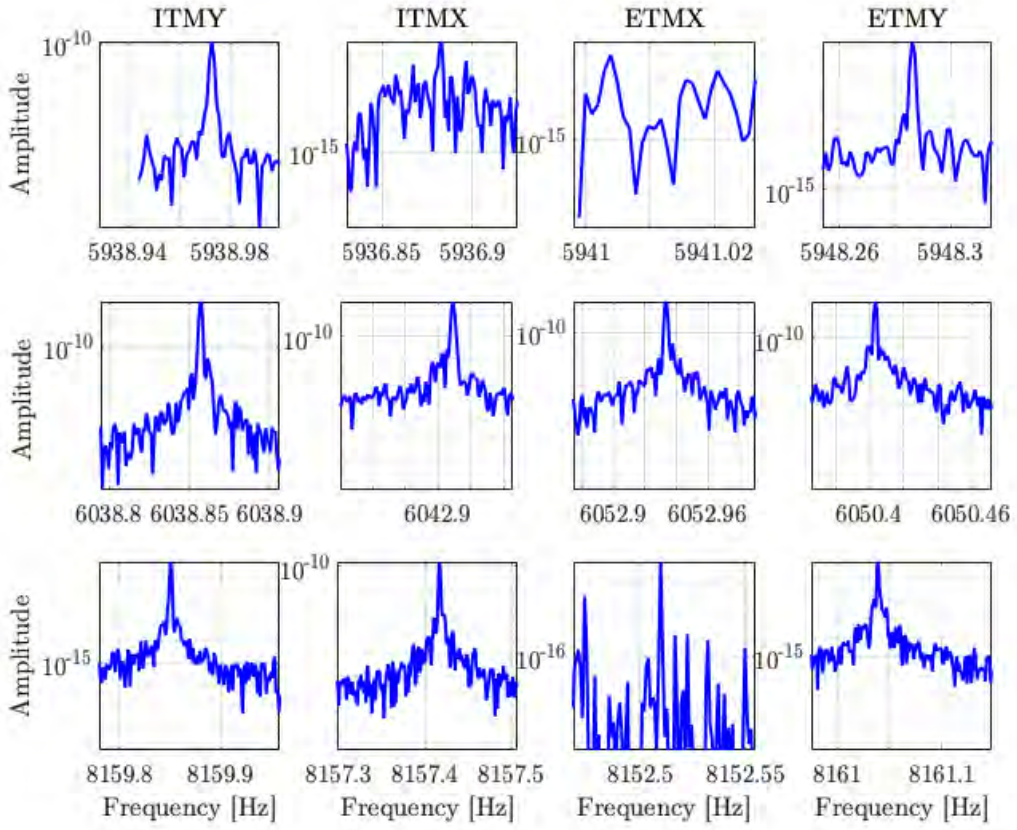

Figure 3.8: Windowed power spectral density of the OMC output showing resonant frequencies of ITMs and ETMs in Hz, concurrent with those listed in Table 3.2.

Each of these frequencies is tracked for the duration of the stable lock period, with the rate of change in mode frequency calculated for the first two hours. Over these two hours, test mass heating is dominated by coating absorption α_{coat}, and temperature fluctuations due to re-radiation and environmental effects do not need to be considered. This tracking method produces a linear trend of frequency shift as a function of temperature which can be compared to modelled predictions. Due to the low signal-to-noise ratio (SNR) of the resonant peaks, other side-band frequencies present in the OMC can sometimes be confused for the resonance,

leading to an incorrect prediction of its frequency.

Since COMSOL model predictions suggest that each frequency should follow a similarly increasing trend over the two hours, the expected $\Delta f/f$ for each frequency in the form of Eq. 3.5 is factored into a more robust method of frequency tracking. This method, developed by the author and C. Blair, takes into account the frequencies predicted from COMSOL and how they are expected to shift within the two hours. This approach employs necessary data windowing and filtering which allows the frequencies of each mode to be more reliably tracked over the two hours. An example of this method is shown in Fig 3.9. Fluctuations in peak tracking then introduce a source of error in the frequency measurement and can influence the overall calculated gradient.

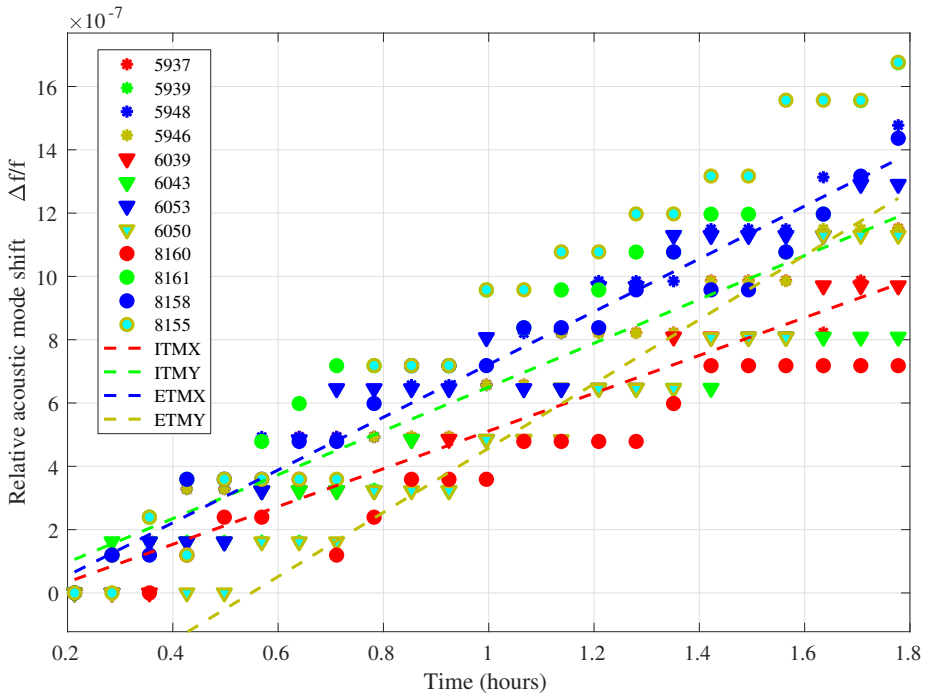

Figure 3.9: Frequency tracking over the same time duration applying the COMSOL correction.

The data in Fig 3.9 is a direct measurement of Δf, which contains information about the temperature of the test mass, and is directly related to how much power was absorbed into the optic. Therefore these measurements then become the basis for what will be referred to as the 'measured absorption estimate', as this comes directly from IFO data.

3.3.0.2 Estimation of Unlocked Frequency Shift

Frequency tracking only provides estimates of the state of each test mass when the interferometer is locked. As the goal of this experiment is to determine whether coating absorption changes over long periods of time, estimates of the absorption of the coating when the interferometer becomes unlocked are also made. As discussed in section 3.3, the modelled cooling trend of the test mass follows an negative exponential decay - mirroring the trend observed when heat is applied to the modelled optic. For the assumption that the heating and cooling rates of the optic are symmetric, the timescales over which the unlocked frequencies can be predicted in this manner must be ≤ 30 minutes. Over this time the exponential decay in temperature of the ETM can be approximated in this manner. This allows the change in frequency, and thus a change in the absorbed power of a test mass optic to be inferred from these predictions when the interferometer loses a stable lock.

However, a different algorithm is required to identify segments of the data where the interferometer is in a unlocked state. A 22-hour segment is taken from the GUARDIAN data presented in Fig 3.6. During this time, the interferometer was fully unlocked for $\sim$ 45 mins. Throughout locked periods, there is continuous heating of the test mass and available data in the OMC for frequency tracking. As the time constant for the heating and cooling of the optics can be considered nominally identical, when the interferometer moves from a locked to an unlocked state at 280 hours it will begin to dissipate heat and follow the trend predicted in Eq 3.6. The model allows for cooling of the test mass; an estimate of the total residual heat stored in the test mass can be calculated, and an estimate of the change in absorption can therefore be obtained from modelled predictions.

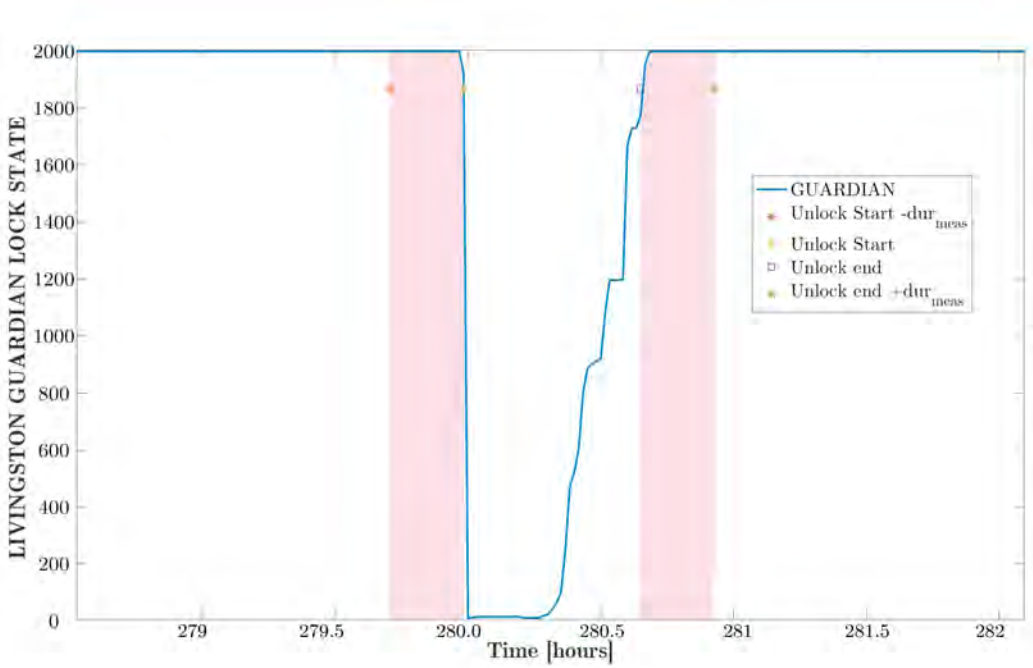

Figure 3.10: Scaled plot of GUARDIAN state shown in Figure 3.6 showing pink regions ± 1000s before and after lock loss where frequency tracking data can be acquired.

An interpolation of the decrease in temperature can be produced if the mechanical mode frequency before and after the interferometer loses/gains lock is known, with ±1000 seconds of data sampled on either side of the unlocked period.

3.4 Estimation of Coating Absorption

To monitor any changes in the detector absorption, the model of heating and cooling predicted from COMSOL can be combined with the current power data inside the interferometer. If the mechanical modes of the optics are also tracked over the same time, this can be used to produce a 'measured' absorption estimate using the methods described in Section 3.3. Given that the amount of heat absorbed into the test mass is proportional to laser power, a constant absorption value of 1.4 ppm will produce frequency shifts which directly correlate with COMSOL predictions. This allows a direct comparison between the constant absorption value modelled in COMSOL and the results produced through the measured absorption trends.

As the total input and circulating power in each arm is recorded by different

subsystems, the circulating laser power inside the interferometer can also be taken into consideration in these calculations. As the system contains pick-off and optical transmission monitors, this data can be used to predict the total interferometer power with installation spec absorption, where

$$P_{\text{sim}} = \alpha_{\text{coat}} \left(\overline{\frac{P^{IM4}_{\text{trans}}}{P^{\text{x-arm}}_{\text{trans}}}} \right) \frac{1}{k} \tag{3.7}$$

and P^{x-arm}_{trans} is the transmitted light power measured inside the X-arm cavity and P^{IM4} is the transmitted power measured on a steering mirror after the input mode cleaner (IMC) [94]. These values are averaged over the time of interest, and k is a scaling factor accounting for the total input laser power and cavity gain.

Power in different subsystems can fluctuate $\pm 5\%$ and affect the final calculated absorption. For each lock classified as 'stable' a zero-pole-gain (ZPK) transfer function was used to predict the level of absorption. The same ZPK method can be used to infer the unlocked absorption. Taking the gradient of the tracked frequencies before and after lock loss, as detailed in Section 3.3.0.2, the estimated frequency response at these times which should follow a linear approximation from the change in test mass temperature.

The total laser power inside the interferometer was analysed using Eq. 3.7. Using the power data, in conjunction with Eq. 3.6, allows a full estimation of frequency shift for the test mass optic to be calculated assuming a constant absorption value. The model predicts the rise and fall of frequency when the interferometer gains and loses lock respectively; consistent with COMSOL predictions. A caveat should be added to this assumption; in Figure 3.10 there is a $\sim$20 minute period where the interferometer is powering up. The power inside each cavity is increasing, but no mechanical mode frequency data is available at the OMC output during this time. This means that there is a small period with circulating power on the ETM optics that will begin to increase the temperature of the system, that is not being taken into account. This will be corrected at a later date.

Taking account of the power data from the OMC, the transmitted power in both ETM-X and ETM-Y, and the monitored gain of the Power Recycling Cavity (PRC) allows a simulation of the frequency response for the optic to be produced. This

simulation is not dependent on measured frequency tracking data but follows the gain and total modelled frequency shift produced from COMSOL. It should be noted that terms which are succeeded by the suffix 'sim' have been simulated in this manner. The results of the simulated frequency shift Δf_{sim} are shown below in Fig 3.11.

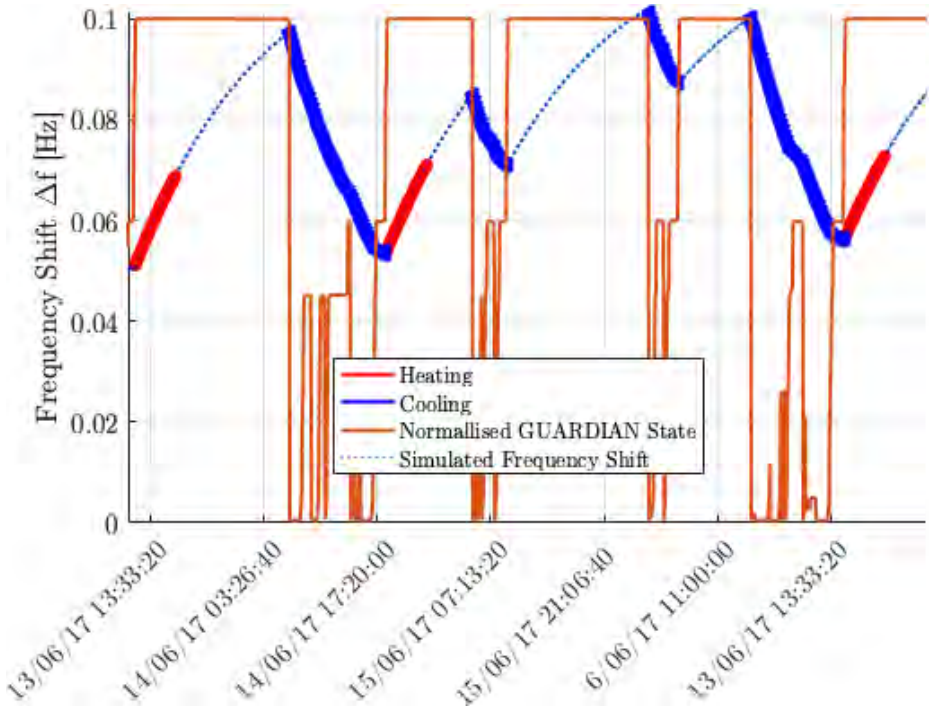

Figure 3.11: Simulated change in frequency Δf_{sim} for several lock cycles between 13-16$^{\text{th}}$ June 2017 accounting for fluctuations in IFO power. Regions which can be further analysed are marked as heating or cooling sections which are calculated using Eq. 3.6. The GUARDIAN lock state has been normalised as a guide to the eye, showing where the heating is being applied to the optic

At this stage, there is now an estimate for frequency shift Δf_{sim} and measured frequency shifts Δf_{meas} for discrete time sections throughout the same period. To make a direct comparison between Δf_{sim} and Δf_{meas}, the same lock stretch identification logic shown in Figure 3.6 is used to identify regions where data exists in both methods. In Fig 3.11 these regions are marked as heating or cooling, and can be used to estimate absorption due to the coating α_{coat}. By comparing the respective gradients of measured and simulated frequency shifts :

$$\alpha_{\text{coat}} = \left(\frac{d\,\Delta f_{\text{meas}}}{dt} \Big/ \frac{d\,\Delta f_{\text{sim}}}{dt} \right) \times P_{\text{sim}}, \tag{3.8}$$

the ratio of frequency gradients can be scaled by P_{sim}, and the IFO power absorbed

by the test mass (see Eq. 3.7), to estimate coating absorption for each lock/unlocked state.

The simulated absorption trends shown in Fig 3.11 were then compared to the COMSOL calculated frequency shift. It was found that the transfer function can correctly replicate the frequency shift of a given test mass over several lock/ unlock cycles. To accurately gauge if there is any change in coating absorption over the course of an observing run, this method could be used to provide a constant coating absorption comparison to the measured frequency shifts of each test mass.

Absorption Predictions for Resonant Modes

Due to the time constraints associated with acquiring the amount of data needed, this method has only been tested between January and September of 2017. Figures 3.14, 3.16, and 3.18 show the optical absorption calculated during locked and unlocked states of end test mass optics at LIGO Livingston during this time while Figures 3.13, 3.15 and 3.17 show results from the input test mass optics . These are discussed for each set of mechanical modes which were analysed. The resultant deformation of each mode shape is shown in Fig 3.12.

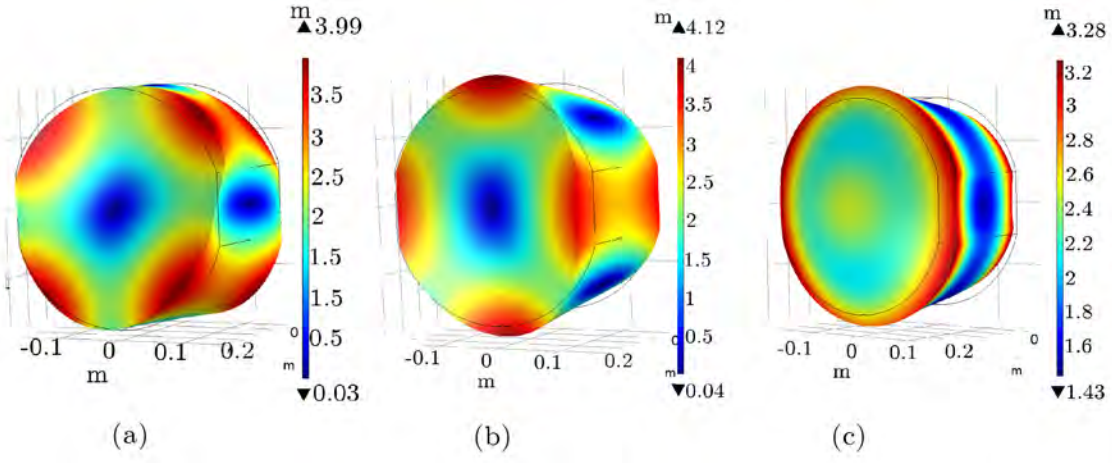

Figure 3.12: Calculated total surface displacement of (a) 5.9 kHz, (b)6 kHz and (c)8.2 kHz mechanical modes using FEA analysis. Regions shown in blue denote regions with minimum/no displacement and red regions are points of maximum displacement.

A review of the data for each mode reveals several key observations over eight months. At the beginning of the year, between January and March, a large spread

in the data was measured. It is speculated that it could be caused by a large number of short lock/unlock cycles. This variation in power could cause an increase in the noise in the predicted absorption of the test masses as the total power inside the interferometer is continually changing. This hypothesis is then strengthened by a number of short lock cycles also being reported in the L1 summary pages between January and March 2017.

No predictions were produced between May and June of 2017, as the interferometer was taken down for inspection and repairs [167]. In this time the Horizontal Access Modules (HAMs) [168], which house the test masses were opened for maintenance.

Under the initial assumption that the coating absorption should remain constant, one can presume to first order, a constant value for the system over the length of the measurement. This allows a least-squares regression to be applied to the data to reveal any underlying trends. In order to mitigate the spread produced by short lock cycles or inadequate frequency tracking, points which are $>1.5\sigma$ from the mean have been omitted.

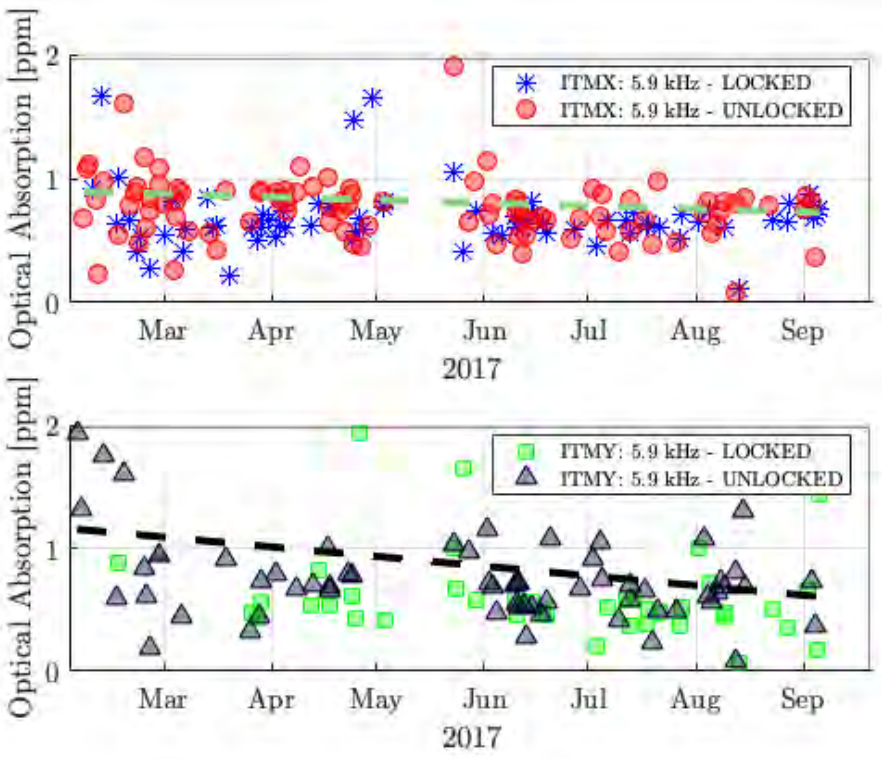

Figure 3.13: Calculated optical absorption of input test masses during locked and unlocked states of test mass optics at LIGO Livingston between January and September of 2017 at 5.9 kHz.

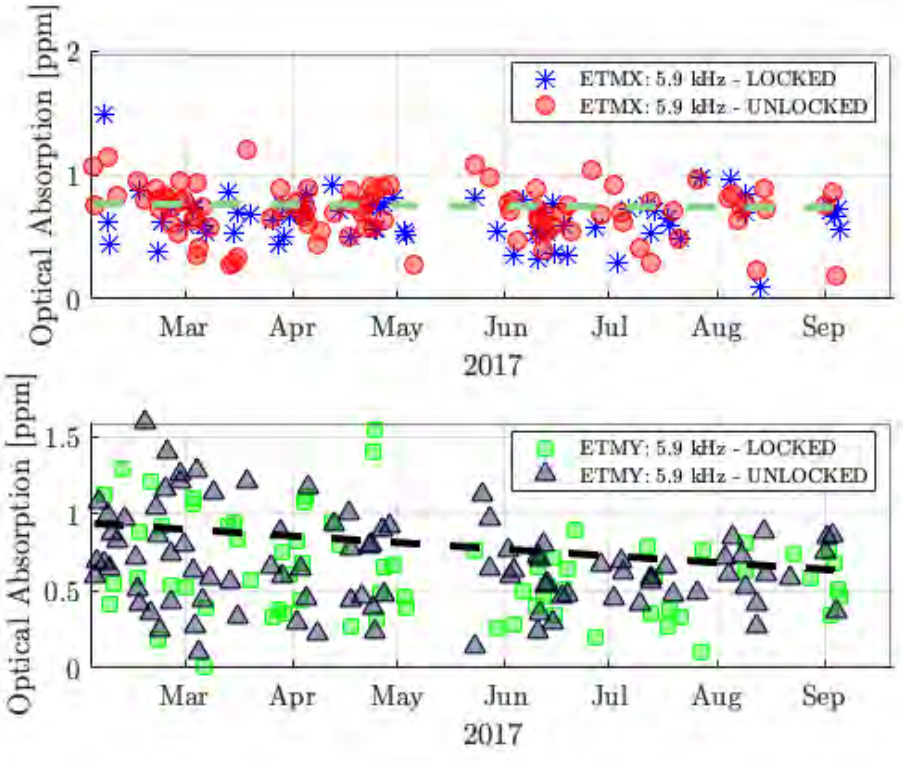

Figure 3.14: Calculated optical absorption of end test masses during locked and unlocked states of test mass optics at LIGO Livingston between January and September of 2017 at 5.9 kHz.

Given that the ITMs and ETMs in both arm should have the same mechanical and optical responses, if the trends in calculated absorption for similar optics do not exhibit the same tendency, changes in absorption can be localised.

Examining the calculated absorption response for the 5 kHz mode across all optics shows similar responses for each detector arm. As seen in Tab. 3.4, the mean absorption of both ITM-X and ETM-X optics are comparable, where the quoted error is the standard deviation throughout the measurement. This shows no signs of increase after the chambers were opened in May 2017. Inspection of the inferred coating absorption from the Y-arm optics shows a relatively large decrease over this time, with a maximum absorption of 1 ppm at the start of the measurement decreasing by 26% over the run. Comparing this to the calculated absorption on the Y-arm optics: ITM-Y at 5.9 kHz shows a similar value of 0.78 ± 0.21 ppm, following a similar trend to the values observed in ITM-X. ETM-Y shows a lower mean coating absorption of 0.71 ± 0.08 ppm, with the lowest standard deviation in measurements at this frequency. A breakdown of $\overline{\alpha}$ for each optic and frequency is detailed below in Table 3.4

Optic	Frequency [kHz]	$\overline{\alpha}$ [ppm]	Spread [ppm]
ITM-X	5.9	0.77	0.13
ITM-Y	5.9	0.78	0.21
ETM-X	5.9	0.78	0.21
ETM-Y	5.9	0.71	0.08

Table 3.3: Summary of mean calculated optical absorption from Fabry-Pérot cavity optics using their respective 5.9 kHz mechanical resonance between February and September 2017.

If instead one assumes that the similarities in coating absorption are a systematic effect of fitting to the data, then it would be more useful to compare the difference in fitted trends between sets of optics. Using the same methods described in section 3.2, the change in frequency for other mechanical modes can also be calculated. Focusing on the mechanical modes of the ETM alone, the total normalised frequency shift ($\Delta f/f$) was produced for each 2DAS mode up to 20 kHz. This shows that the expected ($\Delta f/f$) has a near-linear trend with increasing

resonance frequency. Extrapolating this trend gives an estimate of the expected response of lower eigenfrequencies to temperature, allowing non-2DAS modes to be measured. The absorption of the coating can be estimated from the measurements of each mode using Eq. 3.8. If the level of the coating absorption measured does not match the level produced by COMSOL, then the model assumed estimate could be inaccurate or there could be additional factors which can affect the optical absorption of the test masses which are not included in the FEA model.

Optic	Frequency [kHz]	$\overline{\alpha}$ [ppm]	Spread [ppm]
ITM-X	6.0	0.78	0.11
	8.2	0.38	0.07
ITM-Y	6.0	0.68	0.08
	8.2	0.93	0.11
ETM-X	6.0	0.73	0.06
	8.2	0.78	0.05
ETM-Y	6.0	0.70	0.08
	8.2	0.91	0.13

Table 3.4: Summary of calculated optical absorption from Fabry-Pérot cavity optics using their respective 6 kHz and 8.2 kHz mechanical resonances between February and September 2017..

Reviewing the results for the second mode (6 kHz, see Fig 3.12(b) for mode shape) shown in Figures 3.15, 3.18, there is again a large spread in the predicted absorption values both optics. However, the region with the highest concentration of calculated points in both optics again shows a large overlap between both the X and Y arm optics with the largest density of points between 0.5 ppm and 1 ppm. As the mechanical mode-shape of the 5.9 kHz and 6 kHz modes both have a node of minimal motion in the centre of the optic, if the absorption across the optics surface is homogeneous, there is a minimal change in the interaction between the incoming laser light and mechanical vibrations of the test mass.

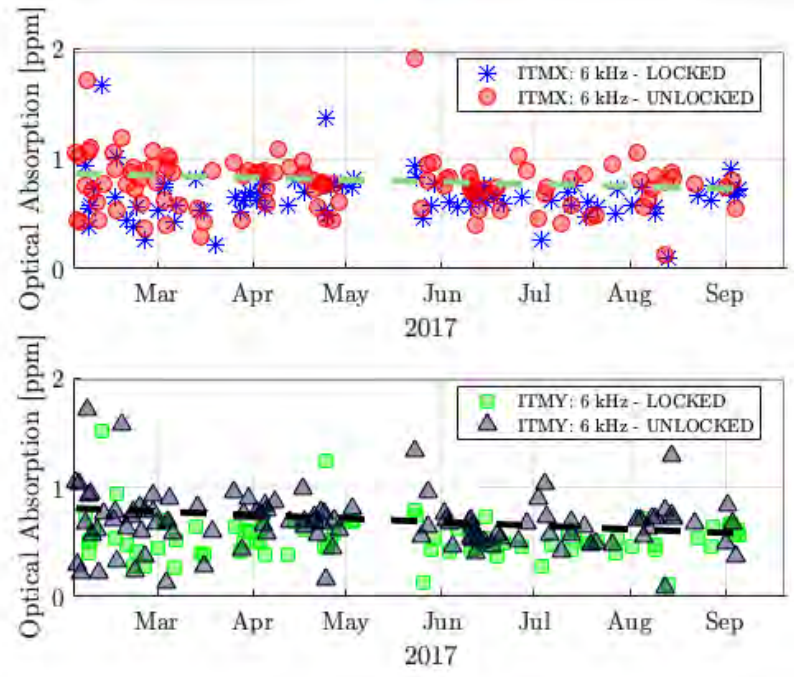

Figure 3.15: Calculated optical absorption during locked and unlocked states of input test mass optics at LIGO Livingston between January and September of 2017. Showing the calculated absorption from the response of the 6 kHz mechanical mode.

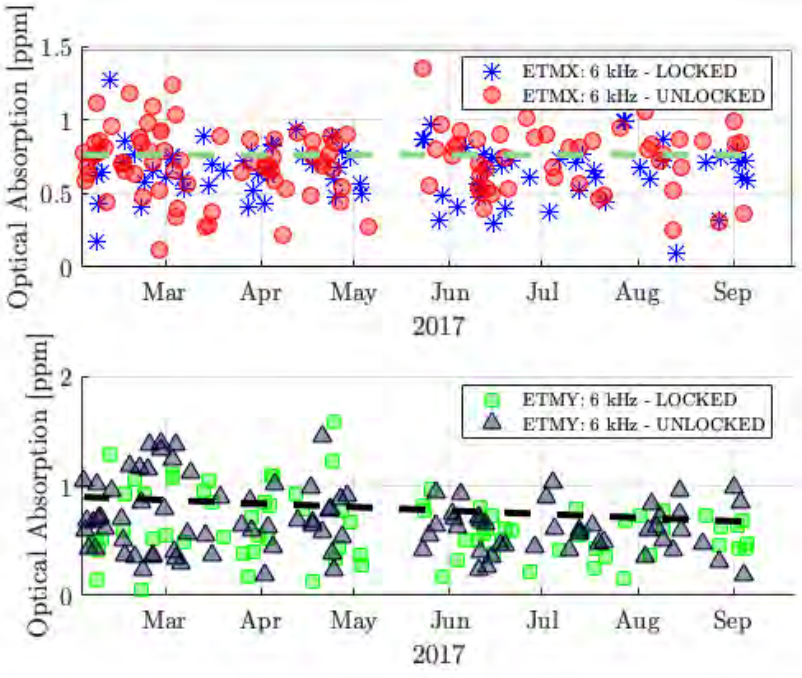

Figure 3.16: Calculated optical absorption during locked and unlocked states of end test mass optics at LIGO Livingston between January and September of 2017. Showing the calculated absorption from the response of the 6 kHz mechanical mode.

However, the higher frequency 8 kHz mode, shown in Fig 3.12(c) seems to exhibit a different trend in ITM-X from the other optics, with a measured absorption ∼0.3 ppm compared to the consistently higher absorption in the other three optics. When the noise on the data is considered the measured absorption of ITM-X at 8.2 kHz is 75% lower than that measured on ITM-Y at the beginning of the measured data. As time increases, the mean absorption values of both optics continue to diverge, whereby the end of the measured data, the mean values differ by 99.6%. As the same analysis on the 8.2 kHz modes of the ETMs over the same periods does not exhibit the same gap in predicted absorption, it suggests that effect is centralised to ITM-X.

FEA analysis of the mechanical mode, shown in Fig 3.12(c), reveals that the surface deformation of this mode has a much higher percentage of energy stored in bulk motion. The surface deformation of such a mode manifests as a volumetric change of the geometry, normal to the path of the IFO beam. As the IFO beam will have larger interactions with such an oscillation, if any degrading points of coating material lie within this region, this effect would provide larger interactions with a mode such as this. It could also indicate that the main IFO beam is not directly centred on the optic. Any small deviations of beam position on the optic would be exaggerated by the large surface deformations of this mechanical mode, which could manifest as an anomalously low absorption measurement.

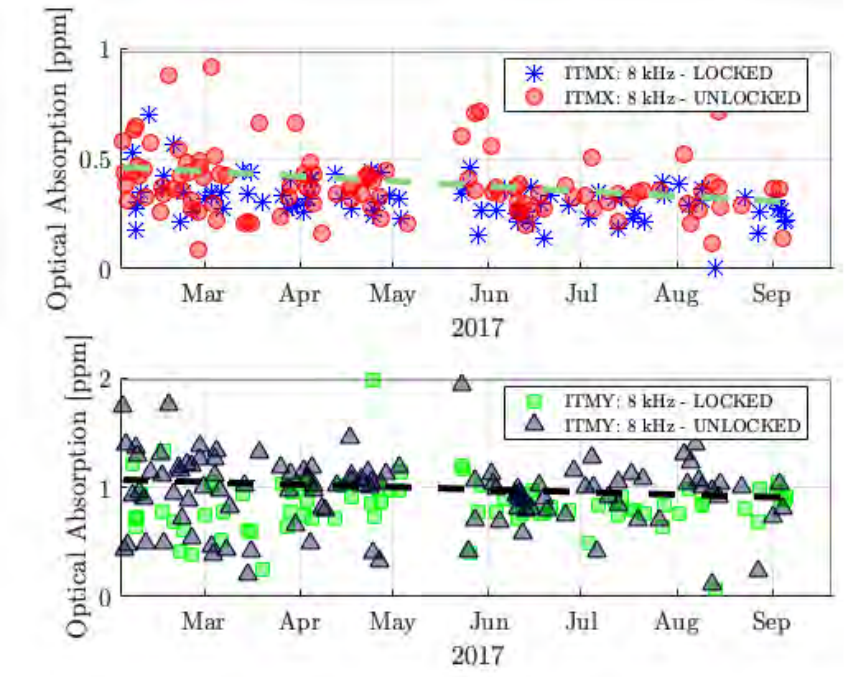

Figure 3.17: Calculated optical absorption during locked and unlocked states of input test mass optics at LIGO Livingston between January and September of 2017. Showing the calculated absorption from the response of the 8.2 kHz mechanical mode.

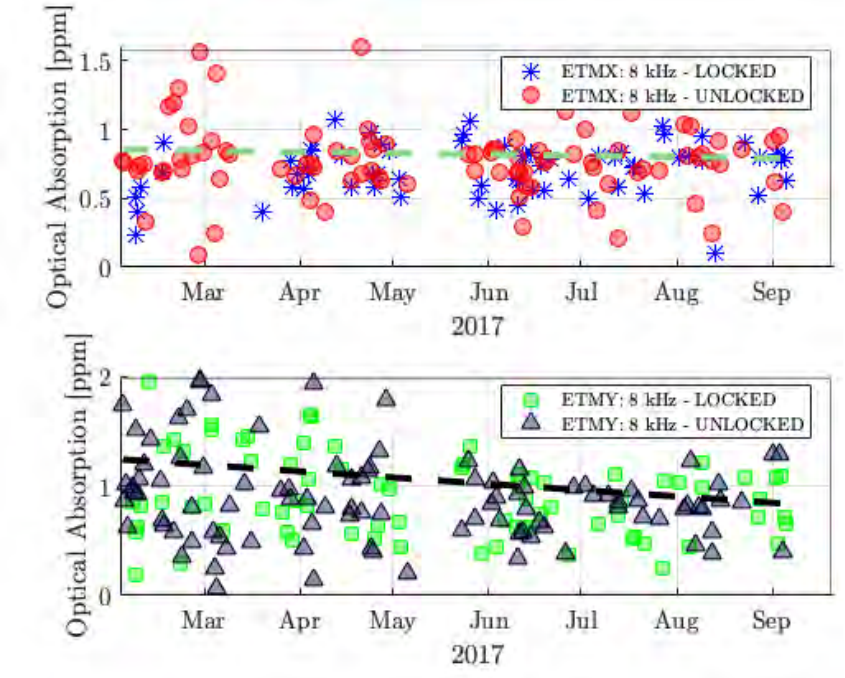

Figure 3.18: Calculated optical absorption during locked and unlocked states of end test mass optics at LIGO Livingston between January and September of 2017. Showing the calculated absorption from the response of the 8.2 kHz mechanical mode.

Inferred Changes in Coating Absorption for each Test Mass

To gain insight into the average trend of calculated absorption values over this time, the coating absorption of each mechanical mode is considered with respect to each optic. The calculated average absorption for X-arm optics, ETM-X and ITM-X shows that the decreasing trend observed in ITM-X is also apparent in ETM-X at 8.2 kHz at a higher level. Comparing the trend in average absorption for 5.9 kHz and 6 kHz modes show level or decreasing absorption across the measurement duration. It is interesting to observe that the average decrease measured on the 8.2 kHz modes for both optics show the same gradient, without measured offset shown on ITM-X. This suggests that the difference in absorption shown on ITM-X could be affecting other parts of the detector, and is not isolated to one optic.

The calculated average absorption of ITM-X, ITM-Y and ETM-X show consistent increasing and decreasing trends from month to month, reaching a maximum absorption in May 2017. However, if the average absorption for all mechanical modes on a given optic is considered, the observed trend for ETM-Y continues off-trend from all other optics. Between May and July of 2017 where the opening of HAM chambers temporarily increased the measured absorption of each test mass, we see that ETM-Y was not affected in the same respected but rather shows an average decrease in absorption. This trend continued until July where the optic then saw a steadying increase in absorption, whereas all other optics in the detector started to decrease.

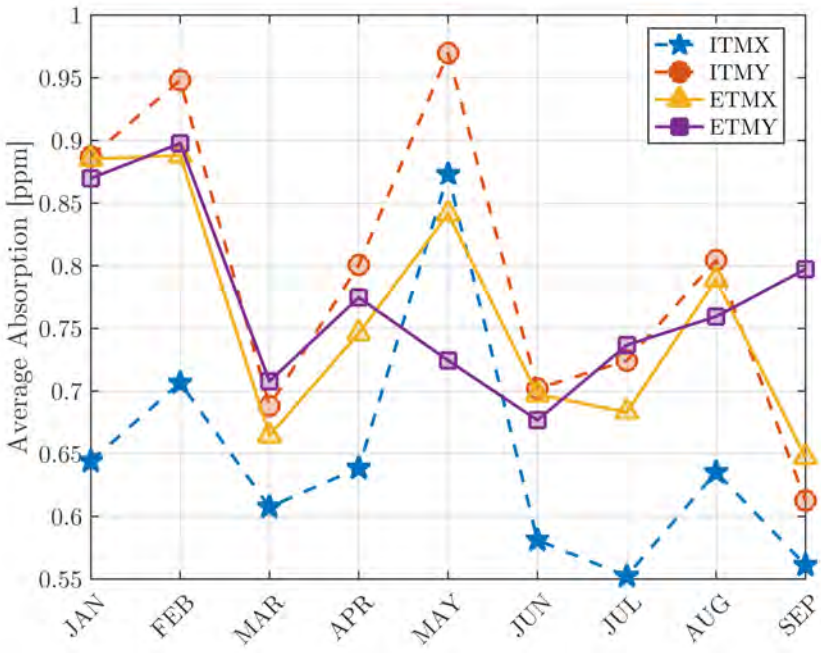

Figure 3.19: Average measured coating absorption from Fabry-Pérot cavity optics from all tracked mechanical modes between Jan - Sep 2017.

In each case, the calculated coating absorption for each optic shows significant variations within a given month due to inadequate measured frequency tracking, manifesting in large error bars when viewed on a monthly average basis. Shown in Fig 3.19 this analysis predicts that there is an overall decrease in the absorption of each optic between January - September of 2017. The calculated absorption of each test mass is at least 44% lower than the best predictions produced by the HWS, with the upper limit of ITM-Y calculated at 0.96 ppm compared the HWS estimated 1.4 ppm. These calculated values show a mean absorption of ITM-Y: 0.79 ± 0.07 ppm, ETM-X 0.76 ± 0.20 ppm and ETM-Y 0.77 ± 0.33 ppm compared with the ITM-X which due to the lower absorbing 8.2 kHz mode produced ITM-X 0.64 ± 0.18 ppm. All of these values measured higher than the 0.5 ppm required by LIGO for optical coating absorption at 1064 nm, however, the average values are subject to frequency tracking accuracy.

Given that three out of four optics in the interferometer are all showing similar trends, the measured changes in coating absorption, which should be consistent between each mirror, must be considered. As the calculations used to provide an

absorption estimate from each mechanical mode frequency factor in the total power fluctuating in the IFO during each measurement, these should be accounted for in Eq. 3.7 and are not of consequence. As these measurements were taken over nine months, the outside temperature of the LIGO facility fluctuates, increasing cooling power in the summer months and increasing heating in the winter to compensate for a change in outside temperature of up to 20 °C. While environmental temperature effects should not dominate the measurements of frequency shift for each locking period of the IFO, it cannot be ignored that the outside temperature of the IFO could be responsible for the measured absorption trend shown here. However, while the general decreasing absorption trend shows similarities to the measured fluctuations outside the IFO, changes measured between May-July are significantly above trend for this data suggesting a significant difference in test mass temperature susceptibility.

To use this method as a means of detecting coating defects, such as point-absorbers [159], the method used for frequency tracking of mechanical modes requires improvement, such that the measurement uncertainty is reduced. The temperature fluctuations of environmental factors also warrant further investigation as artefacts of outside temperature changes is shown in the predicted decreasing absorption of each test mass shown in Fig 3.19. For the Livingston detector, the first instance of a region of defective coating reported on a test mass optic was on ITM-Y and ETM-Y in Oct - 2018, outwith the length of this data-set.

Once improvements have been made, any trend in the measured coating absorption can be discerned with a much higher degree of confidence. As the development of such an absorption monitor is still in the early stages, any deviation in the peak tracking as was shown in Figure 3.10 will produce an overestimation in the absorption trend. Post-processing and data cleaning can be employed to remove such effects; however, this does not rectify the underlying problem. Work is currently underway to implement a multi-frequency line tracking software, created by Edward Daw from the University of Sheffield to remove such a systematic error in the absorption estimate. Initially designed to track the change in frequency of suspension violin modes in the detectors DARM output, if such a tool were to be employed when calculating absorption for the previous observing runs, any trend in the absorption, by extension coating degradation would be viewed directly. With

improvements in line tracking, the related uncertainty produced from this analysis should hopefully provide an increase in the accuracy of the final measured absorption trends.

3.5 Conclusions

The analysis shows that the frequency change of radially symmetric eigenfrequencies of test mass optics at LIGO Livingston can be used to interpret the coating absorption on each test mass. Using FEA, the heating and cooling of a test mass was predicted, allowing the expected trends in frequency to be estimated during lock acquisition and after the detector had become unstable, losing lock. Three eigenfrequencies of each test mass were tracked over nine months, calculating the absorption of each optic during lock acquisition and lock loss using primitive frequency tracking techniques. Comparing the results for the same modes of each mirror allows their response to be compared.

By comparing the changes in optical absorption from three different mechanical mode frequencies over a nine month timescale, different effects were observed. Lower frequencies such as those measured at 5.9 kHz show very similar trends on all optics, with higher frequencies such as those measured 6 kHz starting to show disagreement between masses. At 8.2 kHz, there is a large, statistically significant discrepancy in the calculated absorption for ITM-X when compared with other test masses which is not visible at lower frequencies. This could hint at detector misalignment or could also stem from higher absorbing regions on the test mass, which interact with different modes shape deformations.

While different trends in optical absorption were found for each analysed mode frequency, each optic shows an average decrease in coating absorption between January and September 2017. The absorption of each test mass also shows a clear increase in May 2017 which is correlated with the chamber housing each optic being opened for maintenance. The author suggests that this data then be compared with other monitoring systems in the detector, such as the Hartmann wavefront sensor to corroborate the measured change in absorption

However due to fluctuations in signal amplitude and the interference of side-band frequencies, the tracking software used in this analysis introduces a large spread into the final calculated absorption of values. It is planned that this research will be continued by T. Mistry (University of Sheffield) and C. Blair (University of Western Australia), by implementing better line tracking software to reduce the spread in calculated absorption and using techniques such as Kalman filtering to predict the movement of each mechanical mode based on its previous position.

Further attempts at modelling the expected frequency shift with localised heating, to replicate the effects of point absorbers on test mass coatings is being carried out by A. Brooks at Caltech. This modelling will give more details on how to interpret the change in temperature of a test mass optic with such a spatially dependent heating profile with respect to mode frequency.

It is planned to implement this technique after the installation of new test mass optics inside the aLIGO detector, and track coating absorption from the point of installation. This will give a much larger picture of how the absorption might change after being installed in the detector. This can then be compared to other detectors inside the LIGO/Virgo collaboration, with hopes that the technique can be used to predict the presence of point absorbers in the HR coatings.

Chapter 4

Optical Absorption of Thin Film Coatings

Measurements of gravitational waves require optical coatings with high reflectivity and low optical losses. Highly reflecting (HR) mirror coatings are composed of alternating layers of materials with high and low refractive indices. The reflectivity of a HR coating reflectivity increases with the number of pairs of layers used, and with the ratio in the refractive index between the two materials chosen.

Current requirements for Advanced LIGO (aLIGO) place an upper limit of 0.5 ppm on the optical absorption of each mirror coating at λ_{IFO} =1064 nm. Any laser light which is absorbed by the HR coating layers or substrate will impart heat into the test mass optic, changing its surface curvature (see Fig 4.1) [94, 158]. This deformation results in optical instabilities inside the detector, decreasing the detector duty cycle and limiting the detector's total sensitivity if left unchecked. The aLIGO end test mass (ETM) mirrors are coated in alternating layers of SiO_2 and Ti:Ta_2O_5 where each coating layer thickness is optimised to produce >99.999% reflectivity of λ_{IFO} (see Chapter 2). The absorption of current coatings is $\sim$0.25 parts per million (ppm) of the sensing laser light [169].

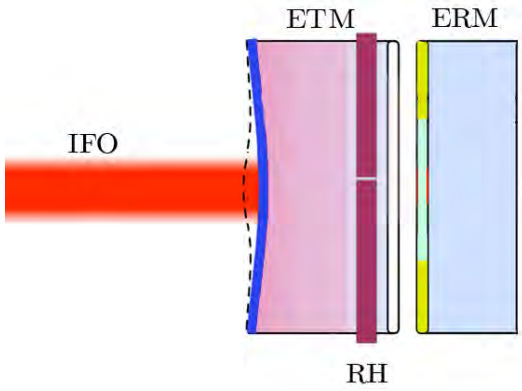

Figure 4.1: Representation of an aLIGO end test mass undergoing surface deformation due to 'self-heating'. Ring heater (RH) and End-Reaction Mass (ERM) also pictured.

As discussed in Chapters 2 and 7, the effects of thermal noise inside a gravitational wave detector can be reduced by operating future detectors at cryogenic temperatures. Under these conditions, the thermal stability of the test masses becomes a much deeper concern. In this section, measurement techniques for optical absorption of thin-film coating materials will be presented. From here, the underlying principles of high reflecting coating layers will be discussed. The applicability of any optical materials produced by different deposition methods will also be assessed, based on their potential improvements to gravitational wave detector sensitivity.

4.0.1 Photothermal Common-Path Interferometry

Photo-thermal common-path interferometry (PCI) is a highly-sensitive technique for measuring optical absorption of optically transparent materials [170]. A high power, pump beam with a small waist, w_{pump}, at a given wavelength is chopped at a known frequency [149], before passing through a sample where it creates a thermal lens due to the thermo-optic effect. The magnitude of this effect is dependant on the absorption, α, and the thermal properties of the material [158]. The pump laser is focused by a f=77.88 mm lens to a waist $w_{\text{pump}} < 35\,\mu\text{m}$ on the samples surface. As the pump light is chopped at a known frequency, the sample periodically heats and cools, creating a localised periodic change in the refractive index of the sample

$$\Delta n = \frac{dn}{dT}\Delta T. \tag{4.1}$$

A second laser beam of much larger waist ($\approx 100\,\mu\text{m}$) and lower power, referred to

as the probe beam, crosses the waist of the pump beam at a small angle and is measured with a photodiode sensor. A small portion of the probe beam passes through the thermal lens created by the pump beam, distorting the shape and changing the phase of this part of the probe beam. This portion of the probe beam interferes with the part of the probe which does not pass through the thermal lens modulating the signal, introducing a phase distortion $\Delta\varphi$ [170] to the beam profile. This interference pattern is imaged onto the photodiode detector.

A lock-in amplifier separates the detected signal into a DC signal and a small AC modulation, using the chopping frequency as a reference signal. Assuming that the thermal lens created in the sample is spherical, the portion of the probe beam which passes through this distortion will experience an axial Gouy phase shift of $\arctan(Z/Z_{\mathrm{R}})$, where Z_{R} is the Rayleigh range which is a function of the focus and propagation of the probe beam.

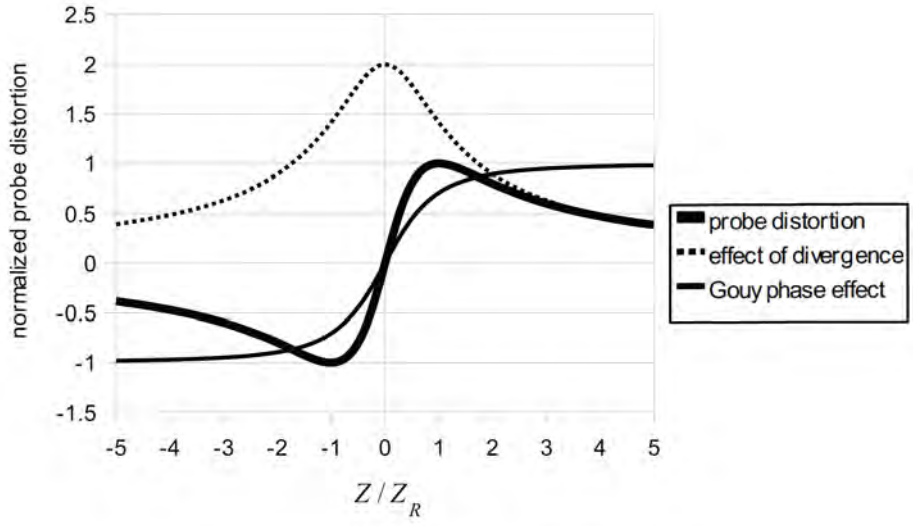

Figure 4.2: Calculated phase distortion of the probe beam propagating in direction Z, normalised to the Rayleigh range Z_{R} for the propagating beam. Figure reproduced from [170].

Fig 4.2 shows the normalised distortion of the probe beam as a function of Z, which is the position along the pump beam, position and Rayleigh range calculated by A. Alexanderovski *et al* [170]. The effects of divergence and the increase in phase velocity from the Gouy phase effect to the probe beam's overall distortion can also be seen. For maximum sensitivity to the weak-wave distortions of the pump beam, these calculations show that the ideal distance for measurements of the probe

distortion is Z_{R}. This can be expressed as a function of the probe waist and its wavelength $Z_{\mathrm{R}} = \frac{\pi w_0^2}{\lambda}$. It is then pertinent for greater sensitivity to $\Delta\varphi$, that the waist of the pump and probe beam cross at an angle θ between 12° and 16° to increase the longitudinal resolution of the measurement [170].

The beams are crossed where $w_{\mathrm{pump}} < w_{\mathrm{probe}}$, where the ratio of $\frac{w_{\mathrm{pump}}}{w_{\mathrm{probe}}}$ defines the longitudinal resolution of the technique. This difference in beam size at the point of measurement is imagined as larger and smaller wavefronts inside each laser beam. As the pump beam is modulated at a low-frequency f, the rate of heating of the sample can be defined in terms of thermal relaxation D of the sample, over a time t

$$D(f, \varepsilon, z) = \int_0^\infty \mathrm{Im}\left(\frac{1}{1 + t - \frac{iz}{1 - i\varepsilon z}}\right) \exp(-ift)\, dt, \tag{4.2}$$

and a geometrical term ε which takes into account the beam radii of the pump and probe

$$\varepsilon = \frac{w_{\mathrm{pump}}^2}{2w_{\mathrm{probe}}^2}. \tag{4.3}$$

In the case where D is defined by f, ε and z, it is pertinent to normalise the modulation frequency f in Eq. 4.2 to the thermal relaxation frequency f_{T} of the material, where the peak absorption α occurs at the pump/probe crossing point. Translating the sample through several positions of z and monitoring the weak wave interactions of the probe beam aligns the surface of the sample to the beam crossing point, where the system is most sensitive. Fig 4.3 shows the layout of the photo-thermal common path interferometer (Gouy Interferometer) setup used for optical absorption measurements. Each pump laser is aligned such that it is centred on the first focusing lens before entering the PCI, ensuring that the pump/probe crossing points for each pump wavelength are consistent.

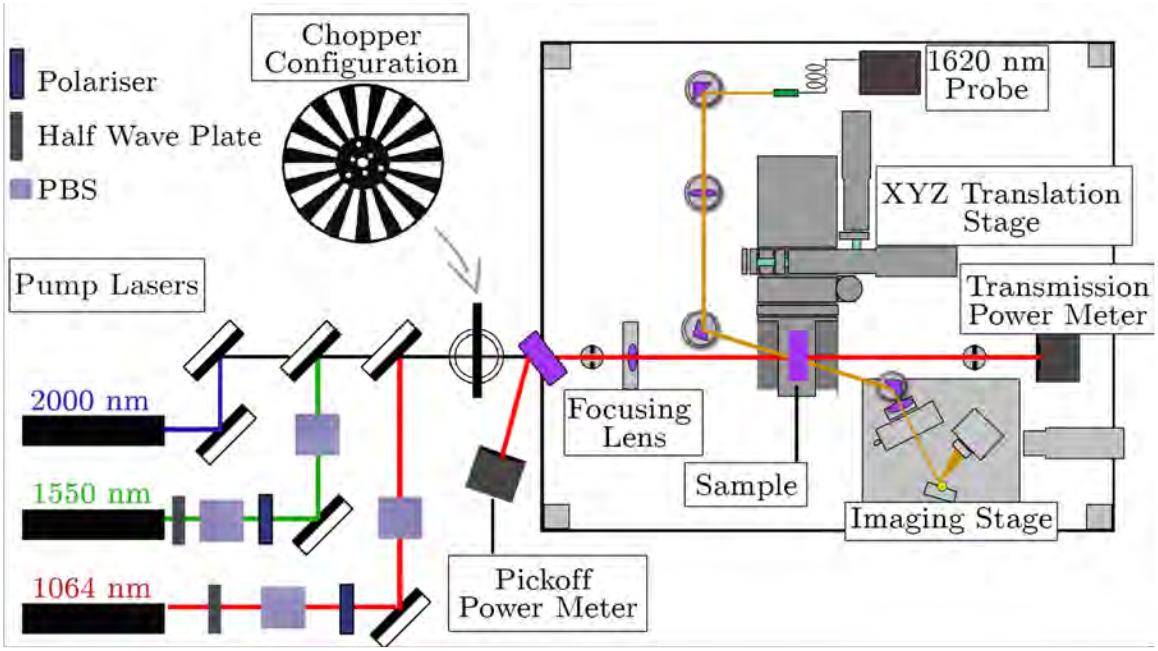

Figure 4.3: A diagram of the PCI setup. Fibre-coupled pump laser wavelengths are denoted by different colours. A solid state probe beam is also shown. Blue squares denote the position of polarising beam splitter cubes (PBS) inside the laser beam path.

To quantify the optical absorption of a coated sample, the system must be calibrated to a reference of known absorption. The calibration reference sample is placed inside the apparatus shown in Fig 4.3 where its surface is held at the crossing point of the pump and probe beams. The calibration reference is a neutral density filter, comprised of fused silica and has an optical absorption of $\alpha = 21.4\%$ at laser wavelengths between 1 µm and 2 µm. The photo-detector used for monitoring the probe beam is placed on a fine control translation stage known as the 'imaging stage', at a position S_{cal}. The sensitivity to the weak wave distortion of the probe is maximised at $\frac{Z}{Z_{\mathrm{R}}} = \pm 1$ from the pump/probe crossing point. In a perfectly aligned system, this allows the weak modulation of the probe (AC signal) and the average probe signal (DC signal) to produce an absorption scaling factor R, if the input pump power P_{I} is known, of

$$R = \frac{AC_{\mathrm{ref}}}{DC_{\mathrm{ref}}} \frac{1}{\alpha_{\mathrm{ref}} P_{\mathrm{ref}}}. \tag{4.4}$$

The magnitude of each signal at a given pump power can then be used to scale the signals observed for samples with different $\frac{dn}{dT}$ responses:

$$\alpha = \frac{AC_{\text{signal}}}{DC_{\text{signal}}} \frac{1}{P_I} \frac{C}{R}. \tag{4.5}$$

Comparison of the measured signal to the expected value from the calibration sample C allows the absorption α to be calculated

$$\alpha = \left(\frac{AC_{\text{signal}}}{P_I DC_{\text{signal}}} / \frac{AC_{\text{ref}}}{P_{\text{ref}} DC_{\text{ref}} \alpha_{\text{ref}}} \right) C. \tag{4.6}$$

By passing the surface of the sample through the beam crossing point the calibrated absorption signal, maximum of 21.4%, shown in Fig 4.4 is produced. The central peak occurs at the maximum distortion of the probe beam recorded during sample translation. On either side of the central peak two artefacts of the interference pattern are visible, colloquially referred to as 'side-peaks'. These can be used as a measure of the alignment quality of the system. The phase shown in Fig 4.4(c) is a consequence of the thermal dissipation of the material with respect to the frequency at which the pump beam is chopped. As the surface of the calibration sample is translated through the beam crossing point a phase of 65° is observed, indicative of a fused silica substrate.

Once the system is calibrated, the calibration sample is replaced with a coated sample of interest. Suppose the thickness or refractive index of the sample differs from that of the calibration sample. In that case, the photo-detector position is changed to account for the optical path difference of the probe laser light inside the sample. Assuming that the original imaging stage position S_{cal} is known the new position of the imaging stage S_{new} on which the photo-detector is mounted is given by:

$$S_{\text{new}} = t_{\text{calib}} + S_{\text{cal}} - \left(\frac{1 - n_{\text{sub}}}{n_{\text{sub}}} \right) \times t_{\text{sub}}. \tag{4.7}$$

As the change in optical path length inside the coating layer can be considered negligible compared to the thicker substrate, the imaging stage's position only depends on the properties of the substrate.

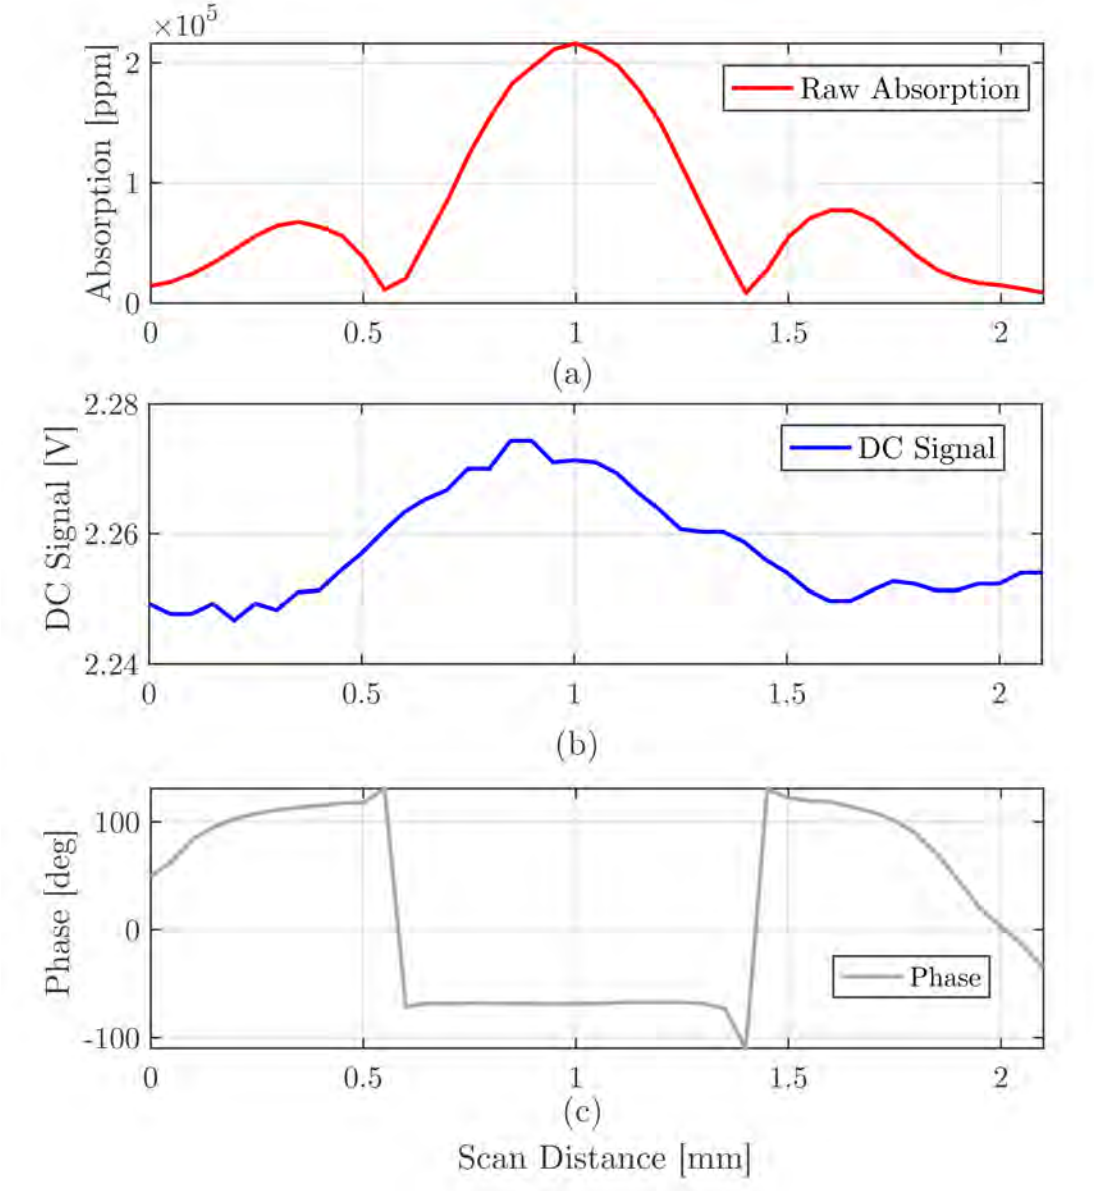

Figure 4.4: (a) Weak wave interaction produced by translating the calibration sample through the beam crossing point. The magnitude of the central peak is the maximum distortion produced by $\frac{dn}{dT}$ and equivalent to the samples absorption α. (b) Average probe signal on photo-detector. (c) Measured phase of probe laser on photo-detector.

The transmitted probe laser is then re-aligned on the photo-detector at its new position, and the new sample is then moved through a range of Z, where the beam crossing point is centred on $Z = 1$. The resulting signals are then scaled using Eq. 4.6.

Usually, multiple points on the surface of the sample are then measured and the average of these measurements produces an average sample absorption $\overline{\alpha}_{\text{sample}}$. After these measurements, the imaging stage is moved back to its original position and the calibration reference re-measured. If the maximum absorption of the signal has changed by $\pm 10\%$, the measurements are repeated[1].

4.0.2 Refractive Index and Attenuation of Light

With proper calibration, the PCI technique provides absorption values as a fraction of the incident pump power which is absorbed by the coating layer, and typically measured in parts per million ($1\times10^{-6} = 1\,\text{ppm}$). This absorption number is directly dependent on the thickness of the coating and the electric field intensity in the coating layer. By taking into account the propagation of the electric field inside the coating layer, a measure of optical absorption which is independent of sample geometry can be produced. In doing so, values can easily be compared to samples of different thickness, or have been deposited on different substrates.

Consider an oscillating electromagnetic field with electric and magnetic components $\overrightarrow{\mathbf{E}}$ and $\overrightarrow{\mathbf{B}}$ respectively, travelling through a medium of permittivity ϵ and permeability μ [171]. The refractive index of the material is then given by

$$n \equiv \sqrt{\frac{\epsilon\mu}{\epsilon_0\mu_0}}. \tag{4.8}$$

Eq. 4.8 states that the velocity of an electromagnetic field passing through a given medium depends on the properties of the material and compares the properties of the speed of light in a vacuum, $\epsilon\mu$ to its speed in free space with properties ϵ_0 and μ_0.

Materials which are optically transparent to a given wavelength of light, interact with the incident light defined by Eq. 4.8, scattering elastically from atoms in the material. Atoms with optical resonances close to the wavelength will no longer scatter elastically and energy is absorbed into the vibrational states of the atom or molecule [171]. Metal oxides and molecules such as H_2O and CO_2 are known to have optical

[1] Re-measuring the absorption of the reference sample allows any changes in alignment of the system to be quantified.

absorption bands in the infra-red spectrum, which can absorb different amounts of light depending on its exact wavelength and the electronic structure of the material. To take into account the attenuation of the light field due to optical absorption For an absorbing material, Eq. 4.8 can then be re-written, taking into account the extinction coefficient k, the imaginary component of n:

$$n = \sqrt{\frac{\epsilon\mu}{\epsilon_0\mu_0}} + ik. \tag{4.9}$$

The extinction coefficient k of each material is wavelength-dependent, providing a measure of the energy lost to the material by optical absorption, which is geometry independent. However, the total absorbed light power in a given layer will depend on both k and the electric field intensity in the layer. The latter is determined by the geometry of the layer

Assuming that the EM wave is propagating in direction z, through a medium of n=1, such as air, the attenuation of the electric field inside the material can be calculated using the work of O. Arnon and P. Baumeister [172]. Eq. 4.10 details the propagation of a light wave through a given material $E(z)$, normalised to the total transverse electric component of the wave. The normalisation condition of both components can then be expressed as:

$$N = \frac{|E(z)|^2}{|E_0^+|^2}. \tag{4.10}$$

The total electric field intensity $\mathbf{E}$ stored in the coating layer can be calculated by integration of N over the coating thickness. Fig 4.5 shows the normalised electric field strength calculated using Eq. 4.10 for three wavelengths of light, 1064 nm, 1550 nm and 2000 nm incident from air (n=1) to a coating layer 500 nm thick(n=1.8).

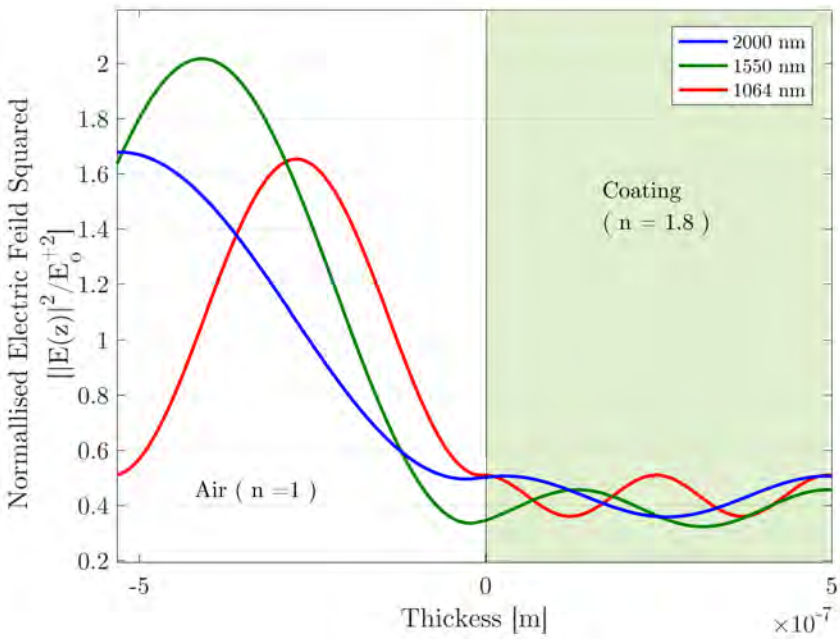

Figure 4.5: Comparison of the normalised electric field strength in a coating layer t=500 nm, n=1.8 at for incident monochromatic light at 1064 nm, 1550 nm and 2000 nm. The coating is deposited on a SiO_2 substrate of $n = 1.44$ (not pictured).

Measured values of optical absorption α, can be converted to k if $\mathbf{E}$ is known. Thus, the total light intensity inside the coating layer can then be used to scale α to produce a unit-less measure of the total light attenuation inside the glass layer

$$k = \lambda \left(\frac{2\pi t^2 \alpha}{\mathbf{E}\lambda} \right) \frac{1}{4\pi}. \tag{4.11}$$

For gravitational wave detector coatings, knowledge of the electric field inside the coating layer is used to dictate the physical thickness of each material layer. The maximum coating reflectivity can be achieved where $t = \lambda/4$. For single-layer coatings, this provides a scaling factor which depends on refractive index and wavelength for measured values of α. This same methodology can also be applied to coatings with multiple layers, again by following the methods described in [172].

In subsequent sections in this chapter, all measured values of α will be quoted in 'ppm' and k where appropriate, for ease of comparison to other measurement systems.

4.1 Ta_2O_5 Single Layers Produced with RLVIP

Reactive Low Voltage Ion Plating (RLVIP) is an energetic coating deposition technique which can produce materials with packing densities comparable to coatings deposited by ion beam sputtering (IBS). Recent studies have shown that RLVIP can produce oxide coatings with significantly lower mechanical loss than the equivalent IBS material [93]. RLVIP deposited coatings are therefore of great interest to the gravitational wave community, both as a potential source of coating thermal noise reduction and the ongoing research into possible links between optical and mechanical properties of coating materials [128].

Coatings made of the same material, produced by different deposition methods can show significant differences in its properties. It has been shown that the deposition rate and energy can be related to the packing density and mechanical losses of the material [143]. The number of free carriers/dangling bonds in the material such as amorphous silicon (aSi) has been shown to decrease with heat treatment, suggesting that the atomic structure could be rearranged into a lower potential state [122]. The optical properties of a coating material can also be changed by heat treating at temperatures higher than the deposition temperature [173].

Therefore, high energy deposition methods which can produce coatings with similar material properties to IBS have been investigated by the author for their optical and mechanical properties (see Chapters 5 and 7), as will be detailed in the following sections.

4.1.0.1 RLVIP Deposition Technique

RLVIP deposition is a novel process that uses a low voltage electric current to vaporise the coating material held in a crucible, producing positively charged ions. An auxiliary high electron density plasma, formed by a hot filament discharge [174], is then injected into the coating material, producing a positively ionised film vapour, resulting in a 'self-bias' potential onto the insulated substrate holder (see Fig 4.6). As the electrons produced by the high current arc have much higher mobility compared to the heavier coating material, they are accelerated away from the emission source towards the top of the coating chamber with

$E_K = (30\,\text{eV} - 50\,\text{eV})$ [6]. The highly energetic bombardment of the substrate with the coating material allows for nucleation and re-deposition in areas that may not be reached otherwise [174] which in theory should produce high packing density coatings without the need for heated substrate deposition. The temperature inside the coating chamber can reach somewhere between 200 °C and 300 °C during deposition [175]. By raising the temperature of the substrate during deposition, this can alter the structural properties of thin-film materials [176, 177].

Due to the high refractive index of Ta_2O_5 [178], and low mechanical loss properties [178] (see Chapter 5) Ta_2O_5 is still of great interest for use in a gravitational wave detector HR coating. In an aLIGO coating, the layers of Ta_2O_5 dominate the mechanical loss and contribute significantly to the coating stack's optical absorption. Therefore any reductions in its mechanical loss or optical absorption would be beneficial to the detector's sensitivity and stability.

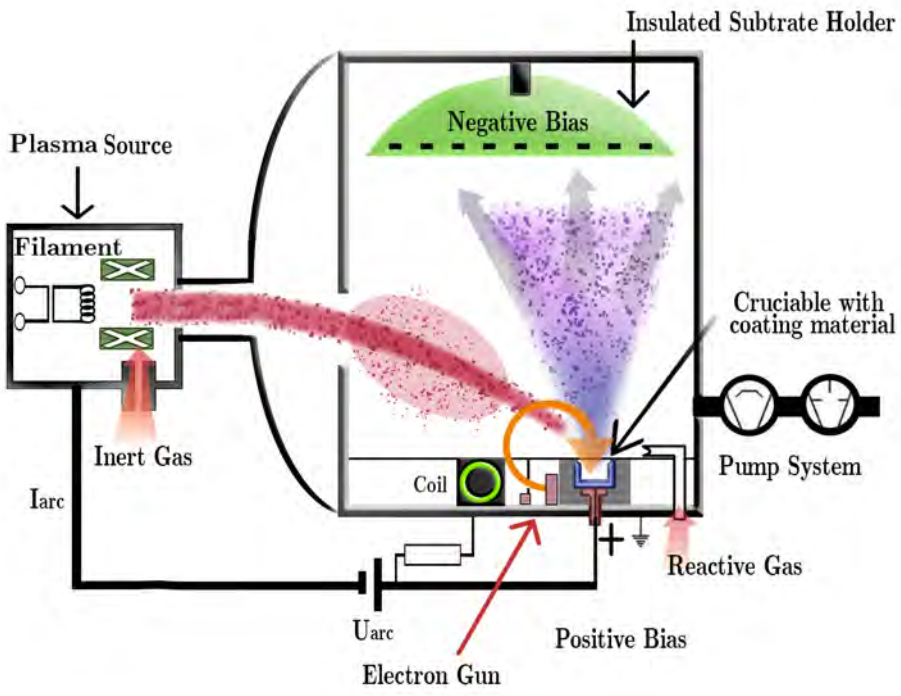

Figure 4.6: Principle of material deposition via ion plating at Tafelmaier [6].

Single layers of Ta_2O_5, 500 nm thick, were deposited onto Corning 7980 (C7980) by Tafelmaier using RLVIP. The absorption of these films was measured at two

wavelengths, 1064 nm and 1550 nm using PCI (see section 4.0.1). At 1064 nm, the absorption of the C7980 substrate to the surface absorption is <1 ppm and therefore it will not dominate over the absorption of the deposited layers.

4.1.0.2 Sample Preparation

The samples' absorption was first measured in the as-deposited state. An average $\alpha = 76\,\text{ppm} \pm 26\,\text{ppm}$ was found for a total of 16 points across the surface of the $\varnothing = 1"$ (25.4 mm) optic. To allow for the multiple studies discussed later in this section to be carried out the sample was then cut into quarters using a diamond blade saw.

During the cutting process, each sample is held in place with bonding wax on the non coated face. The absorption of four points on each new quarter sample were re-measured after being cleaned using acetone and isopropanol to confirm that this process did not alter their absorption. At λ_{pump}=1064 nm, the average absorption after cutting $80\,\text{ppm} \pm 27\,\text{ppm}$, showed no significant difference to the absorption of the original sample.

Measurements at λ_{pump}=1550 nm show an initial absorption of $\alpha = 28\,\text{ppm} \pm 6\,\text{ppm}$, with re-measurement after cutting showing $\alpha = 27\,\text{ppm} \pm 4\,\text{ppm}$. By measuring the sample using ellipsometry, which measures the ratio of p to s polarised light reflected from the sample surface to calculate psi and Delta as a function of wavelength. Fitting to this data allowed the film thickness and refractive index to be inferred.

Ellipsometry measurements of a nominally identical sample (shown in Fig 4.7) revealed that the coating thickness does not deviate across the surface of one sample by more than ±5 nm[2].

[2]Ellipsometry data was assumed using a Sellmeier model of Ta_2O_5

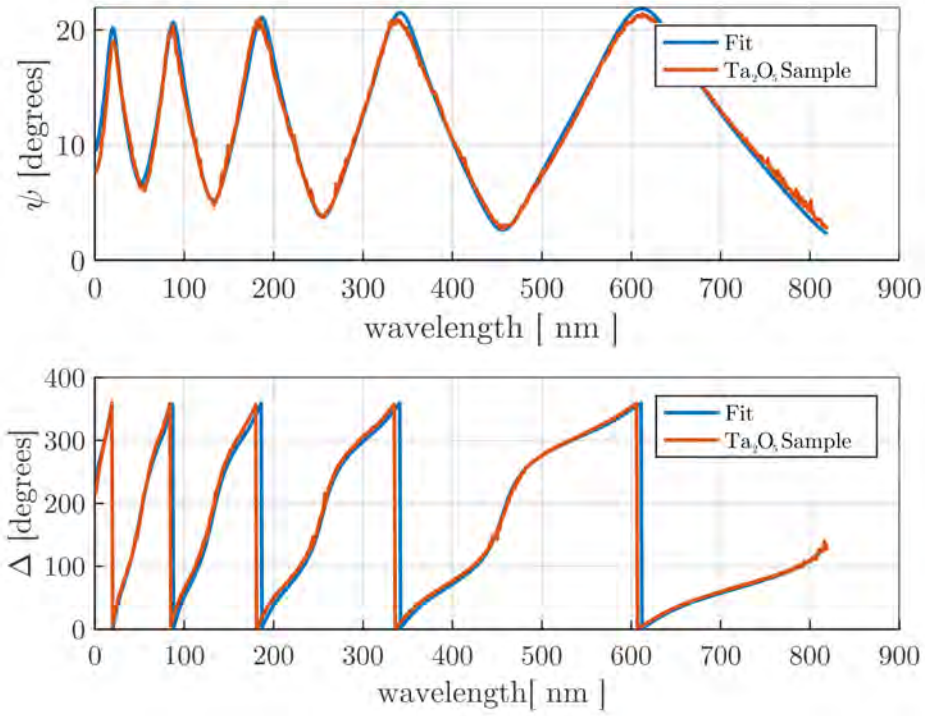

Figure 4.7: Example ellipsometry measurement of RLVIP Ta_2O_5 giving a thickness of 498 nm ± 2 nm.

4.1.1 Effect of the Annealing Temperature on Absorption

To investigate the effect of heat treatment on absorption, a sample of RLVIP Ta_2O_5 was incrementally heat-treated for 4 hours at increasing temperatures. After each heat-treatment step the absorption was measured at 1064 nm and 1550 nm. The results are shown in Fig 4.8. where every point is the average of four absorption measurements at different locations on the sample, with these error bar denoting the standard deviation of the results. At 1064 nm, the absorption decreases consistently with increasing temperature for heat treatments below the deposition temperature (between 200 °C and 300 °C) - by 10% at 100 °C and by 21% at 200 °C with respect to the as-deposited state. For heat treatment at 300 °C and 400 °C, the absorption reduces at a larger rate, with a total reduction of 88% at 400 °C.

At 1550 nm, an absorption decrease similar to that at 1064 nm was observed at 100 °C, but at 200 °C the absorption reverted a value similar to its as-deposited

value. The differing trends at the two wavelengths may indicate two different absorption mechanisms being present in the coating. At 300 °C and 400 °C, the absorption at 1550 nm also reduces significantly, with total reduction 58%, which is a smaller reduction than that at 1064 nm. It is noted that a larger variation in optical absorption is measured from point to point at 1550 nm which could indicate that there are regions of the coating which are more susceptible to absorption of photons at 1550 nm which are not present at 1064 nm - possibly indicating wavelength dependent absorption mechanisms.

The absorption at both wavelengths remains relatively constant for heat treatment at 500 °C and 600 °C, although there is evidence of a small increase at 600 °C, within the bounds of the measurement error.

It is interesting to note that larger absorption changes occur when the heat-treatment temperature exceeds the deposition temperature of around 200 °C. The reason may be that any structural rearrangements of the material,which are possible at this temperature, may be expected to have already occurred during deposition. In contrast, the higher thermal energy available at higher temperatures may be enough to allow the material's structure to rearrange [122], changing its absorption.

As the crystallisation temperature of Ta_2O_5 deposited by IBS is known to be between 500 °C and 700 °C [179], at these temperatures, this may also occur for the RLVIP material. Optical absorption measurements after 700 °C heat treatment were only possible at 1064 nm. This was due to scattering of the pump light by the coating onto the photodiode sensor, which is more sensitive at 1550 nm than at 1064 nm. The high amount of scattering is an indication that the coating had crystallised.

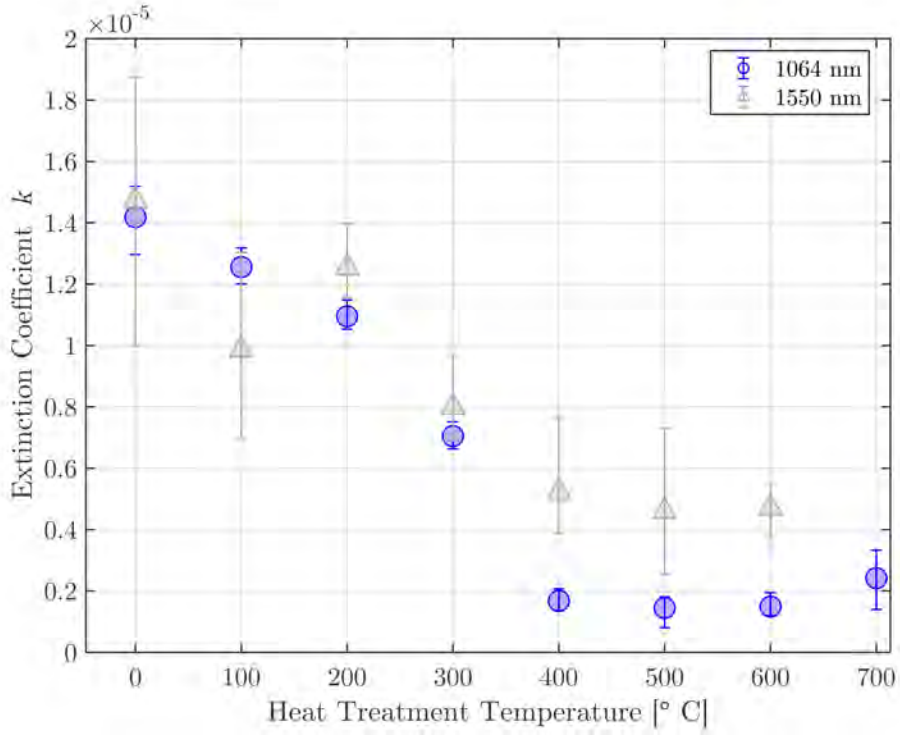

Figure 4.8: Optical absorption of RLVIP Ta_2O_5 at 1064 nm and 1550 nm as a function of heat treatment temperature. Optical absorption is quoted in terms of the materials extinction coefficient k. Error bars denote the standard deviation in measured absorption values over 4 or more points on the sample surface and take into account possible variations in the optical properties of the film (see section 4.1.4.3).

A comparison of the absorption scans measured at 1550 nm for heat treatments at 500 °C and 700 °C, is shown in Fig 4.9. An optical band-pass filter, which attenuates wavelengths far from the probe laser's wavelength (1610 nm), was used in front of the photo-detector of the PCI. As 1550 nm is much closer in wavelength than 1064 nm, more of this light is transmitted by the filter and measured by the photo-detector creating the larger signal shown in Fig 4.9. This makes measurements at 1550 nm much more susceptible to scattered light than measurements at 1064 nm.

For the sample heat treated at 700 °C, closer inspection of the phase shows a value oscillating around zero, indicating that mainly noise from scattering was measured, with little signal from the absorption. A corresponding dip can be seen at ∼8 mm can also be seen, produced from the rear surface of the sample.

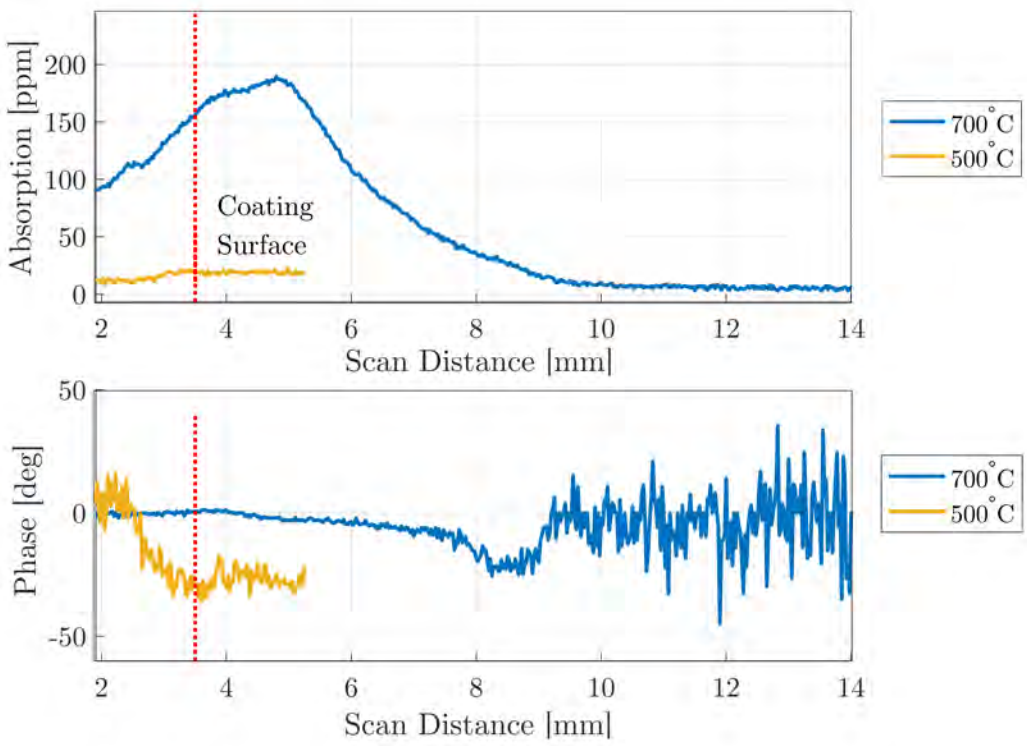

Figure 4.9: Optical absorption of 500 nm thick Ta_2O_5 at 1550 nm heat treated at 500 °C and 700 °C. The high absorption signal of the 700 °C sample compared to the 500 °C sample likely originates most from scattering due to crystallisation.

This study has shown that the optical absorption of RLVIP Ta_2O_5 can be reduced at 1064 nm and 1550 nm by post deposition heat treatment. Minimum absorption for both wavelengths was found at 500 °C, reduced by ∼90% at 1064 nm and 63.5% at 1550 nm compared to the as-deposited state (see Fig 4.8). As the absorption trend changes above this temperature, it is assumed that the crystallisation temperature of this material occurs between 500 °C and 600 °C, is slightly lower than that of Ta_2O_5 deposited by IBS (600 °C - 700 °C), [180]. By sampling the absorption at more than one wavelength, a more detailed picture of the absorption mechanisms inside the material can be obtained. Below 400 °C, the dominant mechanism responsible for absorption is much more sensitive to photons at 1064 nm than 1550 nm. At higher heat treatment temperatures, the sensitivity flips so that the 1550 nm absorption is higher than the 1064 nm absorption.

4.1.2 Effect of Heat Treatment Conditions on Optical Absorption

4.1.2.1 Introduction

Ta_2O_5 coatings can be deposited using Ta_2O_5 target or a combination of metallic Ta target and O_2 gas inside the deposition chamber. In each instance, it has been shown that the final material can be deposited in an 'oxygen-poor' state, where the ratio of oxygen to metal atoms is less than 5/2 [181]. This leads to some tantalum atoms having unpaired electrons, creating defects known as colour-centres [182], which can absorb incident laser radiation.

To investigate the effects of oxygen content on the absorption a comparison of samples of the RLVIP Ta_2O_5 coatings were annealed in vacuum and in air. The absorption results were compared to nominally identical samples annealed in air.

4.1.2.2 Sample Preparation

As in section 4.1, four samples were produced by quartering a tantala coated 1" diameter C7980 disk 6.1 mm in thickness. Two samples labelled *A* and *C* were heat-treated in atmosphere at 400 °C in several steps, for varying lengths of time. A second set of samples, *B* and *D*, were put through identical heat-treatment conditions, but under high vacuum ($\leq 4\times10^{-4}$ mbar). These samples were placed inside the oven at room temperature, and the pressure was reduced, before heating at 400 °C for a set time. After heat treatment, the oven was left to cool to room temperature while under vacuum before the samples were removed. Before heat treatment, all samples and the oven were cleaned with acetone and isopropanol.

4.1.2.3 Absorption at 1064 nm

Samples C and D were subjected to a series of short timescale, identical heat treatments in air and vacuum respectively. Similarly, Sample A (air) and Sample B (vaccum), were put through a series of longer identical heat treatments. The absorption of these samples at 1064 nm as a function of total heat treatment duration is shown in Fig 4.10 and summarised in Tab. 4.1. For heat treatments in air, both Samples C and A showed a significant reduction of 92% and 73%,

respectively, after the first heat treatment step. Following this, the next two heat treatment steps produced an increase in absorption for both samples but the absorption was 84% (Sample A) and 56% (Sample C) lower than in their as-deposited states. Sample C was subsequently heat-treated for a much longer duration (an additional 190 mins) . The absorption dropped to <10% of its initial value[3].

While this behaviour is somewhat complex and the exact mechanisms behind these changes are as of yet unknown, it is clear that heat treatment in air can significantly reduce the optical absorption. It is also clear from this data that subsequent heat treatments at the same temperature may tend to increase optical absorption.

The absorption at 1064 nm of Samples B and D heat treated in vacuum does not follow the trends observed for in-air heat treatment. For Samples B and D, the absorption was found to increase with every heat treatment step. To further investigate, the absorption of samples A-D at 1064 nm as a function of the number of heat treatment steps are shown in Fig 4.11.

When analysed in this reference frame, the similarities between the samples become noticeable. Both pairs of samples, heat-treated in air and vacuum, show an apparent trend in absorption with the number of heat treatment cycles, with vacuum samples (Sample B and Sample D) increasing in optical absorption with little dependence on heat treatment duration.

[3]It should be noted that level of background noise in the measurement system at this time was comparable to the size of the absorption signal for this measurement. As such, the absorption for Sample C after 580 minutes is an upper limit of optical absorption.

Sample	Total Heat Treatment Duration [mins]	k	σ	Conditions
Sample A	As Deposited	1.3×10^{-5}	1.4×10^{-6}	Air
	235	1.1×10^{-6}	1.2×10^{-7}	Air
	550	1.3×10^{-6}	1.2×10^{-7}	Air
	955	2.1×10^{-6}	2.9×10^{-7}	Air
Sample B	As Deposited	1.5×10^{-5}	1.6×10^{-6}	Vacuum
	227	2.5×10^{-5}	2.6×10^{-6}	Vacuum
	534	9.1×10^{-5}	9.7×10^{-6}	Vacuum
	991	1.2×10^{-4}	1.2×10^{-5}	Vacuum
Sample C	As Deposited	1.0×10^{-5}	9.0×10^{-7}	Air
	40	2.8×10^{-6}	2.8×10^{-7}	Air
	100	2.9×10^{-6}	2.7×10^{-7}	Air
	190	4.5×10^{-6}	4.7×10^{-7}	Air
	580	9.4×10^{-7}	1.0×10^{-7}	Air
Sample D	As Deposited	1.2×10^{-5}	1.1×10^{-6}	Vacuum
	40	2.2×10^{-5}	2.0×10^{-6}	Vacuum
	100	8.3×10^{-5}	3.2×10^{-6}	Vacuum

Table 4.1: Summary of measured k values for each heat treatment on Samples A - D at 1064 nm.

Following the first and second heating cycles the absorption of both samples at 1064 nm increased by $40\% \pm 5\%$ and $85\% \pm 1.5\%$ with respect to their as-deposited states. Only Sample D was heat-treated for a third time where its absorption increased by a further 5%. Samples A and C, which were heat-treated in air for different duration's, both show a reduction in absorption after their first heat treatment. The first heat treatment of Sample A undergoing a 7.5 times longer duration heat treatment and producing a ∼60% lower absorption than for Sample C. Between the second and third cycles the absorption of air heat-treated samples increases with respect to each other, with only ∼10% difference in absorption between both samples. As the oven and each sample are cleaned before every heat treatment, a repeatable level of sample contamination seems an unlikely explanation

for the step-wise absorption changes.

4.1.2.4 Absorption at 1550 nm

Fig 4.10 and Tab. 4.2 detail the absorption of samples A, C and D measured at 1550 nm. Due to a sporadic fault with the 1550 nm laser, not all heat treatments measured at 1064 nm have a corresponding measurement at 1550 nm[4]. For the air-annealed samples, higher absorption was measured at 1550 nm for Sample A than Sample C, mirroring the trend at 1064 nm. It is interesting to note that a lower absorption can be produced by heat-treating the sample for over 4 hrs, with a total heat treatment time of 9 hrs resulting in an absorption 40% lower than the absorption measured when heat treating a nominally identical sample in air at 400 °C for a shorter duration (shown in Fig 4.8).

Sample	Total Heat Treatment Duration [mins]	k	σ	Conditions
Sample A	955	1.0×10^{-5}	5.9×10^{-7}	Air
Sample C	40	2.0×10^{-5}	1.2×10^{-6}	Air
	100	5.8×10^{-6}	3.5×10^{-7}	Air
	190	5.3×10^{-6}	3.1×10^{-7}	Air
	580	7.8×10^{-6}	5.4×10^{-7}	Air
Sample D	As Deposited	1.4×10^{-5}	8.4×10^{-7}	Vacuum
	40	2.1×10^{-5}	1.4×10^{-6}	Vacuum
	100	3.0×10^{-5}	1.8×10^{-6}	Vacuum

Table 4.2: Summary of measured k values for each heat treatment on Samples A - D at 1550 nm.

Further measurements were made on Sample D, to see how the absorption is affected after being left in a laboratory atmosphere for a considerable time. The absorption of the samples again increased, by more than a factor two. Heat treatment of Sample D increased its absorption to $k = 7\times10^{-5}$ after 3 hrs (see

[4] At 1064 nm all samples were measured in their as-deposited state. Only Sample D was measured at 1550 nm in this state due to a laser fault over this time.

Fig 4.10). This supports the case that a chemical change was occurring in the samples when heat treated under vacuum, while it is not present in the comparison study using heat treatment in air.

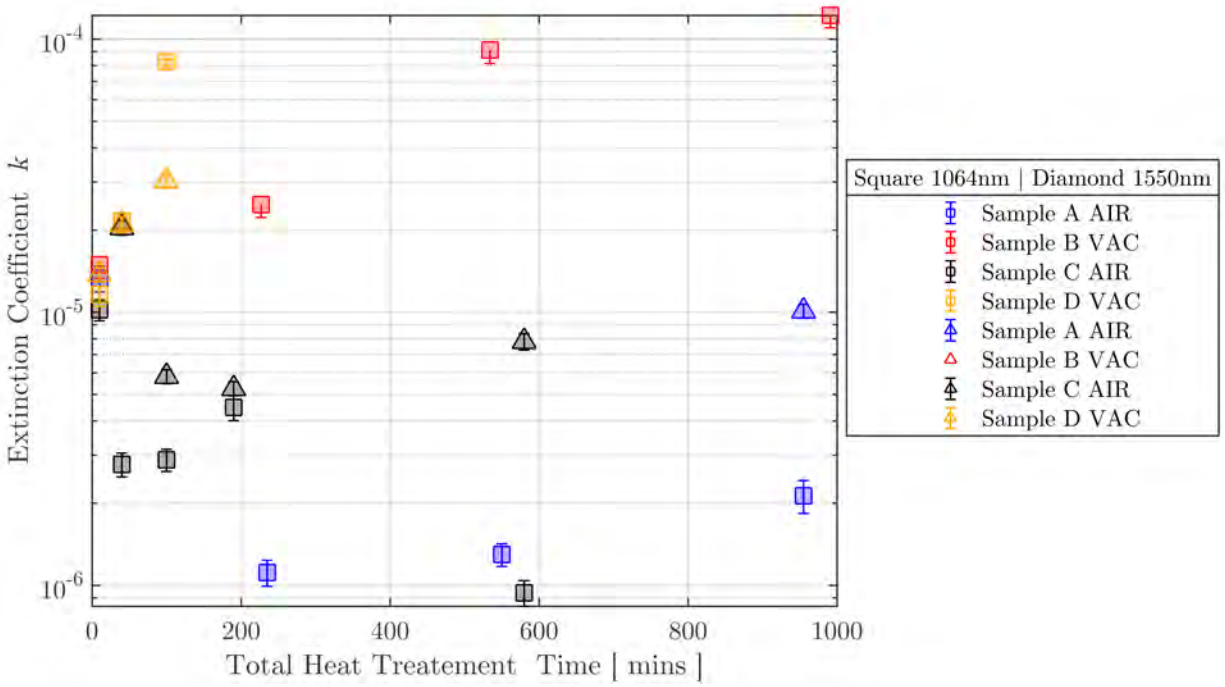

Figure 4.10: Comparison of k at 1064 nm and 1550 nm of RLVIP Ta_2O_5 when annealed at 400 °C in air and in vacuum as a function of total heat treatment duration.

Comparing measurements of Sample D indeed shows that the first heat treatment in vacuum has broadly similar effects on the absorption at 1064 nm and 1550 nm. The second heat treatment cycle again increases the absorption of Sample D by 34% compared to the 46% increase at 1064 nm. After a second, short heat treatment of Sample D, a more considerable difference between the two wavelengths was measured, but the absorption continues to increase at 1550 nm.

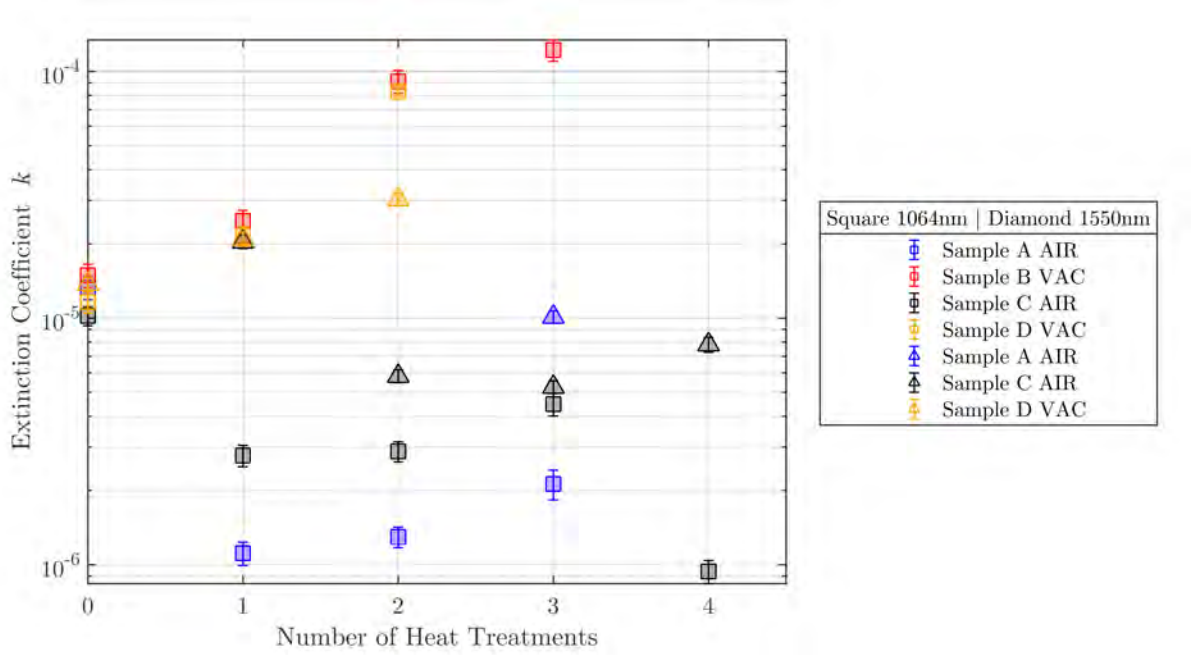

Figure 4.11: Comparison of k at 1064 nm and 1550 nm of RLVIP Ta_2O_5 when annealed at 400 °C in air and in vacuum as a function of number of heat treatment cycles.

This trend becomes evident when looking at the number of heat treatment cycles shown for both wavelengths in Fig 4.11. Sample C, which was heat-treated in air, was measured at 1550 nm after the first heat treatment (total duration 1 hr). This will be used as a baseline for comparison. The second heat treatment of Sample C saw two different trends emerging, producing an increase at 1064 nm, but a decrease at 1550 nm. After the third heat treatment, Sample C's absorption continues to decrease to levels comparable to its 1064 nm absorption. However, its final longer duration heat treatment shows an increase in absorption at 1550 nm, while at 1064 nm the lowest absorption was found. This difference in optical absorption responses to the heat treatment environment and duration at different wavelengths suggests different sources of absorption that are more susceptible to different energies of light.

It is known that heat-treating a material below its fictive temperature will reduce the surface water content [183], decreasing infrared absorption, with the maximum reduction in absorption produced with minimal water content in the coating layer. It would be expected that this effect would be more visible at 1550 nm, but still detectable at 1064 nm. This could explain the initial reduction observed in both air

heat-treated samples, with water being 'driven' out of the coating surface, but fails to explain the sudden increase for further heat treatments. It would be expected that any samples left in a laboratory atmosphere for an extended time would see an increase in their surface water content [183]. Therefore this also cannot fully explain these observations.

4.1.2.5 Effect of Laboratory Atmosphere on the Optical Absorption

To better understand how the measured absorption reflected the changes in the coated samples, each sample was left unaltered in a laboratory atmosphere for more than 300 days, after which the absorption of all samples was re-measured. The results are shown in Figures 4.12 and 4.13.

At 1064 nm samples which were previously heat treated in vacuum (Samples B and D) show a reduction in absorption. Over this time the absorption of Sample B reduced by $\approx$ 50% while Sample D's absorption also reduced by $\approx$ 10%. Further measurements more than 830 days after the samples' last heat treatment showed that the absorption of both B and D continued to decrease compared to their last measured states by 37% and 30% respectively. Re-measurements after 830 days showed a noticeable increase of 26.6%, negating any changes in the previous measurement. As the absorption change is comparable to the spread in measured k, changes on this level could be due to measurement error.

Fig 4.13 shows heat treatments of the same samples measured at 1550 nm. Samples B and D also show decreasing trends in absorption at this wavelength. After 830 days, the absorption of both samples reduced by 57% and 59% respectively. Unfortunately, Sample B was not measured at 1550 nm previous to the long duration due to a sporadic fault with the 1550 nm laser. Sample A (previously heat-treated in air) shows a much larger decrease in absorption ($\sim$60%) over 375 days.

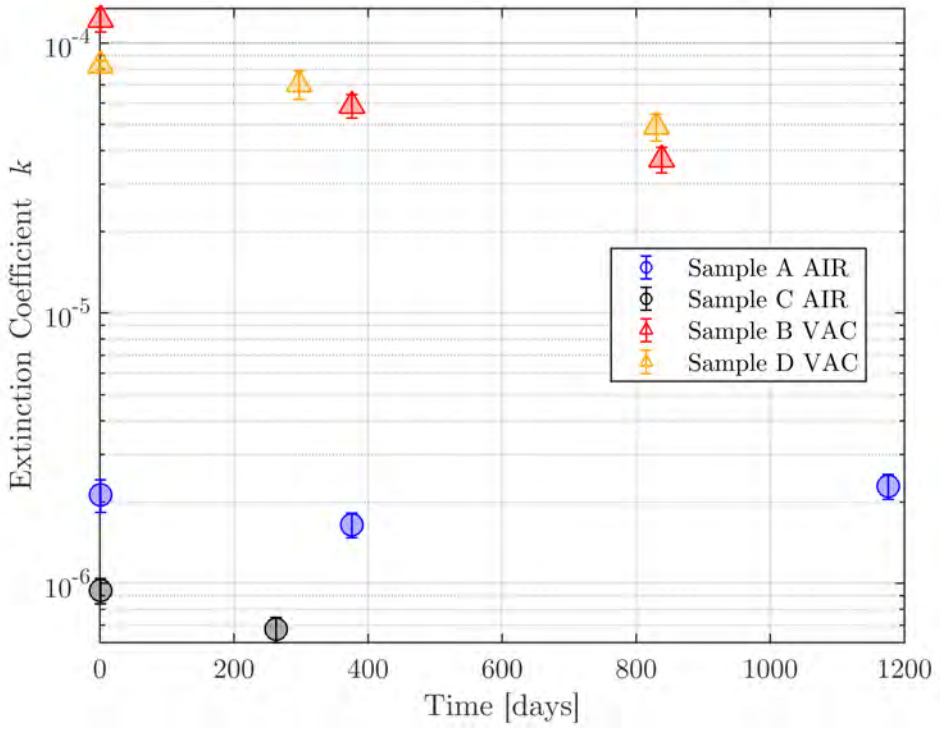

Figure 4.12: Optical absorption of RLVIP Ta_2O_5 samples A-D at 1064 nm with no active alteration, as a function of time. Points at zero mark the last heat treatments for each sample.

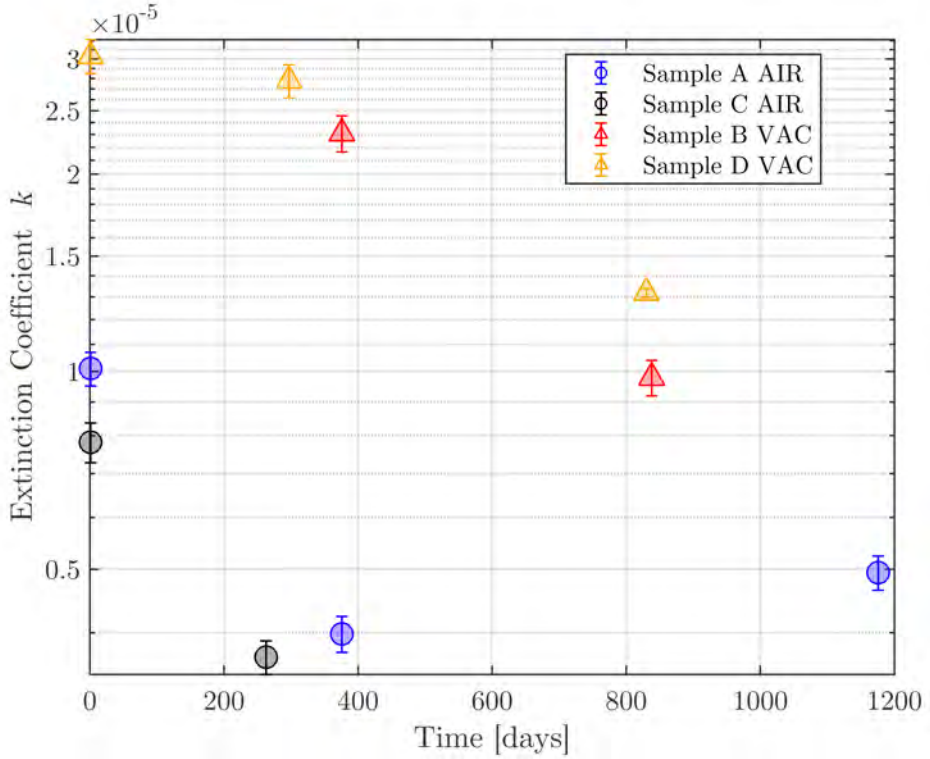

Figure 4.13: Optical absorption of RLVIP Ta_2O_5 samples A-D at 1550 nm with no active alteration, as a function of time. Points at zero mark the last heat treatments for each sample.

However, after being left for an additional 800 days, a ∼20% higher absorption was measured. Interestingly Sample C, which was previously heat-treated for the most prolonged period in the previous section, shows the most considerable reduction in absorption at 1550 nm, decreasing by 52% over 263 days.

The trends observed in these measurements suggest that something is changing in the Ta_2O_5 after subsequent heat treatments. The significant absorption reduction of the samples annealed in vacuum suggests that the mechanisms responsible for a decrease in absorption require reclamation of material from the atmosphere.

In summary, these measurements have revealed several points of interest.

- Samples which are heat-treated in air show a decrease in absorption with longer heat treatment duration.
- In contrast, the same samples which were heat-treated under vacuum show a large increase in their absorption after being heat-treated for the same duration.
- Leaving each sample in a laboratory environment for more than one year results in a significant change in absorption, suggesting that something is changing inside the coating layers, possibly related to atmospheric exposure.
- The optical absorption at 1550 nm changes considerably less than at 1064 nm.

If the water content in the sample was thought to contribute to the optical absorption of these samples, it should be visible as an increase in absorption at 1550 nm compared to measurements at 1064 nm which does not seem to be the case. As the water reclamation and sample contamination are both unlikely culprits for the measured absorption changes in the samples, this suggests a chemical or structural change in the coatings is responsible for such a change in absorption.

It has been previously shown that the dangling bond density of the material can be correlated to the absorption in the infrared spectrum [7]. For Samples B and D, which were heat-treated in vacuum, it is thought the density of dangling bonds in the material can be altered, increasing its optical absorption. As it is currently unknown if the heat treatment temperature of 400 °C is prevalent to the changes in optical absorption, it can be speculated that the mobility of electrons inside the material is

maximised at an intermediate temperature. As the samples were heated and cooled, this would give two points in the cycle at which electron mobility was maximised and could partly explain the correlation between optical absorption and the number of heat treatment cycles.

4.1.2.6 Correlation between O_2 and Absorption

Changing the amount of oxygen bonded in each tantalum molecule can increase free carriers inside the material. To investigate the correlation between oxygen content and optical absorption, these samples require measurements which can characterise the density of elements in the coating layer. As Ta_2O_5 coatings are thought to have been deposited in an 'oxygen-poor' state, it would not seem unreasonable that the optical absorption can be decreased by allowing the material to bond with more atmospheric oxygen.

Rutherford back-scattering (RBS) measurements can be used to identify the density of particular elements in a material. This process uses a stream of alpha particles and a linear accelerator to collide each particle at a known angle with the surface of the coating. The measured change in energy of the alpha particle after the collision can then be used to produce a spectrum containing information about the elements present in the material [184].

Four samples were sent to Université de Montréal for RBS measurements to investigate the correlation between oxygen density and absorption: Samples A and B were unaltered from the previous study (see section 4.1.2); Sample D, which had been re-annealed in vacuum for 1 hour immediatly prior to being sent, and an as-deposited control sample were sent.

Measurements were carried out by A. Lussier and F. Schiettekatte (Université de Montréal) on two separate dates five weeks apart to track the possible changes in the samples over time. Over this timescale, the predicted change in absorption of each sample is minimal; however, if there are changes in oxygen content between each sample in both sets of measurements, this would conclude that the changes in their measured absorption could be correlated to oxygen content. Fig 4.14 shows an example (RBS) spectrum of Sample A showing the number of counts as a function of

energy (keV) where the RBS signal, measured on both dates is compared. The blue dots denote the initial measurement, and the orange points denote the second, later set of data.

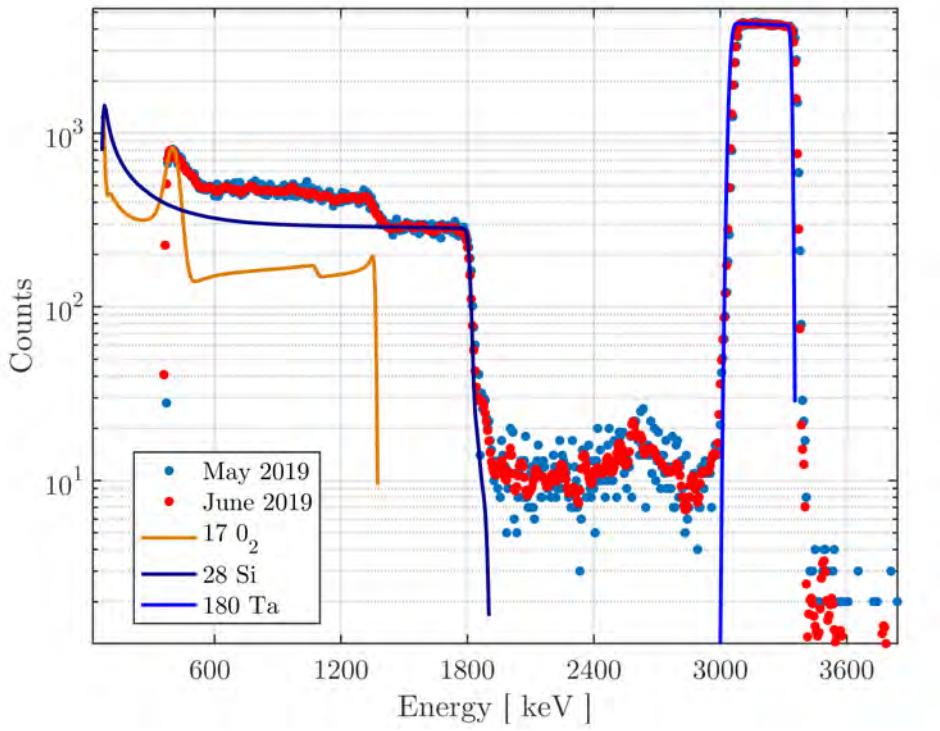

Figure 4.14: Example spectrum of a control sample. Blue points denote the initial measurements, and the orange points show the measurement several weeks after. SIMNRA simulations were used to process channel data, producing expected fits correlated with energies from a given element.

Channels between 300 and 500 (energies between 1800keV and 3150 keV) are indicative of heavy elements such as Fe, Co or Ni being present. However, the resolution of this measurement does not allow for separation of the signals. The heavy element concentration fluctuates between samples and between measurements but does not change outwith the overall error on the measurement. An example of this is shown in Fig 4.15.

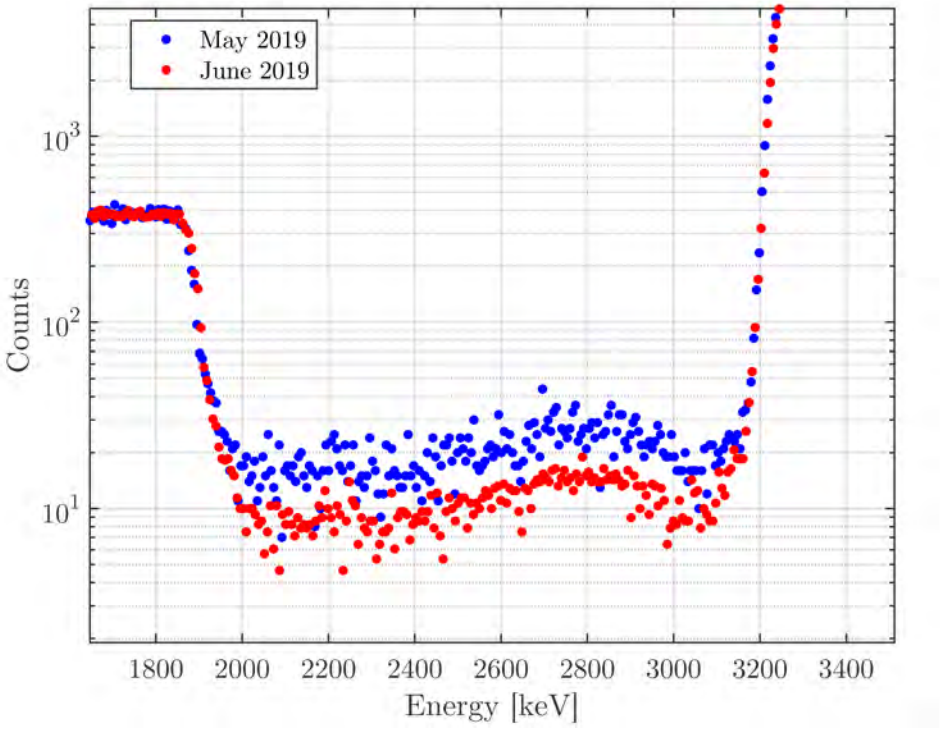

Figure 4.15: Energy range related to heavy elements measured on Sample B, blue dots denote the initial measurement with the orange denoting the later measurement.

In each case, all samples are recorded with initially 0.2% concentration of heavy elements with subsequent measurement of 0.1%. The low concentration of heavy elements recorded is indicative of a signal from the surrounding metallic holders used to contain the samples during the measurement.

Isolating the energy range between 500 and 1500 keV, which corresponds to oxygen, allows any changes in the samples' oxygen content to be analysed. An example of this range for Sample A is shown in Fig 4.16

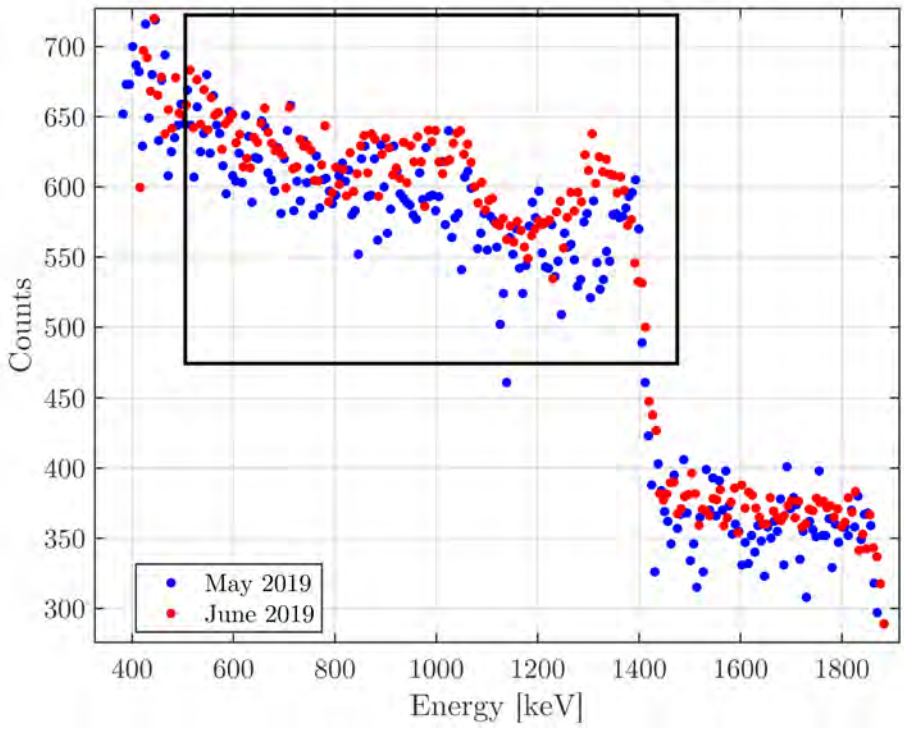

Figure 4.16: Comparison of energy spectra of Sample A for initial and secondary measurements. Marked region of interest indicates energies which can be produced by O_2.

It was noted that over 35 days, a slight increase in oxygen content for Samples A and D and the un-annealed control sample were observed. However, the changes recorded between these sets of measurements are well within experimental error. A full breakdown of all the identified elements inside each coating is shown in table 4.3. Each sample shows a concentration of heavy elements of less than 0.2%, and every sample that was annealed shows a slightly higher concentration of total atoms per centimetre squared than the control. However, the difference is within the error bars of the measurement. A similar observation was made for oxygen and tantala content.

Sample	O_2 (%)	Ta (%)	Co/Fe/Ni (%)	Atoms/cm^2 (1×10^{15})
Control	74.2 ± 1	25.8 ± 0.4	<0.2	4000 ± 30
A	74.3 ± 1	26.1 ± 0.4	<0.2	4020 ± 30
B	73.5 ± 1	26.5 ± 0.4	<0.2	4020 ± 30
D	73.6 ± 1	26.4 ± 0.4	<0.2	4020 ± 30

Table 4.3: Summary of initial element concentrations in RLVIP Ta_2O_5, measured through RBS at 3.9 MeV (May 30th-31st).

Sample	O_2 (%)	Ta (%)	Co/Fe/Ni (%)	Atoms/cm^2 (1×10^{15})
Control	74.3 ± 1	25.7 ± 0.4	<0.1	3960 ± 30
A	73.5 ± 1	26.5 ± 0.4	<0.1	3970 ± 30
B	73.7 ± 1	26.7 ± 0.4	<0.1	3930 ± 30
D	73.3 ± 1	26.3 ± 0.4	<0.1	4200 ± 30

Table 4.4: Summary of repeat measurements of element concentrations in RLVIP Ta_2O_5, measured through RBS sputtering at 3.9 MeV (July 4th).

Interestingly there is a noticeable change in the concentration of atoms for all samples compared to their initial measurements. Sample B, which was last heat treated over two years before this measurement shows a decrease in density, following the same decreasing trend shown via absorption measurements in Figures 4.12 and 4.13. However, this is also shown in Sample A, which was left unaltered for nominally the same length of time. Sample D, which had been re-annealed weeks before all RBS measurements show a noticeable increase in density. Since there is no correlated change in oxygen or tantalum content, it cannot be determined that these are responsible for the changes in absorption measured using PCI. Since there is no noticeable increase in oxygen content in these measurements, no direct conclusions can be drawn vis- *á*vis surface water content in the coatings, as the error associated with these measurements does not allow clear conclusions to be extracted from the data.

Unfortunately, any changes in the oxygen content of the coatings over the 5 week period are below the level which can be detected using this technique. And as such, no reliable conclusions can be drawn as to the origin of the changes responsible for the change in absorption. The PCI technique may be sensitive to changes in atomic structure, or elemental density which cannot be measured via RBS.

Some literature suggests that the observed changes in absorption of tantala coatings could be related to the presence of microvoids/or pores, in the coating layers [128]. There are ongoing plans to test the samples at Stanford University using nano SIMS. This technique uses a focused ion beam to sputter atoms from the coating material, the energy of which can then be used to determine which elements are present in the coating. This technique should have a higher sensitivity than RBS techniques. Plans to investigate the structure of these coatings are also being considered. Methods such as Raman Spectroscopy can be implemented to see how the structure varies depending on the annealing environment. However, due to limitations placed on experimental work globally in March 2020, no further measurements, structural or otherwise, could be carried out. This will be subject to further investigation but goes beyond the scope of the thesis.

4.1.3 Development of 2D Absorption Mapping

PCI measurements, such as those presented in this chapter, typically entail measuring various points on a sample's surface. The resolution of such a scan is defined by the size of the pump waist and the scanning motors' step-size. The sample is mounted on a three-axis stage which allows for micrometre precision movements. At every measurement position, the incident pump power is used to calibrate the absorption signal. Thus any changes in pump power during the measurement will introduce an error. As coatings can vary in optical absorption and thickness across the surface of the sample, information can be lost about spatial variation of the absorption. This section details the steps taken to increase laser stability, and changes made to the PCI set up to decrease the systematic error associated with pump power drift.

4.1.4 Polarisation Stabilisation and Power Correction

During regular operation of the 1064 nm pump laser, the oscillating polarisation translates into a change in output power of ±30% over a period of around 100 seconds as the beam transmits through a polarising beam splitter cube. This fluctuation, in turn, introduces a large source of instability into the PCI measurement where the calibration drifts following this power fluctuation. An example is shown in Fig 4.17. A slow period change in power is observed over the 400-second measurement, with a second, faster oscillating component also apparent. To mitigate this change, the polarised light must first be stabilised to reduce the total error produced by the fluctuating power.

4.1.4.1 Polarisation Stabilisation

The output polarisation of the optical fibre depends strongly on temperature-induced stress in the fibre; changes were made to the system to reduce the effects of stray polarisation that reaches the absorption sample. The temperature of the laboratory was stabilised using an air-conditioning unit.

As the polarisation of the optical fibre output cannot be changed without changing the laser, the polarisation which reaches the PCI must be filtered. In order to decrease the fluctuations in light polarisation which reach the PCI two polarising beam splitters (PBS), with a quarter-wave and half-wave plate were used (see Fig 4.18). The method used proceeds as follows:

- The radial orientation of the fibre coupled laser output was orientated such that ∼90% of laser light is removed by the first PBS.
- PBS1 filters out a large amount of the stray polarised light.
- A second PBS was then placed at a known distance in the beam path. This further filters the remaining ∼10% of stray polarised light.
- PBS2 removes any remaining stray polarised light from the beam path which has not already been removed by PBS1. The use of both beam splitters in transmission produces a laser beam which is predominantly S polarised. The resultant power which is transmitted from the PBS pair then depends on the initial polarisation from the fibre-coupler.

- A half-wave plate in a rotating mount was placed between PBS1 and PBS2.

This allows the user to increase the intensity of the initial laser light (before PBS1) and control the total light power which reaches PBS2, and therefore the absorption sample. Fig 4.17 shows the recorded laser power in the PCI before and after the stabilisation optics were put in place. This shows that the resultant polarisation from the optical fibre has indeed been mitigated As the large slow drift ±30% shown in the previous measurements has been reduced to ±1%.

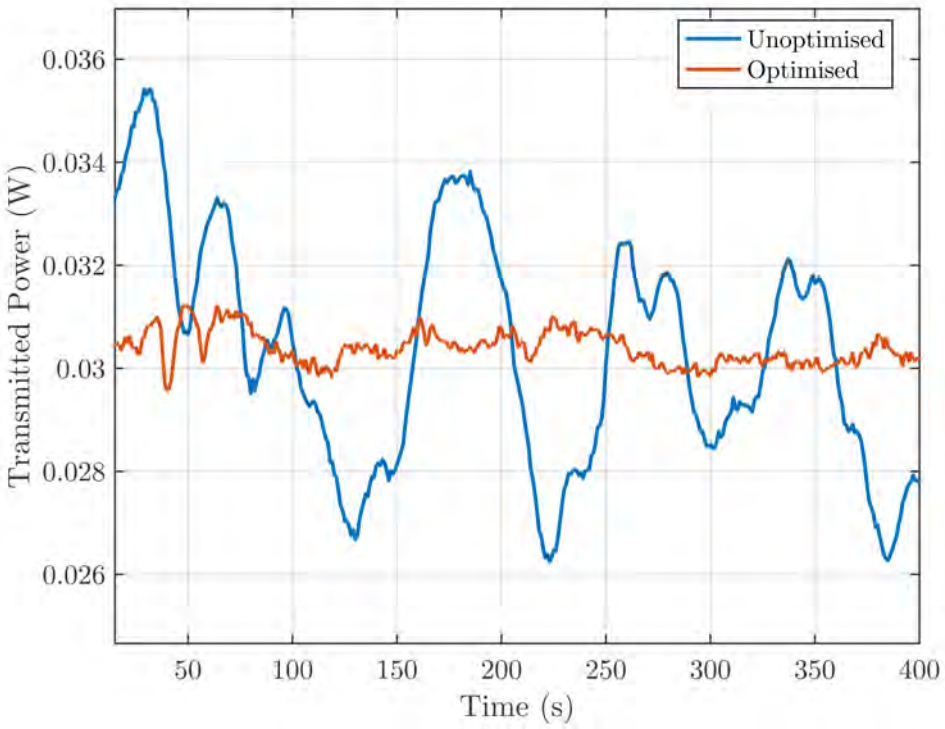

Figure 4.17: Measured output power using the 1064 nm fibre coupled laser, before and after polarisation optimisation.

This power stabilisation scheme cannot be used with the 2000 nm pump laser, due to a lack of suitable optics at this wavelength. Implementation of the beam splitter in the beam path of the high power lasers reduces the statistical fluctuation of their final incident power. Still, as shown in Fig 4.17, the incident power on the sample cannot be considered a constant.

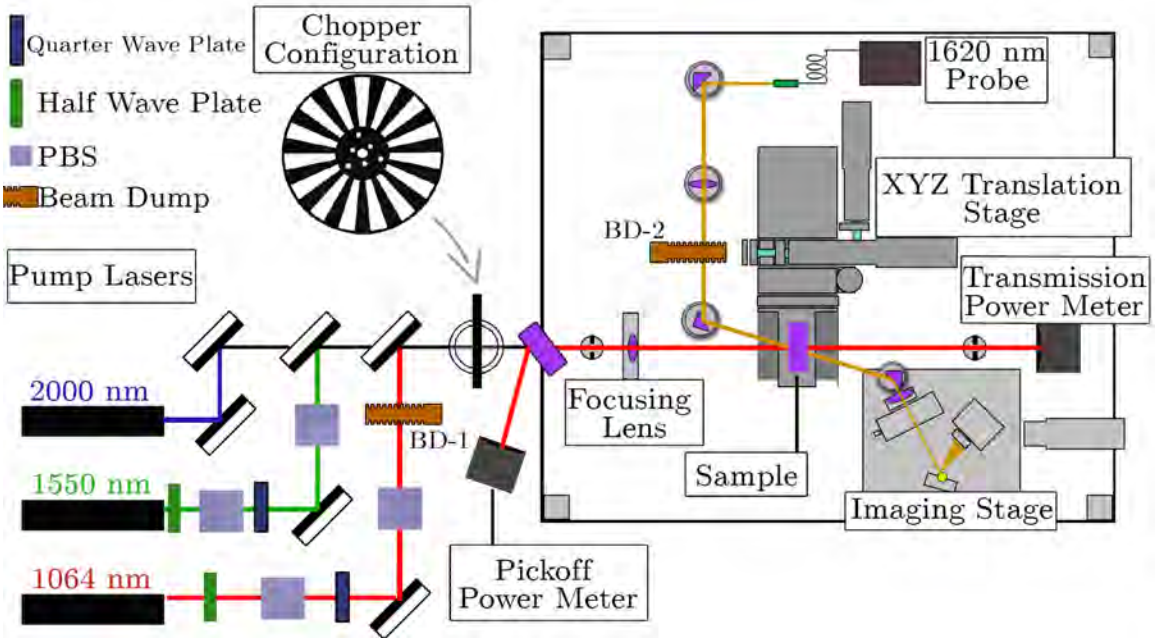

Figure 4.18: A diagram of the PCI setup. Locations of temporary beam dumps (BD-1,BD-2) are denoted in orange, with polarising beam splitter cubes shown as blue squares.

4.1.4.2 Power Fluctuation Compensation

The power transmitted through the sample is recorded by a power meter located behind the sample (see Fig 4.18). As the power incident on the meter is affected by the sample's reflectivity and absorption, the power incident on the sample cannot be accurately estimated from this data - particularly for highly absorbing samples. To monitor the input power, a separate uncoated silica optic (C7979), oriented close to Brewster's angle, was placed in front of the PCI, reflecting a small portion of pump laser light ($\approx 3\%$). The light is reflected at a small angle such that the beam size and position in the PCI is unaltered while maximising power transmission and allowing the reflected beam to be easily measured by a laser power meter. The power reading from this power meter is then un-influenced by the transmission of an absorption sample.

To confirm that the absorption and thermal expansion of the 'pickoff optic' does not change within the expected pump power range, the linearity of the reflected/transmitted powers was tested at increasing pump laser power levels. In each case, the pump laser's output was increased to a given value and allowed to

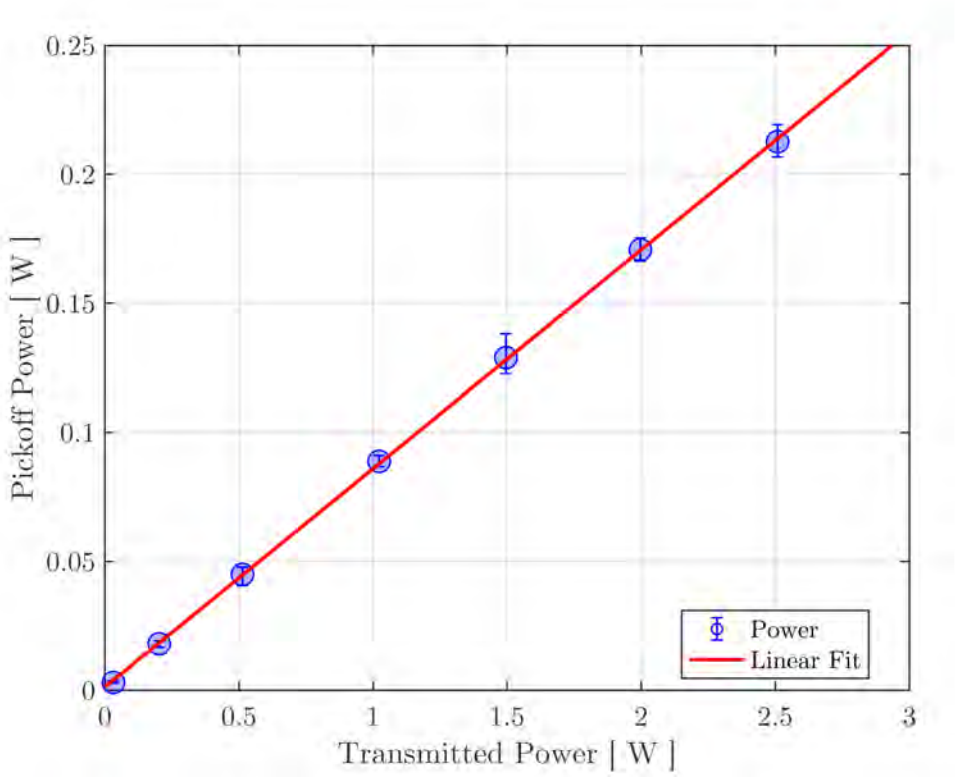

Figure 4.19: Comparison of powers measured from pickoff and transmitted parameters with increasing pump power at 1064 nm.

stabilise. Without a sample in the PCI, the powers recorded on the pickoff and transmitted power meters are compared for ≈ 60 s, allowing effects of the 'pickoff optic' to be tested at different pump powers. The pickoff power and transmitted power were measured for pump laser powers between 30 mW and 2 W. As shown in Fig 4.19, a strong linear correlation was observed.

The errors on each measurement show that there is also a dependence on the stability of the pump power output with different levels of requested pump power. As each power meter's response is linear, it can be assumed that each power meter is functioning correctly and that the measured instabilities in power are produced from the pump output itself.

This shows that the second power meter and pickoff optic can then be used as a reference for the incident pump power, independent of the test sample, without interfering with the measurement process. This allows the level of input power to be

re-constructed from both data-sets. The re-constructed input power can then be used in Eq. 4.6 to correct the calculated absorption values with the actual input power on the sample during a measurement run.

To allow the power correction to be applied correctly, it is essential to synchronise the data from the power meter and the absorption-measurement data, which is produced by a commercial program without any time-stamp. A simple method in which beam dumps (BD1 and BS2 in Fig 4.18) are used to momentarily block the beams and produce a sharp decrease in all recorded signals at the same time. If each sensor's sampling rates were nominally the same, these common artefacts would allow each stream to be synchronised.

Fig 4.20(b) shows an example of how the input power is calculated as a function of the transmitted power and the pickoff power during the first 20 seconds of the measurement (see Fig 4.20), with no sample present inside the PCI. A data synchronisation script written by Joseph Bayley (University of Glasgow) was used to mark and match each of the regions in each data-set to produce a full set of synchronised data. Taking the recorded, transmitted power and calculated input power for this time shows that the calculated power reflects the transmitted power with an average deviation of 0.001%, and can therefore be considered an accurate representation of the pump power.

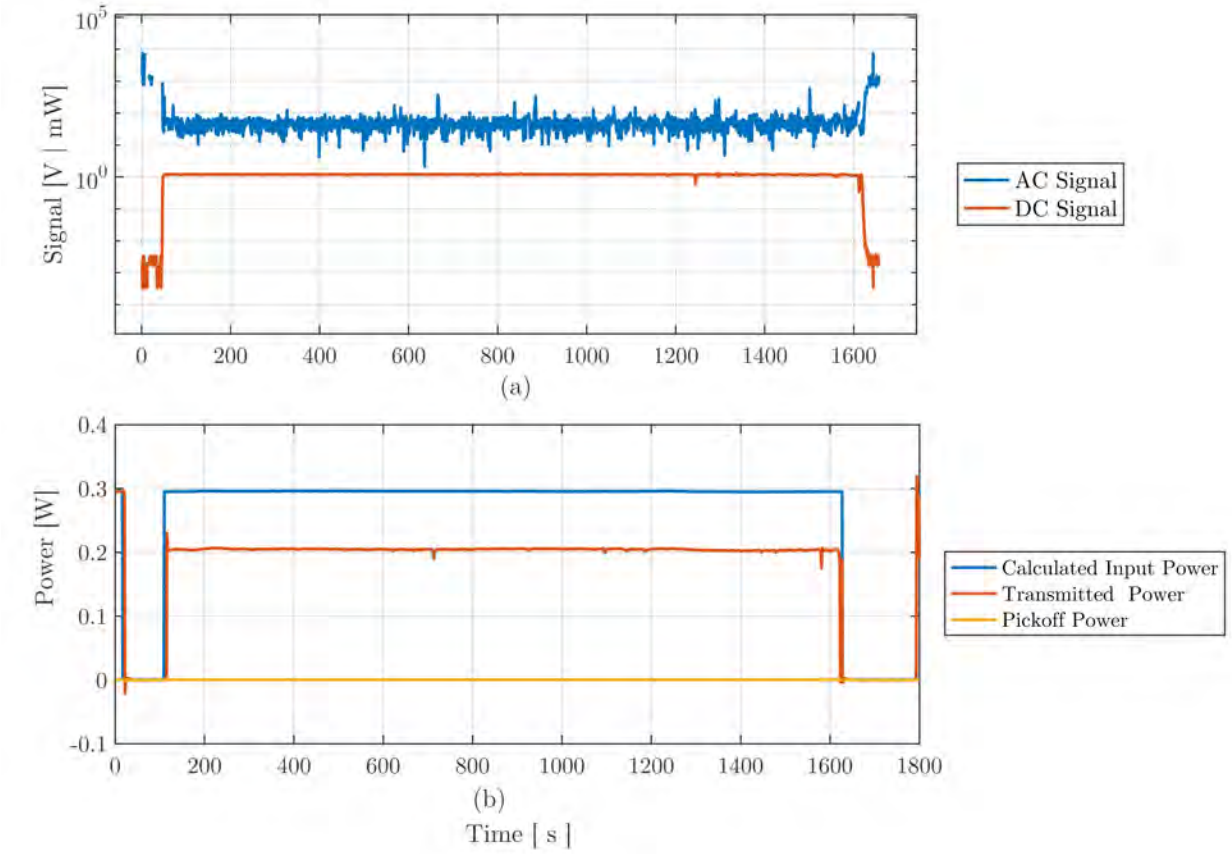

Figure 4.20: (a) AC and DC signals measured with the PCI photo-diode which have been synchronised with data from the pickoff, transmitted power meters shown in (b).

Once the data has been synchronised a surface map of the sample can be produced by moving the absorption sample in the XY plane relative to the pump/probe crossing point. The power correction to the measured absorption data using Eq. 4.6 can then be applied. It is important to note that this scanning method cannot provide longitudinal scan information. Therefore, it is susceptible to small deviations in the samples angle relative to the incident pump laser, and the measurement would need to be corrected. To account for any tilt of the sample with respect to the pump laser and ensure the beam crossing point remained in the coating across the whole map, four longitudinal absorption scans are taken at the edges of the scanning area. The shifting of the coating surface position in z can then be translated into a shift position as the motors translate during the surface map, accounting for the sample's angle.

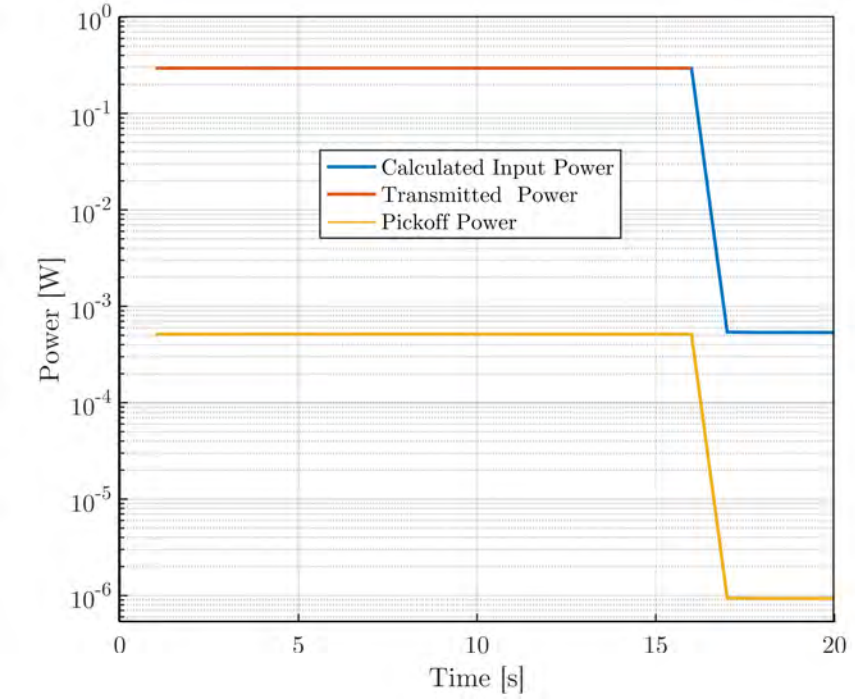

Figure 4.21: First 20 seconds of data shown in Fig 4.20(b).

An example of a final, power-corrected surface map, $6\,\text{mm} \times 6\,\text{mm}$ is shown in Fig 4.22, where the motor step size was $200\,\mu\text{m}$. The colour scale is used to represent the absorption at every point on the surface. The grey regions correspond to positions at the start and end of the absorption scan where the beam was blocked for synchronisation purposes and have been omitted from further analysis.

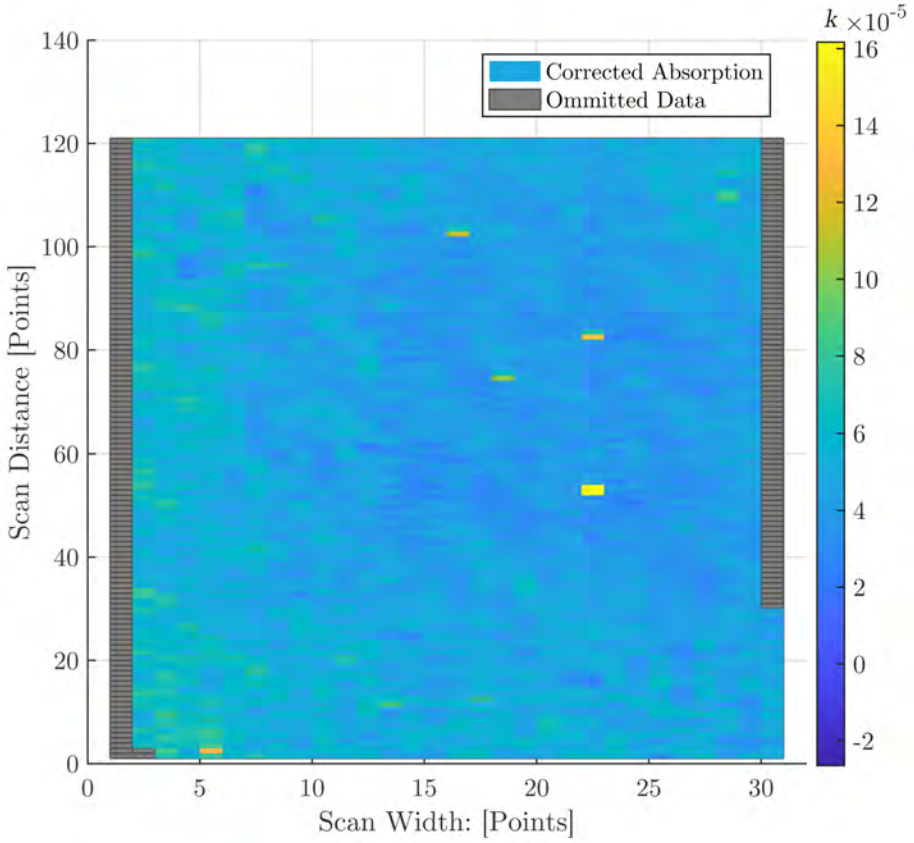

Figure 4.22: 2D map at 1550 nm of Sample D heat treated in vacuum for 3 hrs at 400 °C.

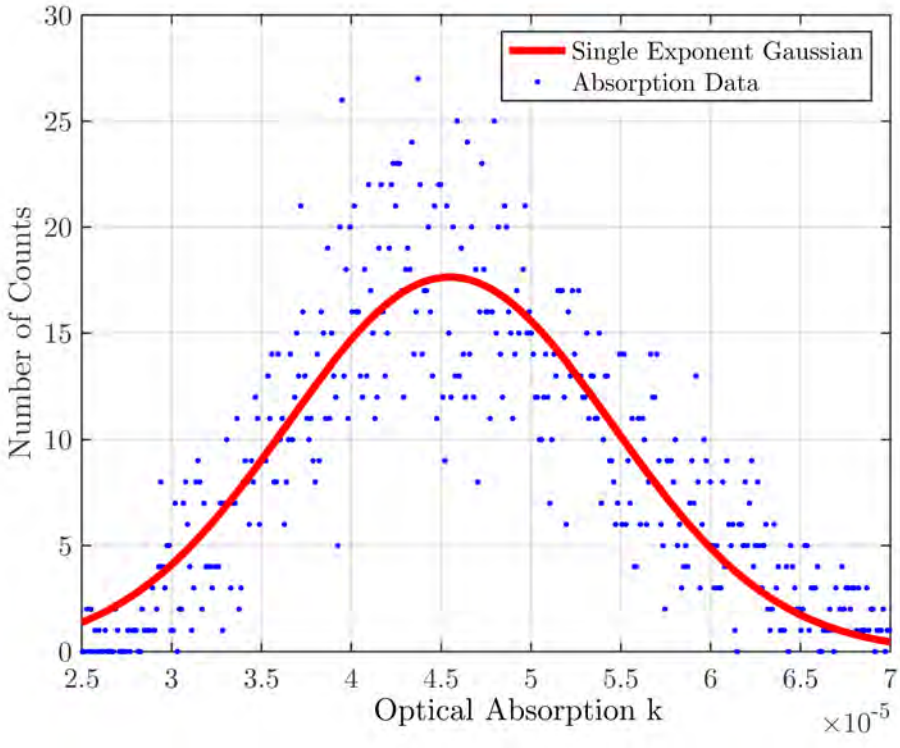

Figure 4.23: Histogram of the absorption data shown in Fig 4.22.

Previous measurements calculated k as 4.55×10^{-5}, when sampling only 4 points. The mapping method allows statistical analysis to be performed on absorption data showing how the absorption varies across the surface of the sample. From a histogram of the data shown in Fig 4.23, k is calculated as 4.53×10^{-5} and the spread in absorption values is the estimated the standard deviation σ=1.29×10^{-5} where σ is calculated from the Gaussian fit of the data.

In Figures 4.24 and 4.25 the average k in each direction is shown. It should be noted that x and y axes are reversed with respect to standard Cartesian coordinates. The error bars on each point take into account the spread and absorption for one line of the surface map. As the conversion from α to k also requires knowledge of the refractive index and thickness of the coating layer, these error bounds are also subject to the uncertainty in these material properties and are calculated using the methods described in section 4.1.4.3. In Fig 4.25 each single vertical line scan is produced successively as a function of time, giving some information on any degradation in the initial calibration.

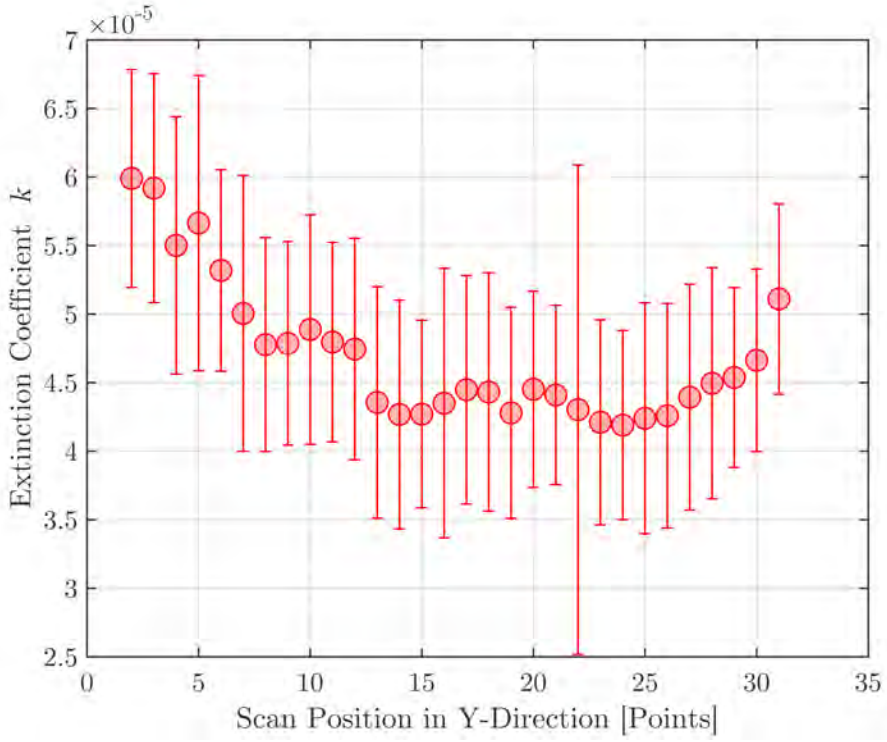

Figure 4.24: Average absorption variation of Sample D heat treated in vacuum for 3 hrs at 400 °C. Error-bars are calculated from the standard deviation in optical absorption for each line across the sample and variations in material properties using the method in section 4.1.4.3.

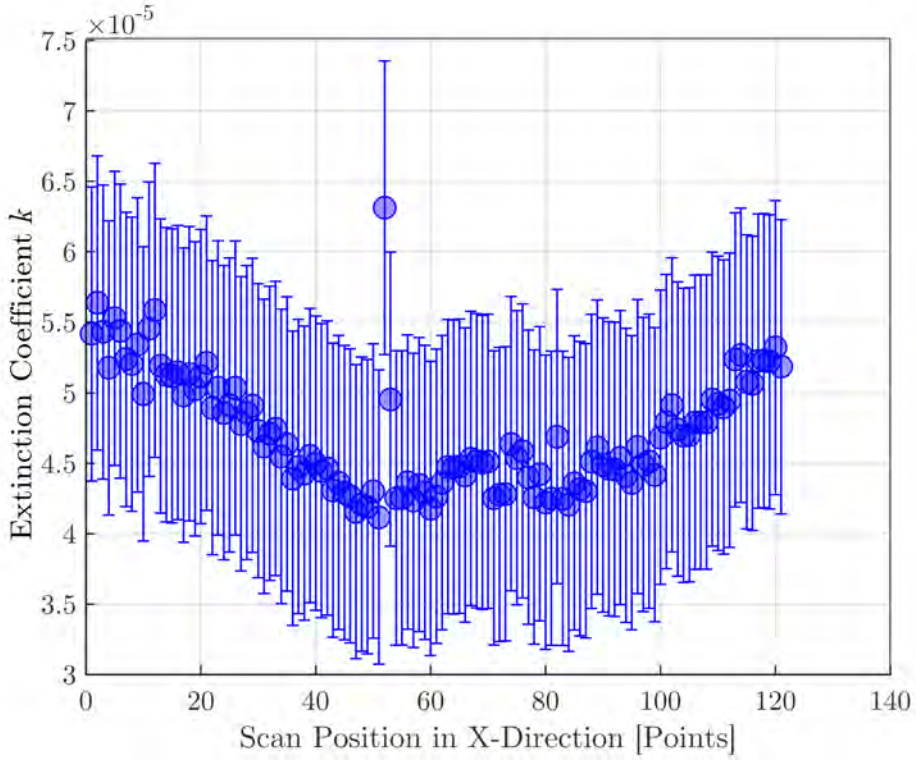

Figure 4.25: Average absorption variation of Sample D heat treated in vacuum for 3 hrs at 400 °C. Error-bars are calculated from the standard deviation in optical absorption for each line across the sample and variations in material properties using the method in section 4.1.4.3.

Temperature fluctuations in the lab, and thermal loading of the optics due to the high laser power, both of which lead to changes in the precise position of optical components. This causes each optic's positions to shift, changing the alignment of the system and thus the calibration. To account for any potential degradation, measurements of the calibration reference before and after the map data was taken is interpolated. As shown in Fig 4.26, the change in calibration as a function of time is approximated as linear with time. In this data, a large gap occurs between the initial calibration measurement and the start of the measurement. In this time the PCI system was being configured for mapping.

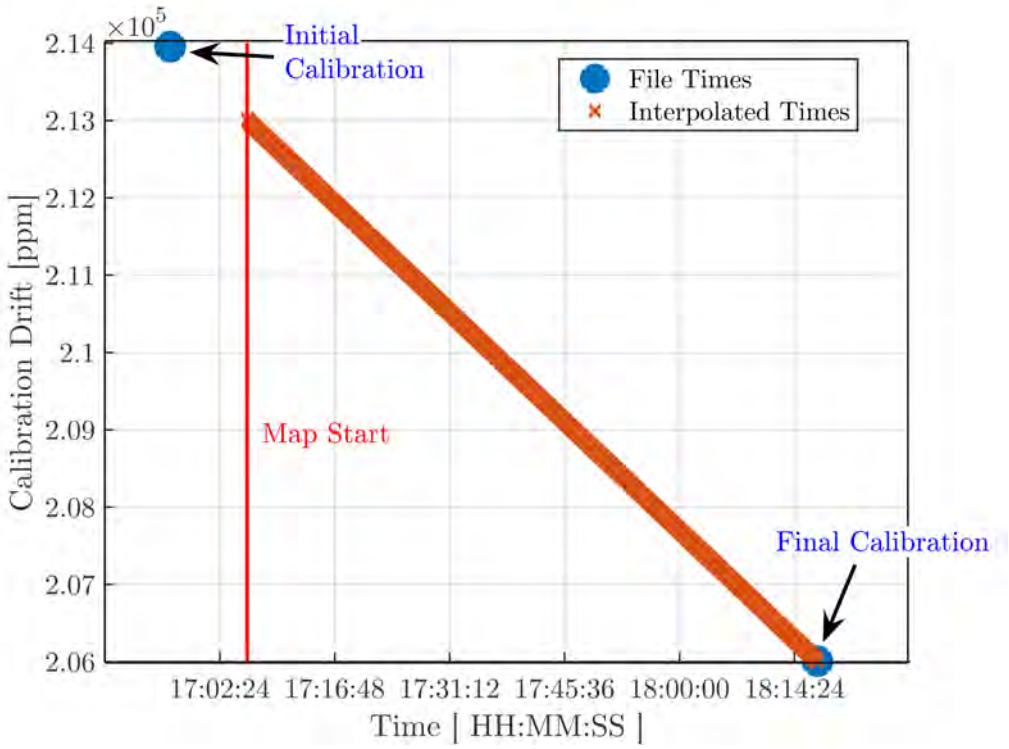

Figure 4.26: Linear interpolation of PCI calibration drift throughout a surface map.

4.1.4.3 Calculating *k* from Measured Absorption

The measured absorption at each point can be translated to a k value, based on the transmission at each point and taking account for the total electric field strength in the coating at each point, as described in section 4.0.2.

A parameter space estimation script written by Zeno Tornasi (University of Glasgow) uses upper and lower bounds placed on the thickness and refractive index of a coating to calculate a corresponding electric field intensity inside the coating. At a given point, the absorption (in ppm as measured by PCI), and the input and transmitted power on the sample are recorded. As thickness and refractive index are often not accurately known, both parameters are varied within reasonable bounds. The electric field intensity (see Eq. 4.10) is then calculated for thickness and refractive index pairs resulting in the measured transmission ratio. From this, the attenuation coefficient, k can be obtained. The accuracy of the final k produced by this analysis is directly related to how accurately the coating's thickness and refractive index are known. Coatings produced by commercial vendors usually show

rather small variations in coating thickness across the sample, reducing the searchable parameter space and the final 'spread' on the calculated k value. In contrast, coatings produced by research institutions may show larger spreads.

An example of one measurement point analysed with this method is shown in Fig 4.27. Each point is representative of a given set of absorption and refractive index values. The x and y axes represent the ranges given for both parameters. The 'branches' represent the refractive index and thickness combinations that can result in the measured transmission. It can be seen that this narrows down the parameter space significantly.

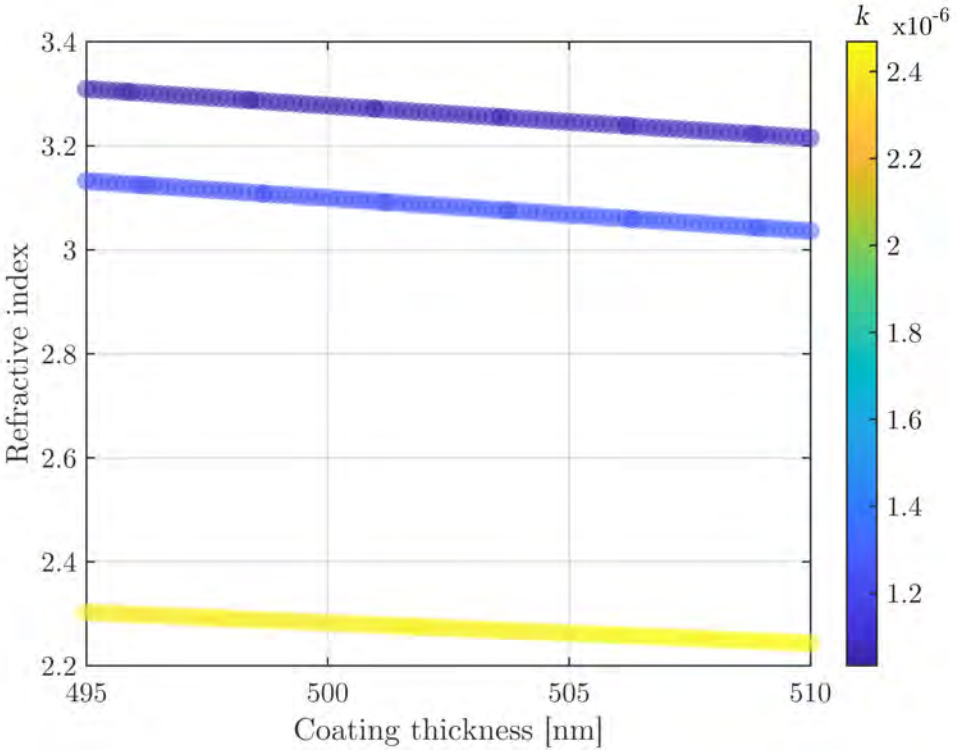

Figure 4.27: Calculated parameter space for one point measured on an RLVIP Ta_2O_5 sample at 1550 nm with 14 ppm raw absorption.

Measurements of the coating thickness were carried out using ellipsometry which shows a thickness of $t_{coating} = 500\,nm \pm 5\,nm$, with a refractive index $n = 2.17 \pm 0.05$. These results narrow down the parameter space, allowing all but the lowest branch to be eliminated. Re-running the analysis with these tightened bounds produces $k = (2.28 \pm 0.02) \times 10^{-6}$. This parameter space estimation can be used to calculate k

for all measured points of a surface map, as shown in Fig 4.22.

4.2 Multi-Material Coatings Produced with RLVIP

Multiple layers, which have low optical absorption and a high contrast in refractive index $C = \frac{n_H}{n_L}$, are required to achieve the $> 99.999\%$ reflectivity of an end test mass mirror. In aLIGO and AdV, multiple alternating layers of SiO_2 and titania-doped tantala (Ti:Ta_2O_5) deposited by Ion-Beam Sputtering (IBS), are used for their favourable optical and mechanical properties. The contrast in their refractive indices, $n_{\text{Silica}} = 1.44$ and $n_{\text{Tantala}} = 2.12$, allows the high reflectivity to be achieved with approx. 38 layers of material. In aLIGO this coating produces an optical absorption of $\sim 0.5\,\text{ppm}$ [159] with the Ti:Ta_2O_5 layers producing the highest contribution to α_{coating}.

If a high refractive index material such as aSi ($n = 3.5$ - 4) [185, 186] would be used in place of Ti:Ta_2O_5, the total number of layers required for high reflectivity could be reduced significantly, thus reducing the coating thermal noise. Amorphous silicon (aSi) is a high-index material ($n = 3.5$ - 4) which has a significantly lower (up to a factor of 4) mechanical loss than Ti:Ta_2O_5 [130, 187], making it a potentially very attractive replacement coating material. In addition to direct thermal noise reductions due to lower loss, replacement of Ti:Ta_2O_5 with aSi would allow the required reflecitvity to be obtained with fewer layers, reducing the coating thickness and so further reducing thermal noise. However, the optical absorption of aSi is significantly higher than that of Ta_2O_5. While it can be significantly reduced when using a longer wavelength of light e.g. 2000 nm [188], and much progress has been made in reducing the absorption [7, 149, 189] through optimising deposition methods and appropriate heat treatment, it remains too high to allow Ta_2O_5 to simply be replaced by aSi.

A novel type of coating design, known as 'multi-material coatings', has been proposed to take advantage of the low absorption of some SiO_2-Ti:Ta_2O_5 layers, to reduce the transmission of laser light into lower layers consisting of more highly

absorbing SiO_2-aSi (by using aSi in the lower part of the coating stack where there is little light power to be absorbed). This type of design which has been theorised by others [134, 146, 188, 190], uses more than two coating materials to produce a HR coating stack. This design allows some of the thermal noise benefits of aSi to be obtained (due to the low loss, and the high index leading to a thinner coating), while preventing the absorption of the aSi layers contributing significantly to the total coating absorption. These materials have been previously investigated as single-layer coatings, but using aSi layers in a multimaterial coating stack is a novel concept [191].

In this section a multi-layer coating stack comprised of SiO_2-Ta_2O_5-aSi layers deposited using RLVIP (see section 4.1.2.6) by Tafelmaier [175] will be characterised for optical absorption at 1064 nm, 1550 nm and 2000 nm. The key results from the measurements at 2000 nm have been published in Physics Review Letters [192]. Fig 6.1 shows the layer configuration of the multi-material coating investigated here, made from 5 bi-layers of SiO_2 and Ta_2O_5, followed by 5 more bi-layers in which the Ta_2O_5 is replaced by aSi. The five bi-layers of SiO_2-Ta_2O_5 will be referred to as the 'Upper-stack', the five bi-layers of SiO_2-aSi the 'Lower-stack' and the total coating as the 'Full-stack'. The thickness of each coating layer is equivalent to $\lambda/4$ optical thickness (see section 4.0.2) at 2000 nm, providing the maximum reflectivity at this wavelength.

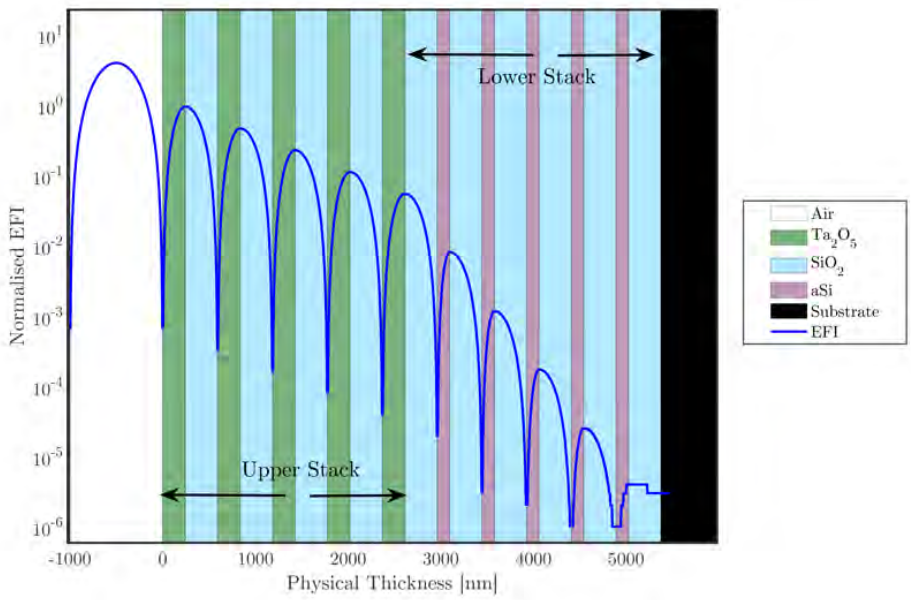

Figure 4.28: Electric field intensity inside the *Full Stack* coating design. Light propagating from left (air) to right (silica substrate).

Fig 6.1 shows the calculated electric field propagation through the coating stack, revealing the reduction in laser light intensity as it is reflected by the *Upper Stack* layers. After the light reaches the *Lower Stack* layers containing the high absorbing aSi, the total electric field intensity should have fallen by >99%, thus reducing the impact of the high absorption of the aSi layers (i.e. an aSi layer will absorb 99% less power here than it would at the top of the coating stack). Transmission measurements of the final deposited coatings show that the actual reduction is closer to ∼96% [192].

From previous measurements of optical absorption of single layers of the same materials used in the coating stack, it was shown that the as-deposited level of absorption can be reduced by heat-treating the material above its deposition temperature. Absorption measurements of single layer Ta_2O_5 are detailed in section 4.1 where the lowest optical absorption was produced after heat treating the material to 500 °C for ≈3 hrs (see Fig 4.8). In a similar study carried out by J. Steinlechner *et al.* it was found that the optimal heat treatment temperature for aSi was also ≈ 500 °C [149]. However for a multi-material coating stack, composed of three different materials, it is still difficult to predict potential changes with heat

treatment exactly.

4.2.1 Absorption with Heat Treatment

In this section, measurements of the effect of heat treatment on the absorption of the *Upper Stack*, *Lower Stack* and *Full Stack* coating are reported. Two separate studies were carried out:

- 'Direct heat-treatment' where the sample was taken from the as-deposited state to an elevated temperature for three hours (e.g. AD $\rightarrow$ 400 °C)
- 'Incremental heat treatment' where a sample was heat treated in 100 °C steps until it reached the same temperature (e.g. AD $\rightarrow$ 100 °C,$\rightarrow$ 200 °C etc.) for the same duration.

Samples which were directly heat-treated to a given temperature were measured, by Jessica Steinlechner for optical absorption at a laser wavelength of 2000 nm, while samples which were incrementally heat treated were measured at 1064 nm, 1550 nm and 2000 nm by the author and by another PhD student, Graeme McGhee, working under the author's supervision.

Fig 4.29 shows the measured absorption of each stack at 2000 nm following direct heat treatment. The *Lower* and *Full* stacks, deposited on C7980 glass, when heat-treated at 100 °C and 200 °C show no noticeable change in measured absorption. This is likely to be due to the coatings reaching temperatures of up to $\approx$ 200 °C during deposition, and suggests any effect of heat treatment at this temperature may already have occurred during deposition. Heat treatment above this level for both the *Full* and *Lower* stacks result in the absorption reducing to a minimum at 500 °C. At this temperature, the *Lower Stack's* absorption decreased by a factor of 7 compared to the as-deposited measurement, to a value of 181 ppm. The *Full Stack* shows a larger decrease by about a factor of 18 at 500 °C. It should be noted that in sampling multiple points across the surface of the *Full Stack* coating after heat treatment at 600 °C, the absorption was found to vary in some cases by more than 70%, a possible indication of coating crystallisation.

The *Upper Stack*, which was expected to have the lowest absorption, was deposited on Corning C7979 glass to keep the contribution of the substrate absorption low.

The absorption of the *Upper Stack* was only measured in the as-deposited state and at the optimum heat treatment temperature of 500 °C as the functional form of its absorption with heat treatment was not expected to change from previous single layer measurements (see section 4.1).

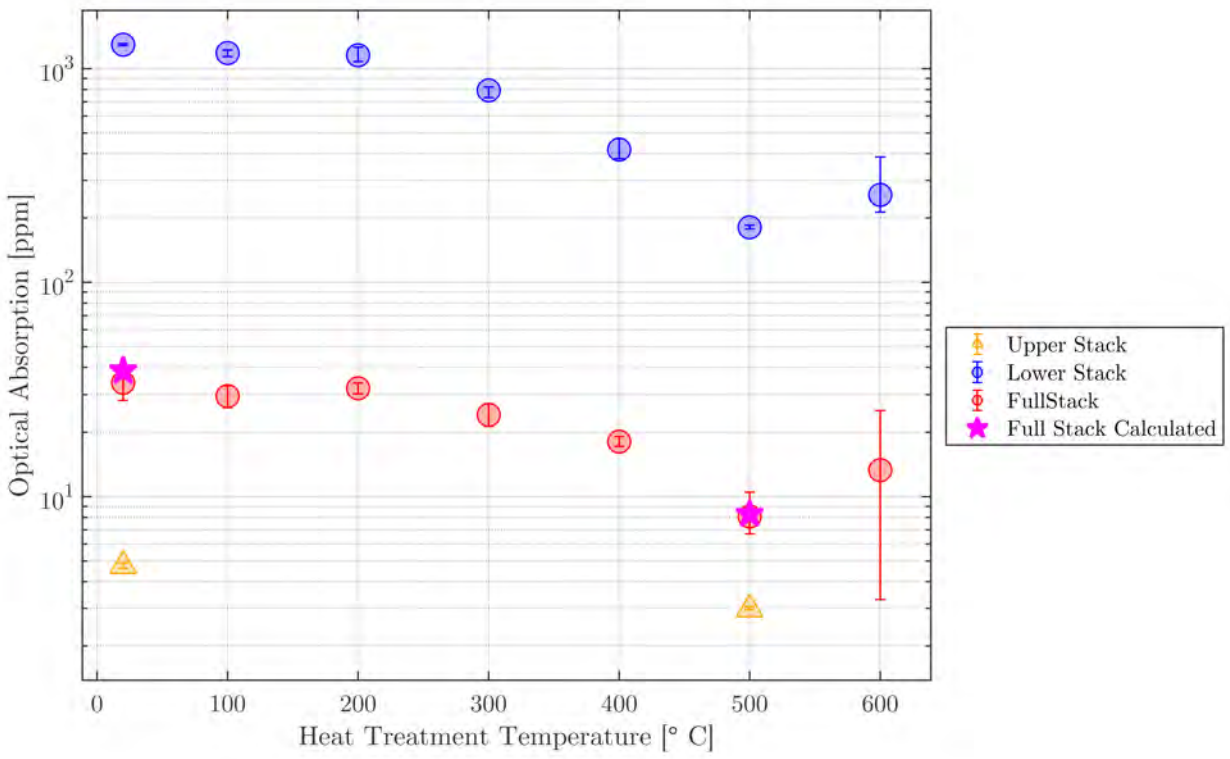

Figure 4.29: Absorption at $\lambda_{\text{pump}} = 2000\,\text{nm}$ of the *Full Stack*, *Upper Stack* and *Lower Stack* coatings with direct heat treatment temperature. Values for 'Full Stack Calculated' were calculated from the separate measurements on the *Upper* and *Lower* stacks.

To investigate if there were significant changes in the absorption of the *Full Stack* coating relative to the two-component stacks, the expected optical absorption of the *Full Stack* was calculated from measurements of the *Upper Stack* and *Lower Stack*, as-deposited and at 500 °C. In each case, the calculated values show considerable agreement with the values measured on the *Full Stack* coating, suggesting no additional interface-effects due to combining the two partial stacks. Carrying out the same calculation using absorption values previously measured for a single layer of each material, where the absorption of aSi dominates over that of SiO_2, predicts a more considerable decrease in absorption with heat treatment. With the absorption of SiO_2-Ta_2O_5 in previous works showing no signs absorption

related to interface effects when structured in an HR coating stack, this leaves the aSi layers for further investigation. Measurements of aSi single layer coatings produced by ion-beam sputtering, a comparatively high energy deposition method, have also shown that a single aSi layer reduced by about a factor of 2 more in absorption with heat treatment than an aSi - SiO_2 bi-layer [7]. This suggests that surrounding layers can suppress the ability of aSi to change with heat treatment and may explain why materials containing aSi show less absorption reduction with heat treatment than expected from single-layer measurements of aSi.

The optical absorption of the 'direct heat treatment' samples was then compared to identical samples that were incrementally heat-treated and measured at 1064 nm, 1550 nm and 2000 nm using PCI.

Measurements of the *Upper Stack*, *Lower Stack* and *Full Stack* at 1064 nm, incrementally heat treated, are shown in Fig 4.32. At this wavelength, each coating layer's thickness is no longer optimised for the reflectivity of the incident electric field. As a result, there is a larger light field inside the high absorbing aSi layers in the *Full Stack* and *Lower Stack* coatings (see Fig 4.30). The increased contribution to the absorption of the aSi layers is made clear in both the *Lower Stack* and *Full Stack* measurements at this wavelength which are more than 10 times higher than the *Upper Stack* material which does not contain aSi. Heat treating each sample at 200 °C induces a negligible change in the *Upper Stack*, curiously heat treatment at this temperature causes the absorption of the *Lower Stack* and *Full Stack* coatings to increase by 50% and 12.6% receptively. Given that this absorption increase only occurs in coating stacks which contain aSi layers, it seems plausible that this heat treatment caused the optical absorption of these layers to increase.

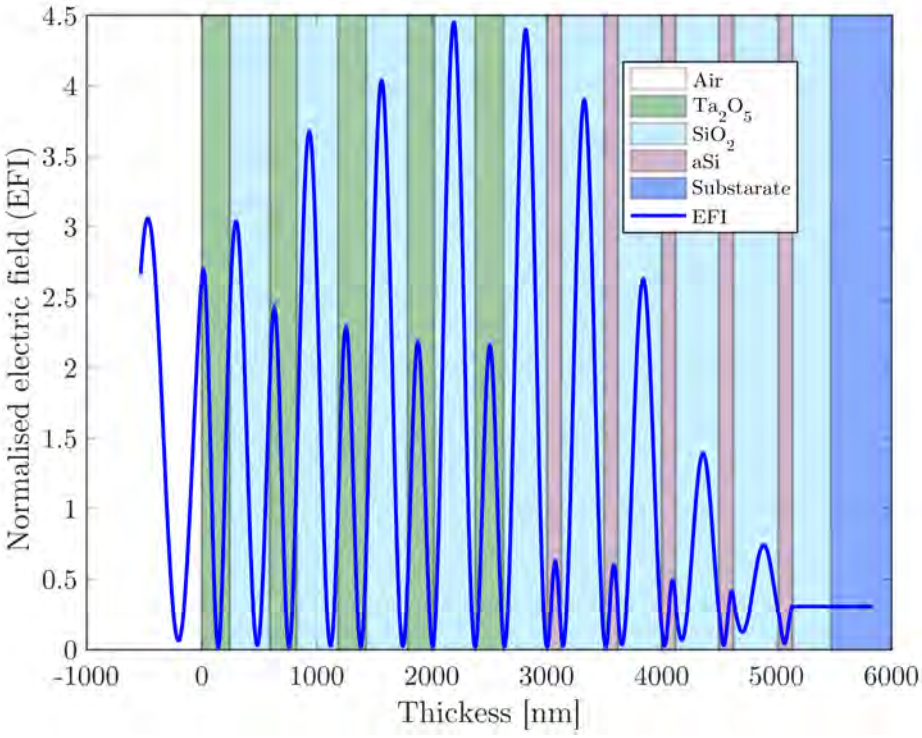

Figure 4.30: Normalised electric field intensity of the *Full Stack* coating at 1064 nm on a silica substrate.

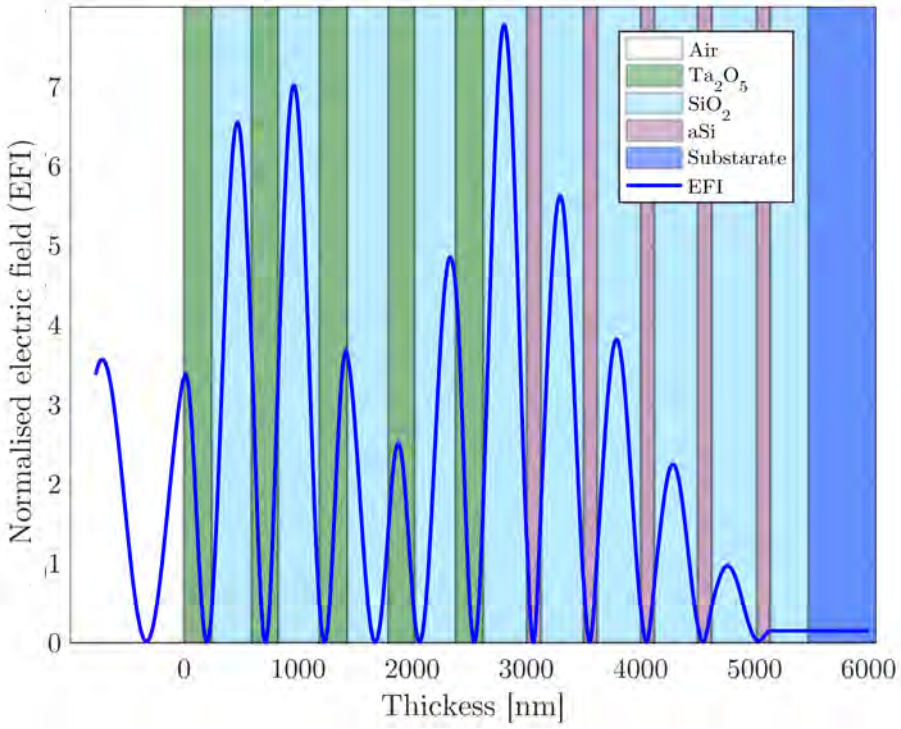

Figure 4.31: Normalised electric field intensity of the *Full Stack* coating at 1550 nm on a silica substrate.

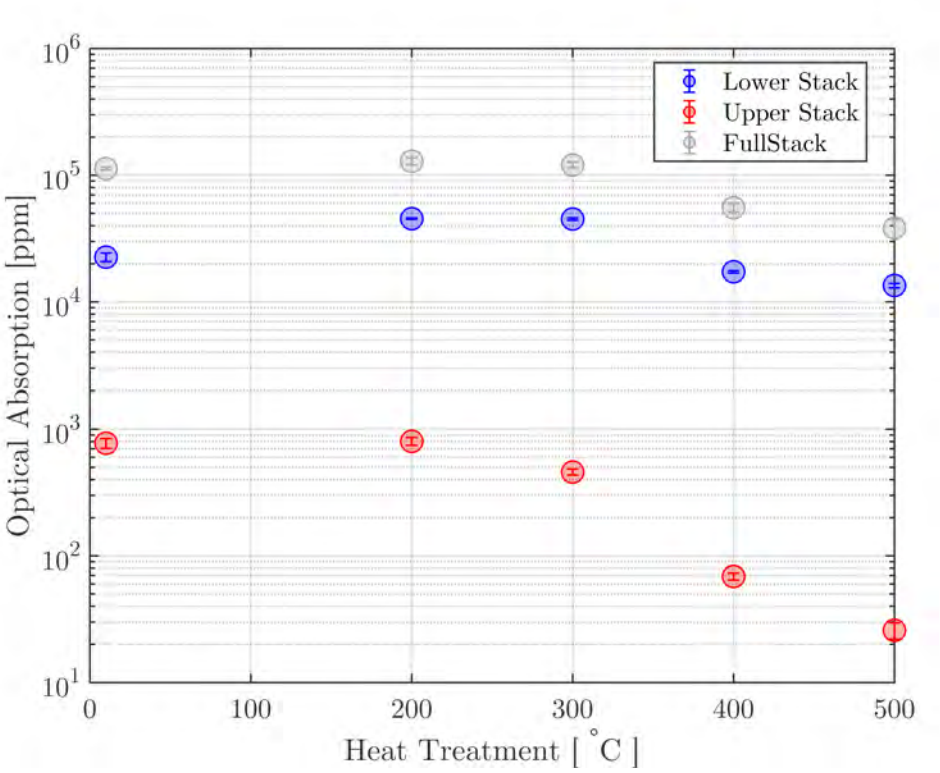

Figure 4.32: Optical absorption for at 1064 nm for the *Upper Stack*, *Lower Stack* and *Full Stack* after incremental heat treatment.

After heat treating each sample to 300 °C the absorption of the *Upper Stack* and *Full Stack* coatings decreases by 43% and 6.5% respectively, with the *Lower Stack* coating remaining unchanged within experimental error. In the *Full Stack* coating at this wavelength the electric field intensity in Ta_2O_5 layers is more than 4 times higher than for the aSi layers. This makes the absorption of the *Full Stack* liable to two different effects at this wavelength; an increase in k of both Ta_2O_5 and aSi compared to at 2000 nm; and a change in the EFI in each layer. Each stack was then heat-treated a further two times, with all stacks continuing to decrease in absorption at 500 °C.

Comparing these measurements to the absorption of each sample at 1064 nm, the observed trend with heat treatment of the *Upper Stack* coating mirrors the observations at 1550 nm and 2000 nm, with a minimum value found at 500 °C. The behaviour for the *Full Stack* and *Lower Stack* coatings, however, differs from results measured at the other wavelengths with the *Full Stack* showing the highest absorption. This is again an effect of the non-optimised layer thickness at this

wavelength. At 1550 nm, the electric field intensity in all aSi layers is reduced, lowering their contributions to the measured optical absorption. The first heat treatment of the *Lower Stack* again shows an increase in absorption, similar to the observations made at 1064 nm. With each subsequent heat treatment, the absorption shows a decreasing trend, with the lowest recorded value occurring at 500 °C. It is interesting to note that a comparable absorption level was measured on the *Lower Stack* at 300 °C.

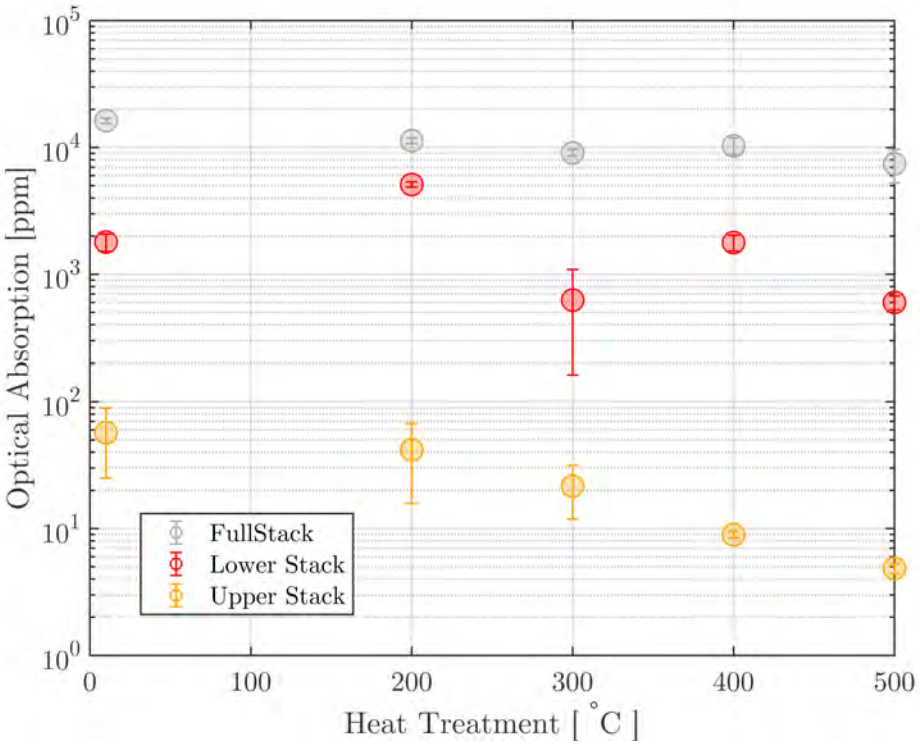

Figure 4.33: Optical absorption for at 1550 nm for the *Upper Stack*, *Lower Stack* and *Full Stack* for incremental heat treatment.

The *Full Stack* coating shows a broadly similar trend to the observations made at 1064 nm. It is interesting to note that at 300 °C this coating stack decreases by ∼20%, before increasing after the subsequent heat treatment at 400 °C by ∼12%. This change in absorption is analogous to the behaviour of the *Lower Stack* coating, with the expected change in absorption muted by the dominating absorption to the Ta_2O_5. As both the *Lower Stack* and *Full Stack* coatings exhibit this behaviour at 1550 nm, but the *Upper Stack* does not, this suggests that this may be a real effect of the aSi layers.

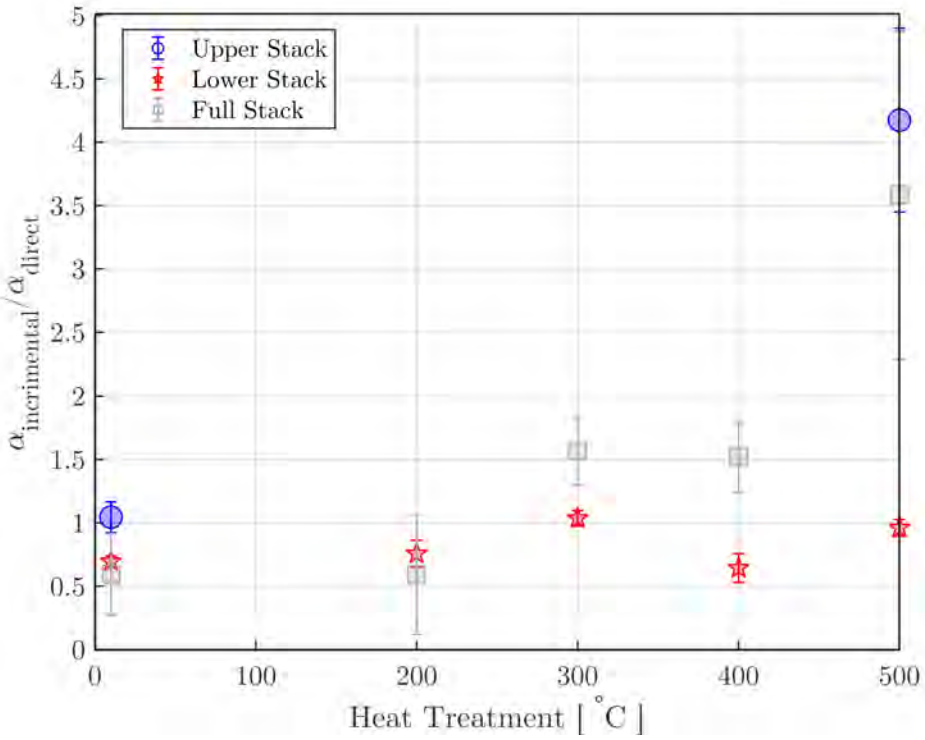

Figure 4.34: Ratio of optical absorption at 2000 nm for the *Upper Stack*, *Lower Stack* and *Full Stack*, measured after direct heat treatment and incremental heat treatment.

Fig 4.34 shows the ratio of absorption at 2000 nm taken on the *Upper Stack*, *Lower Stack* and *Full Stack* via direct and incremental heat treatment where any value larger than one indicates a higher absorption measured for incrementally heat-treated samples. Measurements of the as-deposited *Lower Stack* and values at 200 °C and 300 °C show little to no change outwith measurement error for both heat treatments with values varying around one. The large error noted for the *Full Stack* stems from the spread of absorption values from the incremental heat-treated sample. As the *Upper Stack* coating was only previously measured as-deposited and at 500 °C only these measurements can be compared to the previous heat treatment study.

Coating Material	Wavelength (nm)	Minimum Optical Absorption (ppm)	k	Optimum Temperature ($^{\circ}C$)
Full Stack	1064	$(3.8 \pm 0.6) \times 10^4$	$5.0 \times 10^{-3} \mp$	500
	1550	$(7.4 \pm 0.2) \times 10^3$	$2.5 \times 10^{-3} \mp$	500
	2000	19.0 ± 9.0	$2.0 \times 10^{-3} \mp$	200
	2000	8.0 ± 2.0	$6.5 \times 10^{-4} \mp$	500 (direct)
Lower Stack	1064	$1.35 \times 10^4 \pm 540$	2.1×10^{-3}	500
	1550	600 ± 75	4.2×10^{-4}	500
	2000	174 ± 10.8	3.4×10^{-4}	500
	2000	180 ± 3.5	3.6×10^{-4}	500 (direct)
Upper Stack	1064	26 ± 3.8	2.1×10^{-6}	500
	1550	4.9 ± 0.4	7.5×10^{-7}	500
	2000	5.0 ± 0.6	2.2×10^{-6}	AD
	2000	3.0 ± 0.1	1.3×10^{-6}	500 (direct)

Table 4.5: Summary of optical absorption of multimaterial coating stacks measured at 1064 nm, 1550 nm and 2000 nm. All absorption values are from incremental heat treated samples unless stated otherwise. Values marked with '$\mp$' are calculated assuming absorption of aSi dominates.

The results from both the direct and incremental heat treatment studies are detailed in Table 4.5. Comparing the measurements made at 2000 nm for both studies for the *Lower Stack* coating produces very similar results with the measured absorption within the measurement error for both cases. The *Upper Stack* at 2000 nm shows large differences for the direct heat treatment and the incremental heat treatment studies. Comparing measurements carried out at different wavelengths reveals that the increase in absorption measured at 2000 nm does not affect all measurement wavelengths in the same manner, with 1064 nm and 1550 nm producing the lowest absorption values at 500 °C.

To allow comparison of our measurements with other coatings, it is also necessary to calculate the extinction k of the coatings. For aSi, k was estimated from the absorption of the *lower stack* under the assumption that the absorption of SiO_2 is negligible compared to aSi. Similarly, k of Ta_2O_5 was estimated from the absorption of the *Upper Stack*. The electric field intensity in the *Full Stack* was simulated using these k values to predict the absorption of the *Full Stack* from the two partial stacks.

At 2000 nm, the energy of each photon allows for optical absorption in O-H molecules in the surface and bulk of each sample. To minimise the contributions of

bulk water content the lower absorption *Upper Stack* coating was deposited on Corning C7979 silica glass, and *Full Stack* coating which was deposited on C7980, with its much higher coating absorption expected to dominate over the effects of bulk water in the sample substrate [193]. The increasing trend of optical absorption at 2000 nm with heat treatment temperature could suggest that after each heat treatment, each coating stack's water content increases, thus increasing the optical absorption. It is then open to speculation that multiple heat treatment cycles could increase each material's porosity, leading to the higher optical absorption, which does not occur in samples that have been directly heat-treated through a single heating cycle. To test this hypothesis, a subsequent study has been proposed by the author which uses two single layers RLVIP Ta_2O_5 samples and a sample of the *Upper Stack* coating on C7979. These samples will be subjected to the same incremental heat treatment steps carried out in this investigation and subsequently stored under different conditions, with one sample stored in the laboratory, consistent with this study and the other single layer sample and *Upper Stack* sample stored in a sample desiccator, minimising their contact with atmospheric water. Comparing the measured optical absorption at 2000 nm with the results produced in this study will help quantify if surface water is correlated with the absorption of these samples.

This study confirms that the multimaterial concept works as expected, and is self-consistent with the *Full Stack* following the expected behaviour as predicted from its component stacks at a wavelength of 2000 nm. While the concept of multi-material coatings have been proposed by others [134, 146, 188, 190] using different materials, these measurements are the first experimental verification that the absorption of aSi layers can be reduced by depositing them below layers of SiO_2-Ta_2O_5. This confirmation has several implications for third-generation gravitational wave detectors who require ≤ 0.5 ppm optical absorption in an ETM coating [78]. These materials could be used to reduce the mechanical loss and thickness of an end test mass coating, having a meaningful effect on the level of coating Brownian noise and overall sensitivity of the detector (see Chapter 5). It has been shown that a multi-material coating stack consisting of SiO_2-Ta_2O_5 bi-layers and SiO_2-aSi bi-layers can be used to reduce optical absorption of aSi layers with the *Full Stack* coating absorbing 8.1 ppm, 95%, lower than the absorption of the *Lower Stack* coating. Each coating was heat treated up to temperatures of 600 °C

with no signs of delamination or deleterious effects observed in the coatings. This result coupled with the higher refractive index of 3.73 and lower mechanical loss of aSi (see Chapter 5) would allow the >99.999% reflectivity requirements of an ETM mirror coating to be achieved with 20 layers of coating material instead of 19 bi-layers (38 layers) of SiO_2-Ti:Ta_2O_5 which are currently in operation in aLIGO detectors, if the absorption can be further reduced. The thermal noise implications in a room temperature aLIGO detector for these coatings is discussed in more detail in Chapter 5. For a cryogenic detector operating operating with the primary interferometer laser at 2000 nm and at temperatures < 123 K the optical absorption of these coatings is expected to further decrease by as much as a factor of three [188], producing a similar level of optical absorption to current SiO_2-Ti:Ta_2O_5 coatings at 293 K.

4.3 $LaTiO_3$ and Zr:Ta_2O_5 Deposited by IBS

As part of a research program targeted at improved coatings for the advanced LIGO+ (A+) and advanced Virgo+ (V+) detectors, several coatings were deposited using ion beam sputtering by collaborators at University of Strathclyde and The University of the West of Scotland. Each material was deposited as a potential replacement for the high index Ti:Ta_2O_5 layers currently used in aLIGO and AdV, which dominate coating thermal noise due to their high mechanical loss. Any reduction in the mechanical loss and absorption of these coating layers would improve the sensitivity and stability of the detector

The IBS deposition system at University of Strathclyde utilises an ion beam formed by injection of argon gas (Ar). Argon is injected via a dosing valve into a quarter-wave resonant microwave (2.46 GHz) cavity, and is ionised via an electron cyclotron resonance (ECR) process [194]. The plasma is then extracted and focused into a beam using electrostatic optics; the liberated material then condenses on a remote substrate to form a growing film on the desired substrate. The range of ion energies achievable using this process is far wider than with a conventional radio frequency-inductive coupled plasma (RF-ICP) ion beam source, and acceleration potentials of up to 20 kV are possible. It is thought that the low deposition rates used in this system can produce coatings with a higher coordination of atoms in the

final coating, compared to other deposition methods[5]. It has been speculated that this is linked to reduced optical and mechanical losses [7].

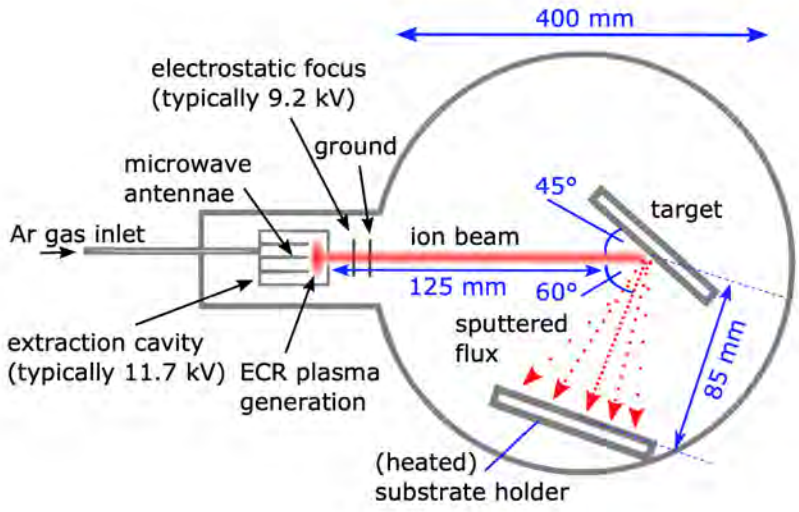

Figure 4.35: Schematic representation of IBS deposition system at Strathclyde University [7].

4.3.1 Optical Absorption of $LaTiO_3$

Lanthanum Titanate ($LaTiO_3$), was identified by R.Birney (University of the West of Scotland) as a potential high-index material to replace Ta_2O_5. To test the materials' optical properties, the optical absorption of a single layer coating deposited on JGS-1 glass[6] ($t \approx 0.2$ mm) was measured at 1064 nm and 1550 nm.

4.3.1.1 Absorption at 1064 nm

Figures 4.36 and 4.39 show the absorption measured at 1064 nm and 1550 nm respectively. The measurements of both wavelengths identified a region in the top left-hand side of the surface map with a much higher absorption compared to other areas of the sample. This spread was then confirmed by viewing the absorption for each vertical scan taken by the PCI, where the average k value shows an increase

[5]It is assumed that some level of self-ordering occurs as a consequence of the slow dep rate. It is thought to enable impinging atoms to find the most energetically favourable arrangement before the next 'layer' of atoms arrives at the surface.

[6]Similar to C7980 substrates used previously.

with each vertical scan of the sample (see Fig 4.38). As k is a thickness independent value, this suggests that the electric field's intensity inside the coating layer is non-uniform across the sample. However, the reverse argument can also be made that the sample's thickness is invariant and is it k which is changing. As the ECR deposition plume is relatively narrow, achieving uniform thickness', even over relatively small areas is challenging. Since the deposition of these coatings, the chamber was upgraded to include a substrate rotation stage which should improve coating uniformity.

Given that previous measurements of commercial coatings made by the author have shown a Gaussian distribution of absorption values (see Fig 4.22), the same fitting procedure was attempted for this measured data and is shown in Fig 4.37. However, due to the large in-homogeneity of the measured absorption, produced by either thickness or k variance, a single Gaussian distribution is no longer a good approximation of the data. A two term Gaussian function (the summation of two Gaussian distributions) allows a better representation of the distribution, with the highest density of points occurring between $k = 4\times10^{-4}$ and $k = 6\times10^{-4}$. By viewing changes in the phase of the light throughout the scan, an estimate of what material the light had passed through (substrate or coating) could be produced. Measurements of the uncoated silica reference sample used for PCI calibration produces light with a phase between -65° and -69° (see Fig. 4.4). Regions of the map which produce a phase within this range are more likely to contain no coating and have been removed from the final analysis. Comparing the histogram to the surface map suggests that the distribution in absorption values at lower values of k are accurate measurements, which could stem from a region of coating which is significantly different in thickness compared to the rest of the coating surface.

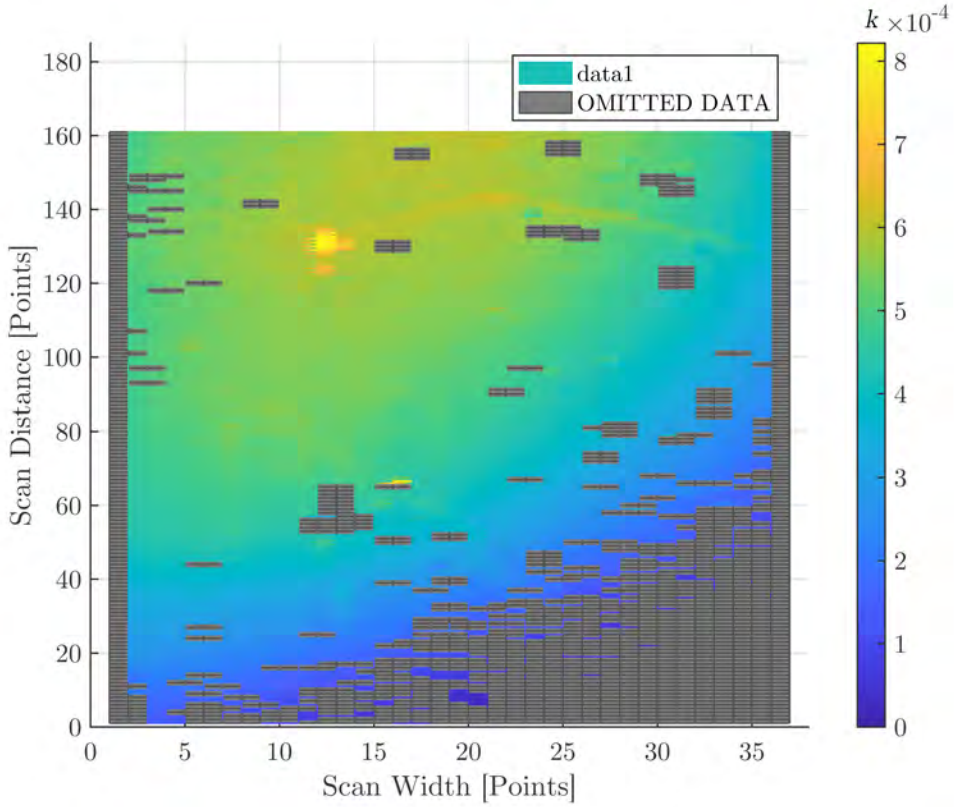

Figure 4.36: Absorption map at 1064 nm of IBS $LaTiO_3$ deposited on JGS-1. Measurements where the probe power or measured phase decreases below tolerable levels are omitted. The mapped area is equivalent to 9 mm × 8 mm.

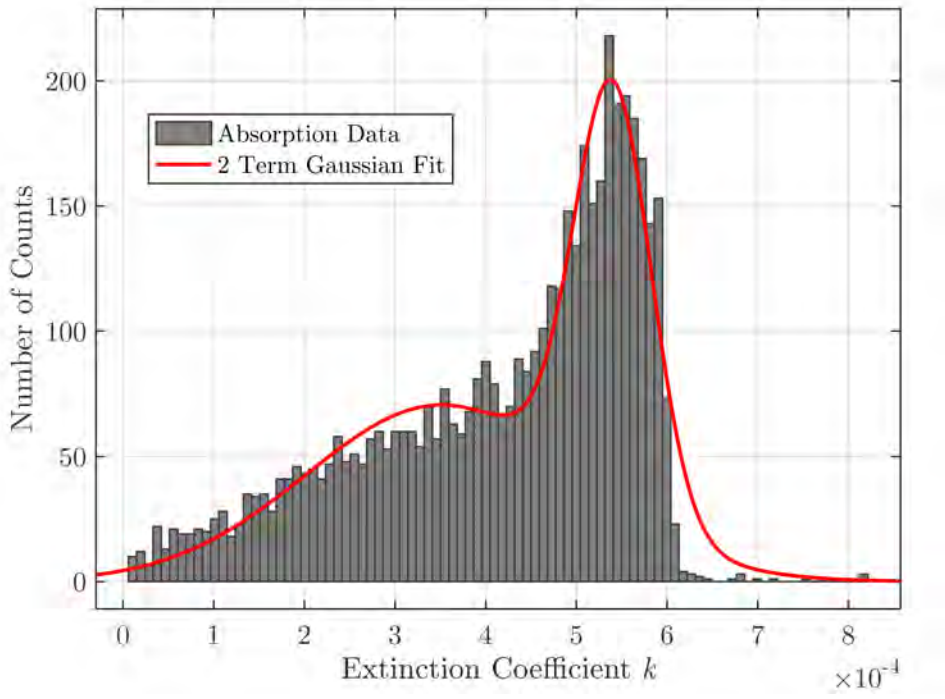

Figure 4.37: Histogram and Gaussian fit of the absorption data shown in 4.36.

4.3.1.2 Error bars

Using the methods described in section 4.0.1, the electric field intensity at each measured point in the coating layer can be calculated. As no coating thickness measurements were provided with the sample, the coating thickness was presumed to be anywhere between 100 nm and 3 µm thick, with literature values suggesting its refractive index could lie between $n = 1.8$ -2.2 [195] in the infrared. This creates a wide parameter space over which the electric field intensity can be calculated. This, in turn, affects the calculated absorption values due to these wide bounds placed on thickness and refractive index, giving a larger associated spread in k.

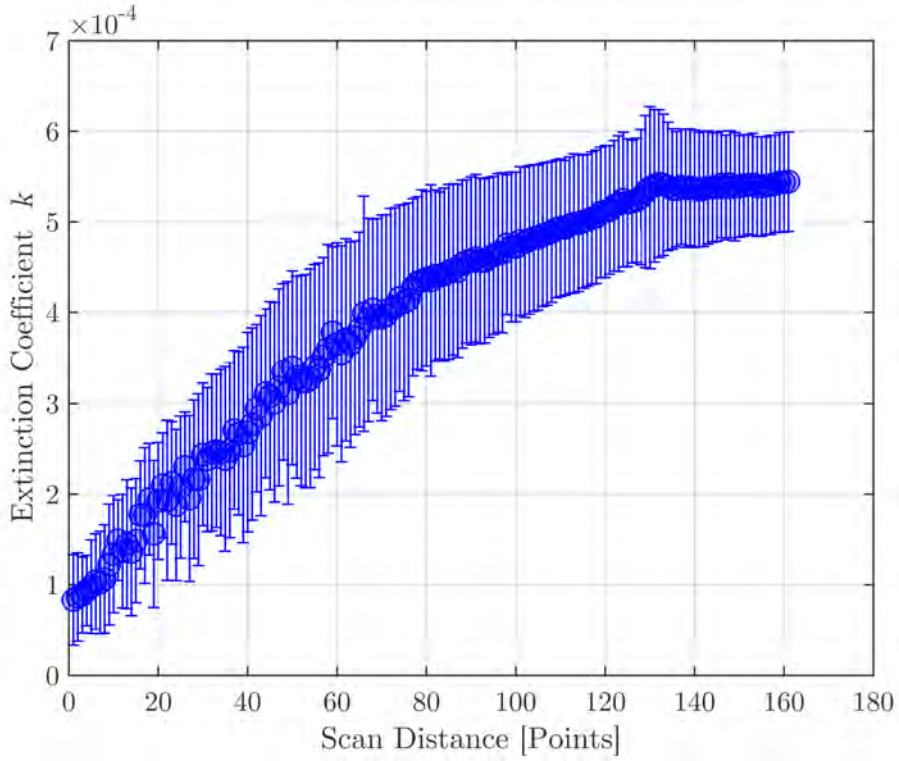

Figure 4.38: Average calculated k at 1064 nm for each vertical scan of the IBS $LaTiO_3$ sample deposited on JGS-1. Error bars denote the standard deviation of k for each vertical scan.

The mean measured absorption at 1064 nm is $k = (5 \pm 5.91) \times 10^{-4}$ (1074 ppm ± 1272 ppm), The lowest absorption measured on the presumably thinner portion of the sample shows a much lower absorption of $k = (4.03 \pm 9) \times 10^{-5}$ (105 ppm ± 300 ppm), which is overshadowed by the non-uniformity in the coating layer. Without a greater understanding of how the coating layer's thickness varies, the bounds placed on these calculations cannot be made more stringent.

4.3.1.3 Absorption at 1550 nm

Examining the coating shows regions that are predicted to have an optical absorption of around 100 ppm, where the coating thickness is thought to be inadvertently optimised to 1550 nm, or the k of the sample is substantially lower. Regions of distinctly higher absorption are measured predominantly in regions where the coating is assumed to be thicker. This scan also reveals a 'spider-web' like structure on the coating surface, which shows a more uniform absorption within the structures. While these patterns are also present at 1064 nm, they are less pronounced at that wavelength. This surface structure may originate from surface scratches or defects in the sample. It is also possible that artefacts from the coating deposition are present in the coating surface, such as voids or small regions of contaminants that are absorbing or scattering light at this wavelength.

Without changing the sample's position inside the PCI, the sample was remeasured at 1550 nm. At this wavelength the absorption appears much more uniform with a spread $\mathrm{k} = (1.16 \pm 1.2) \times 10^{-4}$ ($206\,\mathrm{ppm} \pm 220\,\mathrm{ppm}$). The variation in the absorption across the sample's surface is still present, as clearly visible in Fig 4.39. Under the assumption that there are three different regions across the sample surface which contribute to the absorption, the sample can be described as follows: The first (largest) region covers the majority of the coating surface, which is nominally the same thickness; this could be closer to the plasma plume during deposition producing a uniform region.

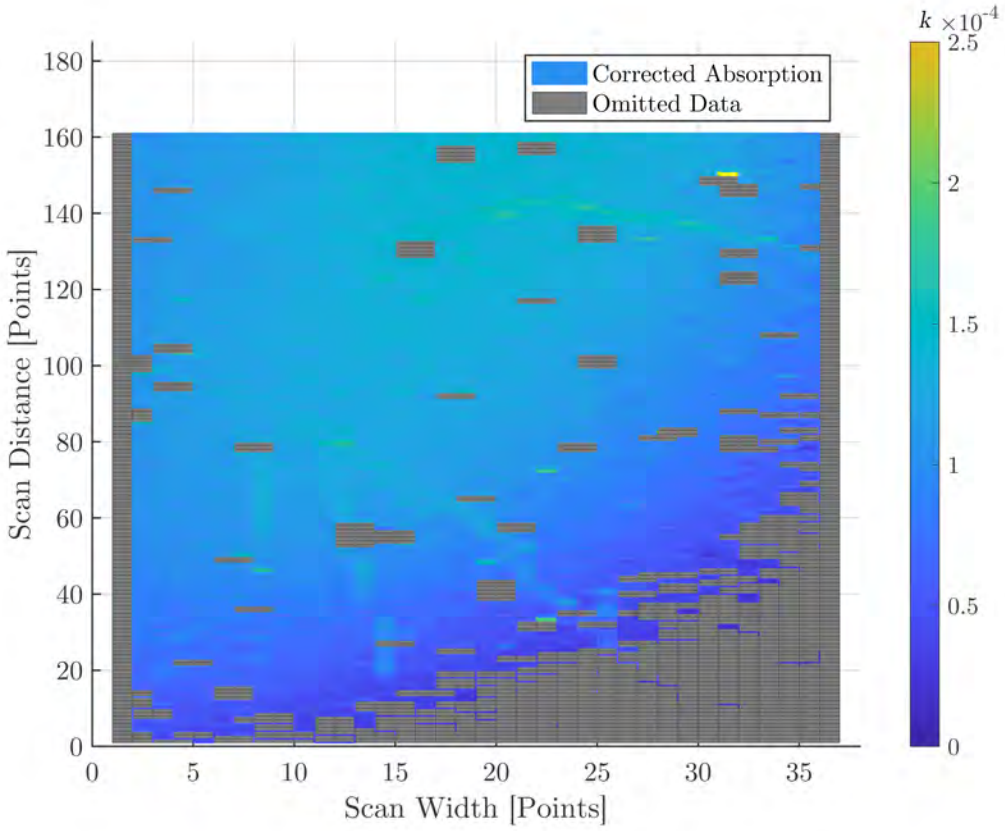

Figure 4.39: Absorption map at 1550 nm of IBS $LaTiO_3$ deposited on JGS-1. Measurements where the probe power or measured phase decreases below tolerable levels are omitted. The mapped area is equivalent to 9 mm × 8 mm.

The second region is then a region of material with a large difference in coating thickness producing a large number of lower absorbing points at $k = 1.8 \times 10^{-5}$ where the coating thickness is changing as it moves further from the centre of the plasma plume. If the same methodology is followed as with the previous wavelength, the sample's potentially thinner region shows an average absorption of around 100 ppm. At this wavelength, some differences in the coating surface, are more defined than when measured at 1064 nm.

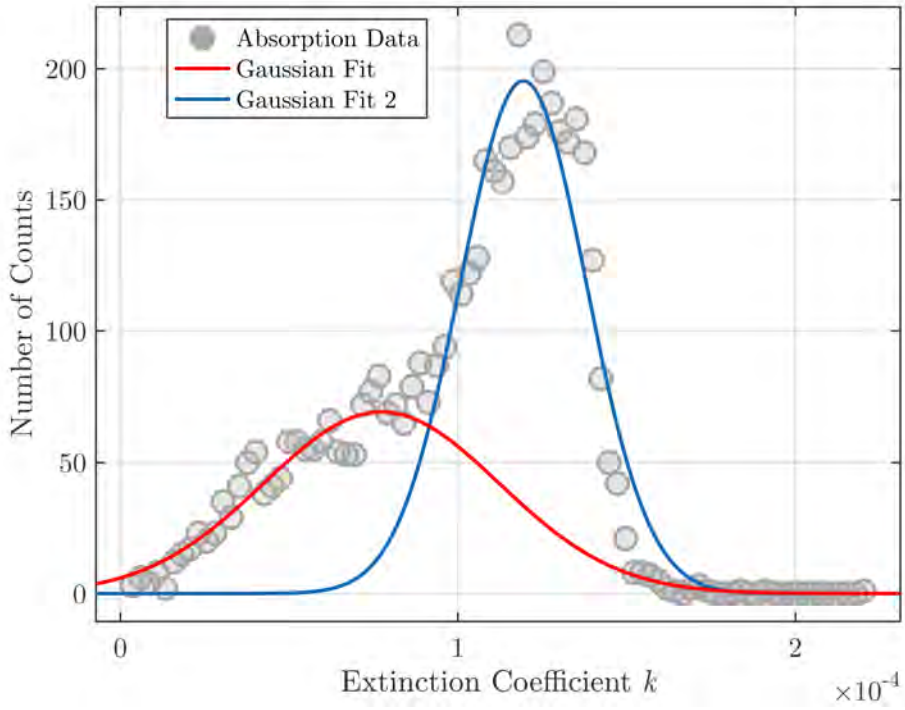

Figure 4.40: Histogram and Gaussian fits at 1550 nm of the absorption data shown in Fig 4.39.

4.3.1.4 Summary

$LaTiO_3$ deposited by IBS at Strathclyde University has been measured to have large absorption variations across the surface of the sample which are thought to be caused by an inhomogeneous coating thickness. Taking the coating portions assumed to be thinner from each measurement show that at 1064 nm and 1550 nm the coating absorption is on average around 100 ppm at both wavelengths. As other high refractive index coating materials such as Ta_2O_5 have similar magnitudes of absorption in the as-deposited state, it is believed that the absorption could be reduced by heat treating of the coating layer. However, these measurements are outwith the scope of this thesis and will be carried out at a later date.

4.3.2 Zr:Ta_2O_5

Modelling of the molecular structures of Ta_2O_5 layers by Prasai [152] and others [196] suggests that its mechanical loss can be reduced by doping Ta_2O_5 with materials such as titania (TiO_3) [145] and have since been used in aLIGO ETM

coatings. While this doping reduces the mechanical loss of Ta_2O_5, the Brownian thermal noise produced by the layers of Ti:Ta_2O_5[7] still limits GWD sensitivity between 50 Hz and 150 Hz [198]. To reduce the level of coating Brownian thermal noise produced by these layers, other dopants for Ta_2O_5 such as Zirconia (ZrO_2) are being investigated. However, unless the doping process also produces a coating with low optical absorption, it cannot be easily implemented in a ETM HR stack without decreasing detector stability.

Previous studies of Zr:Ta_2O_5 coatings produced by techniques such as IBS or Magnetron sputtering (MS) show that these materials can crystallise when heated to temperatures between 600 °C and 800 °C [137, 199], with the lowest levels of mechanical loss found between 300 °C and 600 °C [137]. Recent developments in coating deposition techniques [143, 144], have also shown that by increasing the temperature of a substrate during coating deposition, commonly known as 'elevated temperature deposition', increases the mobility of atoms in the coating layer. Where higher atom mobility can lead to an 'ideal glass' state [143] with a lower density of two-level tunnelling states, connected with a reduction in mechanical loss.

In this section a joint study which investigates the effects of Zr doping concentration, post-deposition heat treatment, and other deposition parameters on IBS layers of Zr:Ta_2O_5 layers carried out with Svetoslava Angelova, (University of Strathclyde) and the author will be detailed.

Each of the coating layers were deposited at 24 °C (room temperature), 200 °C and 400 °C to avoid crystallisation, and allow the effects of post-deposition heat treatment of each coating to be observed. In each deposition run, the sputtering angle remained constant to within ±10°, and the travel distance between the ion source and the sample holder remained fixed. The chamber pressure was 8.5×10^{-5} mbar, increasing to $(4.0\pm0.3)\times10^{-4}$ mbar after a small volume of Ar gas was bled into the chamber during deposition.

The goal of these experiments was to analyse if doping Ta_2O_5 with ZrO_2 can produce coatings with lower mechanical and optical losses than current Ti:Ta_2O_5 layers, when subjected to elevated temperature deposition and post-deposition heat

[7] Doping concentration of TiO_3 to Ta_2O_5 is approximately 18% [197].

treatments. The optical absorption of these samples was measured by the author, and is presented in the following sections.

The samples studied are detailed in Table 4.6. The total deposition time for each coating was between 24 and 25 hrs, with SiO_2 cantilever and JGS-1 absorption samples both coated in each run.

Sample ID	Short ID	Deposition Time	Deposition Temp [°C]	Zr Concentration [%]	$I_{\text{extractor}}$ [mA]	I_{focus} [mA]
I17H2209	'A'	24hrs	24	0.36	0.640	0.185
I17H3008	'B'	24hrs	24	0.36	0.450	0.237
I17J2609	'C'	25hrs	200	0.21	0.667	0.236
I17L0911	'D'	24hrs 50min	400	0.21	0.738	0.101
I17L0809	'E'	24hrs	200	0.21	0.686	0.103
I18A1714	'F'	24hrs	24	0.36	0.493	0.043

Table 4.6: Summary of deposition parameters of interest for Zr:Ta_2O_5 samples.

4.3.2.1 Absorption Measurements

Figures 4.41 and 4.42 show the measured optical absorption of all samples listed in Table 4.6 at 1064 nm and 1550 nm. The absorption of the samples deposited at room temperature ranges from 12 ppm to 43 ppm at 1064 nm. It is thought that the large variations in absorption originate from relative inhomogeneities in the thickness of the coatings as observed on the samples discussed in section 4.3.1. Sample F, deposited at room temperature, was heat-treated for 5 hours at 600 °C. Interestingly the magnitude of the absorption did not change within experimental error, but a reduction in spread by 78% was observed. For each of the samples deposited at room temperature, the Zr doping concentration was estimated to be ∼36%.

The spread in the absorption of the samples deposited at 200 °C and 400 °C is similar to that of the sample deposited at room temperature. Measurements of samples deposited at 200 °C show an average increase in optical absorption compared to those deposited at lower temperatures. Sample C was heat-treated for 5 hrs at 800 °C. The absorption increased from (27 ± 11) ppm to (37 ± 5) ppm. The spread in absorption decreased by ∼55%. The samples deposited at this

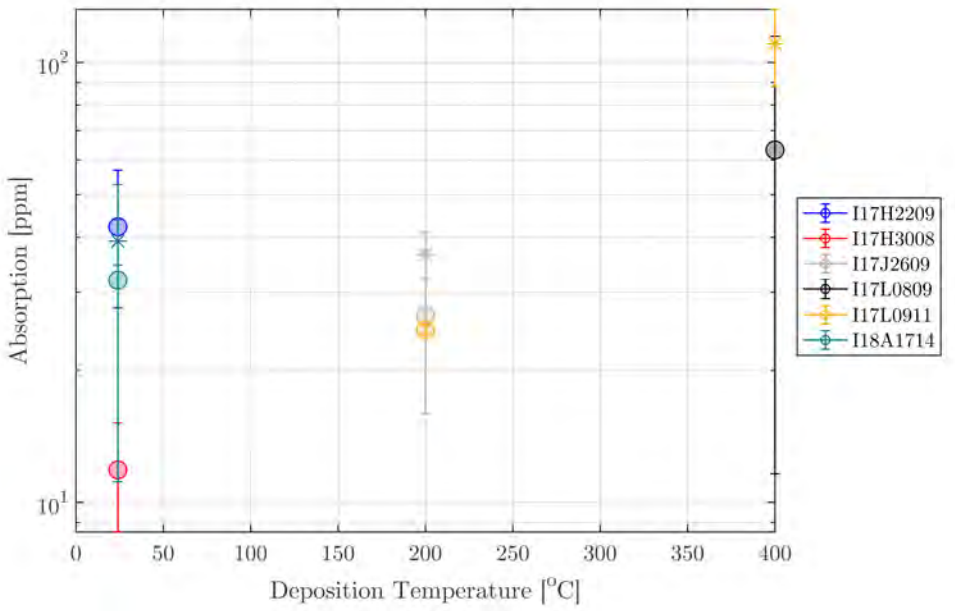

Figure 4.41: Measured optical absorption of IBS Zr:Ta_2O_5 samples as a function of deposition temperature at 1064 nm.Samples in their as-deposited states are marked by '∘', post deposition heat treatments are marked with '*' of the same colour.

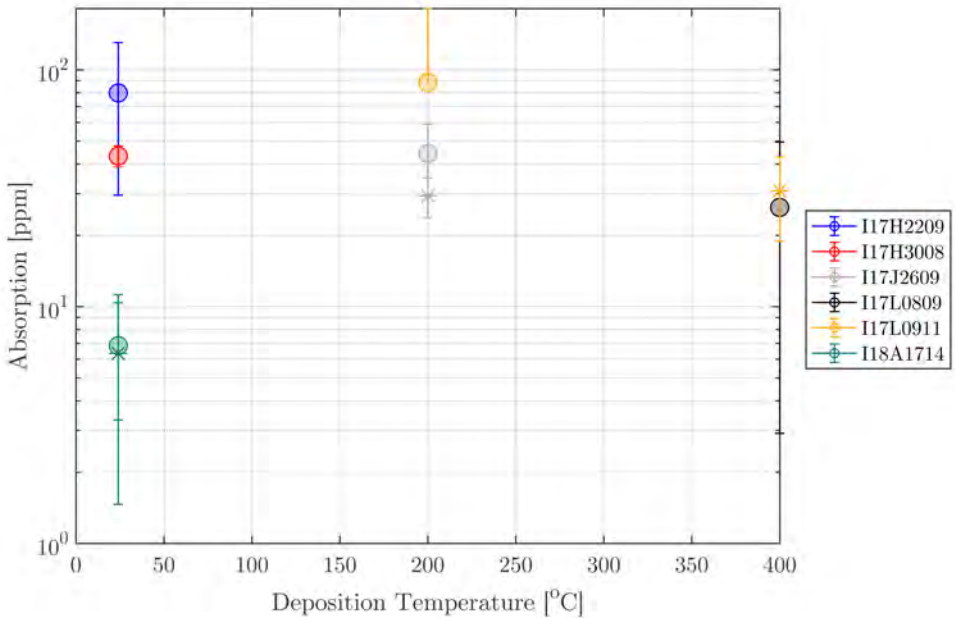

Figure 4.42: Measured optical absorption of IBS Zr:Ta_2O_5 samples as a function of deposition temperature at 1550 nm. Samples in their as-deposited states are marked by '∘', post deposition heat treatments are marked with '*' of the same colour.

temperature had a lower average Zr doping concentration of ∼20%.

Sample D, the only sample deposited at 400 °C, showed an absorption of (63.2 ± 51.6) ppm. After heat treatment at 750 °C for 5 hrs, showing an increase in average absorption with a significantly reduced spread. At 1064 nm, Sample B shows the lowest absorption of (12 ± 3) ppm. At 1550 nm, Sample F was found to have the lowest optical absorption of (97 ± 4) ppm. It is interesting to note that in some cases the absorption of the coating at 1064 nm and 1550 nm changed with respect to one and other after heat treatment. For Sample C, an increase of 28% at 1064 nm for was found after heat-treatment at 800 °C. However, at 1550 nm the absorption of Sample C reduced by 34%.

This could be due to a change in thickness or refractive index of the sample after heat treatment, changing the electric field intensity in the coating layer. It could be produced by a chemical change in the material that favours lower energy optical transitions [128]. However, without sufficient measurements of coating thickness variation across the sample, this change cannot be experimentally verified.

4.3.2.2 Correlations Between Absorption and Deposition Parameters

To gain a greater insight into the absorption differences of these coatings, the coating deposition parameters beyond those investigated above were analysed. Two variables, the currents used to control the beam 'focus' and the acceleration of the ion beam onto the target material, showed a large enough variation to perform statistical analysis concerning coating absorption. I_{ext}, is the current flowing through the ion extraction system (see Fig 4.35), accelerating the ionised plasma towards the material target. Higher values of I_{ext} is correlated with increasing the deposition rate[8]. This value can be affected by the number of secondary electrons produced in the sputtering process, sputtered from the material target and back-scattered onto the extraction cavity. Therefore deviations on the order ± 0.1 mA may not necessarily be correlated to a change in deposition rate but may be due to this effect. The number of secondary electrons in the system is directly measured by I_{foc}, which is the current flowing between the electrostatic focus and

[8] I_{ext} can vary between target materials but in general higher I_{ext} tends to leads to a higher deposition rate.

ground plates, after the ECR plasma source. A higher current on these plates can be used to deviate the ion beam's trajectory but is far more susceptible to impingement by secondary electrons in the system, and is therefore used as a measure of system stability during coating deposition. These values were recorded at the beginning of each run and considered nominally constant throughout coating deposition.

Figures 4.43 and 4.44 show the same measured absorption values for each sample, as a function of extractor current I_{text} and focusing current I_{foc}. In each case, the colour coding remains consistent with other Figures shown in this section.

This data shows a correlation between absorption and I_{ext}, has a weaker dependence on I_{foc}. Fig 4.43 shows that for higher deposition rates, coatings deposited at higher temperatures have higher absorption and larger spread compared to samples produced at lower deposition rates. For coatings deposited at room temperature, this increase is also apparent. Heat treating samples to 800 °C for 5 hrs, after elevated temperature deposition at 200 °C shows a reduction in absorption at both wavelengths. Deposition at 400 °C and heat-treating the sample at 750 °C for 5 hrs, showed a marginal increase at 1064 nm and a large increase at 1550 nm.

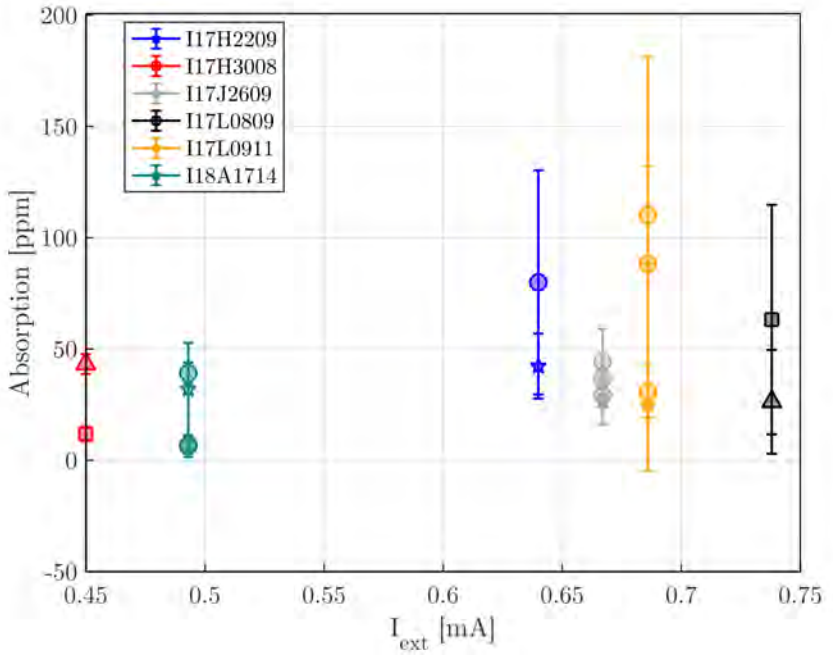

Figure 4.43: Measured optical absorption at 1064 nm of IBS Zr:Ta_2O_5 samples as a function of I_{ext}, ion extractor current. ◯ = 1064 nm As Deposited, = 1064 nm Heat Treated, □ = 1550 nm As Deposited, △ = 1550 nm Heat Treated.

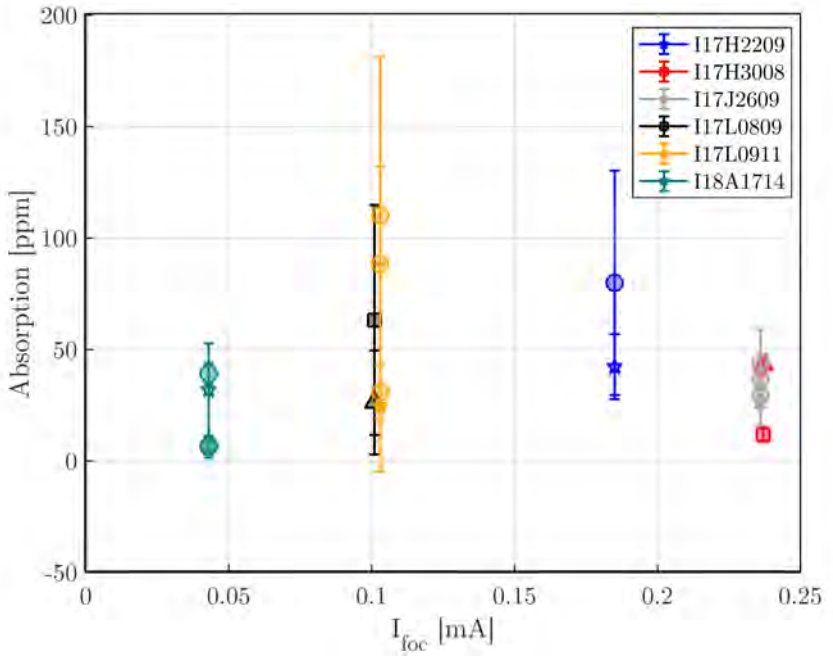

Figure 4.44: Measured optical absorption of IBS Zr:Ta_2O_5 samples as a function of I_{foc}, ion-beam focus. ◯ = 1064 nm As Deposited, = 1064 nm Heat Treated, □ = 1550 nm As Deposited, △ = 1550 nm Heat Treated.

Coatings deposited at room temperature, with low rates of deposition exhibit the lowest optical absorption of all samples, with very different absorption characteristics at 1064 nm and 1550 nm. Sample F had worse absorption than Sample B at 1064 nm (32 ± 21 ppm compared to 12 ± 3 ppm). At 1550 nm the story is reversed with Sample F having 7 ± 4 ppm at 1550 nm and Sample B with 43 ± 4 ppm respectively. This analysis shows that coatings deposited at room temperature with values of I_{foc} at 0.23 mA and 0.043 mA, the extremes of the measured values, resulted in the lowest absorption. Values between these two extremes resulted in higher absorption at both wavelengths irrespective to deposition or heat treatment temperature.

The difference in optical absorption at 1064 nm and 1550 nm in Figures 4.43 and 4.44 reveals a distinct change in the absorption characteristics for sample I17H3008 (shown in red) moving from one of the lowest absorbing samples at 1064 nm, to one of the highest absorbing samples at 1550 nm (see Figures 4.41 and 4.42). Correlating this change with the lower value of I_{foc} <0.1 mA shows that the system was much more stable during coating deposition. It can be speculated that this stability produces a coating which is more susceptible to optical transitions at 1550 nm, absorbing comparably less light at 1064 nm. All coatings produced above this threshold, regardless of deposition temperature and heat treatment, adhere to this boundary, absorbing less light at 1064 nm.

The results from this study show that coating deposition values that are not customarily disclosed from commercial or institutional vendors can directly influence the optical properties of the final deposited film. With small deviations in current and system stability of the coating system producing coatings with different absorption values when measured with PCI at 1064 nm and 1550 nm. These results are added under the caveat that each coating's thickness can affect the total electric field intensity. Therefore, its absorption could affect the conclusions of these results and should be experimentally verified at a later date. This would allow simulations of electric field intensity inside the coating layers to be produced, allowing variations in coating thickness between different samples removed.

4.4 Conclusions

The stable operation of a gravitational wave detector relies on highly reflecting mirror coatings. By measuring the optical absorption of these materials as single layer materials using PCI and with knowledge of the attenuation of a propagating light field through the material, the extinction coefficient, k, of the material can be calculated. Stabilising polarisation fluctuations in the high power pump beam, variations in laser power and the resultant error this produces is reduced. Increasing the stability of the PCI system allowed for 2 dimensional mapping techniques to be developed, where through continuous monitoring of power fluctuations, variations in power dependent absorption can be removed. Using this technique the optical absorption of coating materials deposited by RLVIP and IBS was measured, where investigations into deposition parameters and the effects of post deposition heat treatment were analysed.

Through altering the surrounding atmosphere in which single layers of Ta_2O_5 were heat treated, the optical absorption of the material could be altered. It was observed that samples heat treated in atmosphere decreased in absorption, and the absorption samples heat treated in vacuum increased. By varying the heat treatment duration a trend between optical absorption, and heat treatment duration was found in both sets of Ta_2O_5 samples. From these measurements, it follows that the optical absorption of RVLIP Ta_2O_5, which is deposited in an 'oxygen-poor' state is strongly influenced by the conditions in which it is heat treated. By leaving the samples in a laboratory atmosphere for over a year, measurable changes in the optical absorption of the samples were observed suggesting that the initial changes in absorption after heat treatment were being undone over these timescales. It has been suggested that the changes in absorption could be correlated to a change in oxygen content of the materials or could be influenced by the adsorption of molecular water over long timescales. However in order to confirm these results, further experimentation is needed.

Optical absorption measurements of IBS coatings produced from the University of Strathclyde have shown that coating deposition values which are not customarily disclosed from commercial or institutional vendors can directly influence the optical properties of the final deposited film. These results suggest that values which

dictate the deposition rate and stability of the coating chamber during coating deposition can have different influences on the final absorption of the material at 1064 nm and 1550 nm.

To reduce the effects of Brownian thermal noise in a gravitational wave detector, a novel coating design has been developed which allows the high optical absorption of aSi layers to be greatly reduced in a HR coating stack. It has been shown that a multi-material coating stack consisting of Ta_2O_5-SiO_2/SiO_2-aSi layers can be used to reduce optical absorption of aSi layers with the *Full Stack* coating absorbing 8.1 ppm, 95%, lower than the absorption of the *Lower Stack* coating. These materials could be used to reduce the mechanical loss and thickness of an end test mass coating, having a meaningful effect on the level of coating Brownian noise and overall sensitivity of the detector. This result coupled with the higher refractive index of 3.73 and lower mechanical loss of aSi would allow the >99.999% reflectivity requirements of an ETM mirror coating to be achieved with 20 layers of coating material instead of 38 bi-layers of SiO_2-Ta_2O_5 which are currently in operation in aLIGO detectors, if the absorption can be further reduced.

Chapter 5

Room Temperature Mechanical Loss Measurements Using a Gentle Nodal Support

As discussed at the end of Chapter 2, sensitivity of current gravitational wave detectors is limited by thermal noise produced by the mirror coatings applied to the test-mass optics between 50 Hz - 150 Hz. The mechanical dissipation of a coating layer ϕ, and thickness of a coating material directly contributes to the level of coating thermal noise, with optical properties such as refractive index indirectly contributing as this influences the material thickness required.

The development of coatings with lower thermal noise is critical for the next generation of detectors and detector upgrades, and studies of coating mechanical loss at room temperature [198, 200] and at cryogenic temperatures [78, 104] are a key part of this work.

This chapter details the techniques used to study the mechanical dissipation of coating materials at room temperature for current gravitational wave detector upgrades such as LIGO A+ [198] and the planned Einstein Telescope- High Frequency detector (ET-HF) [201].

5.1 Disk Resonators

Cylindrical samples are widely produced to high optical standards from high grades of silica have shown very low mechanical losses [176,177] when suspended from points of minimum resonant motion 'nodes' to reduce the interactions between the mechanical loss sample and its suspension method [176]. The resonant frequencies of a cylindrical body were calculated by Chree [160] in 1886. A simplified version of this treatment, allows the resonant frequency of rotationally-symmetric modes ω_{sym} to be calculated using the Youngs modulus E, the density ρ, and the Poisson ratio ν of the geometry

$$\omega_{\mathrm{sym}} = \beta_{\mathrm{sym}} \sqrt{\frac{E}{\rho(1+\nu)}}, \tag{5.1}$$

where β_{sym} is a spatial parameter related to the dimensions of the disk [160]. Chree stated that the displacement of the cylinder could be approximated by sinusoidal standing waves, whose order is related to the magnitude of the frequency, passing through the geometry. It can be shown that the frequencies of non-radially symmetric modes of a free cylinder can be calculated in terms of two bounded Bessel functions I_n and J_n [202]

$$\omega_n = \frac{\lambda^3 I_n'(\lambda) + (1-\nu)n^2\left[\lambda J_n'(\lambda) - J_n(\lambda)\right]}{\lambda^3 I_n'(\lambda) - (1-\nu)n^2\left[\lambda I_n'(\lambda) - I_n(\lambda)\right]} \tag{5.2}$$

where the roots or eigenfrequencies of Eq. 5.2 λ can be calculated in terms of the mode number and the radius of the disk:

$$\lambda = r - \frac{m+1}{8r} - \frac{4\left(7m^2+22m+11\right)}{3(8r)^3} - \cdots \tag{5.3}$$

where $(m = 4\,n^2)$ and $(r = (\frac{\pi}{2}(n+2\,s))$ [202]. As these analytical calculations quickly become complicated and time consuming, finite element analysis (FEA) was used. Using FEA packages such as COMSOL or Ansys, a 3-dimensional representation of a cylinder can be modelled, in order to calculate disk mode frequencies at which its modes will occur.

5.1.1 Mechanical Resonances and Energy Dissipation

For each resonant frequency of a geometry, its surface deformation differs following the amount of energy stored in two different components of the samples motion in a

Cartesian coordinate system (i, j, k):-

The dissipated energy attributed to bulk motion (E_{bulk}), shown in Eq. 5.4. This is expressed in the form of an integral with respect to the volume of the geometry:

$$E_{\text{bulk}} = \int \frac{1}{2} K \theta^2 dV. \tag{5.4}$$

The magnitude of the resultant deformation is defined by the materials bulk modulus K, scaled by the first component of a 3×3 strain matrix S_{ii} denoted by θ^2 for that material.

The second component details the energy dissipated through shear motion (E_{shear}), a pure shape deformation that can also be calculated similarly, with the movement dictated by the shear modulus μ multiplied by two components of the 3×3 matrix S

$$E_{\text{shear}} = \int \mu \xi_{ij} \xi_{ij} dV. \tag{5.5}$$

Terms denoted by ξ_{ij} are in reference to the respective components of the strain matrix S_{ijk} scaled by a Kroneker delta function, g_{ij} such that

$$\xi_{i,j} = \frac{1}{2} \left(S_{ij} + S_{ji} \right) - \frac{1}{3} g_{ij} S_{kk}. \tag{5.6}$$

It was also shown in works by Hong et al. [132] that the strain matrix S_{ij} can be represented in terms of displacement in each axis:

$$S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right). \tag{5.7}$$

For each resonant mode, the total energy stored/ dissipated by the geometry during deformation is equal to the sum of these two components:

$$E_{\text{T}} = E_{\text{bulk}} + E_{\text{shear}}. \tag{5.8}$$

Deformations due to resonating motion are a combination of bulk and shear motion. Depending on the distribution of deformation of an on-resonance sample, the amount of energy stored in E_{shear} and E_{bulk} differs. For gravitational wave detector mirror coating applications, materials which have a higher component of

dissipation through shear motion (perpendicular to the sensing laser) will result in a lower contribution to coating Brownian thermal noise and is thus more favourable (see chapter 2).

Fig 5.1 shows examples of the surface deformations of different modes of a thin cylinder of thickness $t < r$, calculated from FEA with COMSOL. It is common to characterise the mode shape of each resonance based on the number of nodal lines m and the number of nodal circles n in the notation (m,n) [166] where

- a nodal line: line of symmetry across the flat surface of the cylinder and
- a nodal circle: An unbroken radial 'ring' of little/no deformation across the flat surface of the cylinder.

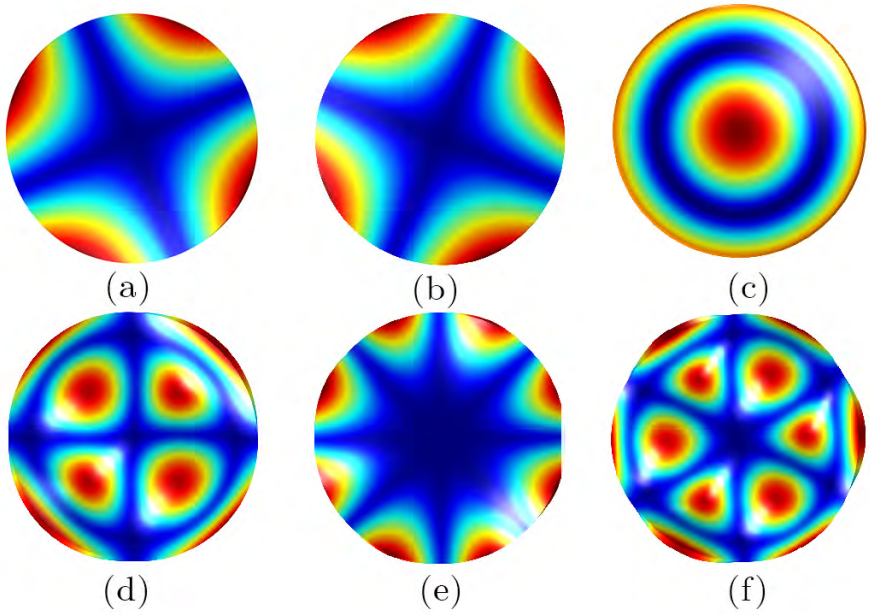

Figure 5.1: Surface deformations of mode-shapes of a cylinder calculated with finite element modelling (COMSOL) listed in ascending frequency. Regions of large displacement are denoted in red, contrasting regions with little/no displacement, are shown in blue.

Of the six surface deformations plotted in Fig 5.1 each mode-shape then takes on the following notation:**(a)**: (2,0), **(b)**: (2,0), **(c)**: (0,1), **(d)**: (2,1), **(e)**: (4,0) and **(f)**: (3,1). In the following sections, it will be shown that the ability to measure the mechanical loss of a given mode shape depends on the experimental apparatus.

For samples composed of isotropic materials such as SiO_2, there is only one, value for the Young's Modulus density and Poisson ratio. This isotropic value results in pairs of modes which are indistinguishable in frequency and relative deformation (i.e. Fig 5.1**(a)** and 5.1**(b)**). The relative orientation of these mode-shapes occurs at 45° rotated with respect to one and other, about a geometry centred coordinate axis. FEA predicts these modes to manifest at nominally the same frequency, but due to geometry imperfections and different measurement apparatus each frequency is normally separated by $\leq$1 Hz (see section 5.1.3).

For anisotropic materials such as cSi, there are different sets of elastic properties, which change depending in which crystal axis, the deformation occurs. This, significantly changes the frequency at which the resonance will occur without changing the overall shape of its deformation. The percentage of strain energy stored in E_{shear} and E_{bulk} for each mode shape allows each deformation to be classified into three different 'mode families' as shown in Fig 5.2.

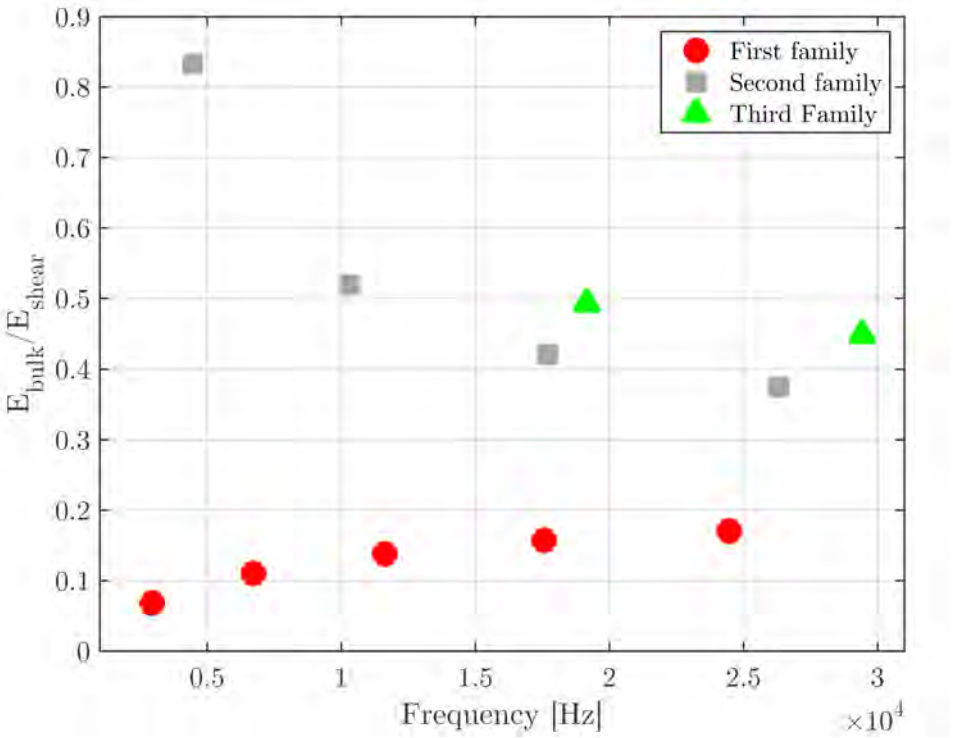

Figure 5.2: Calculated ratio of $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ for an uncoated $\varnothing = 3$",$t = 2.7$ mm silica disk as a function of frequency.

First Family modes ($m \geq 2,0$) form nodes and anti-nodes of deformation around the circumference of the geometry and have $E_{\text{shear}} \gg E_{\text{bulk}}$. Second Family modes have a larger component of volumetric deformation, i.e. $E_{\text{shear}} > E_{\text{bulk}}$ and their deformations form nodes and anti-nodes across circular face of the disk ($m \geq 0,1$) these are colloquially called, referenced as 'drum' modes. 'Third family' modes ($m \geq 1,2$) have a slightly higher distribution of bulk and shear energies compared to Second Famlily modes. The respective frequencies and values of $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ for a selection of mechanical modes from Fig 5.2 are given in Table 5.1.

To accurately calculate mechanical loss $\phi(f)$ of a given mode-shape, this must be taken into consideration. As the mechanical loss of a geometry is defined as the energy dissipated over a given time, if the energy lost during excitation is dependent on two different dissipation mechanisms, this suggests that there are intrinsically different associated values of ϕ depending on the deformation of the sample. As current experimental methods can only measure the total mechanical loss of the coating for a given mode, a combination of experimental results and FEA is used to separate the total loss into these components.

It can be shown that using the method derived by Hong [132], $\phi(f)$ can be calculated using a variant of the coating loss equation which takes into account the dissipated energy into components of $\mathrm{E}(f)_{\text{bulk}}$ and $\mathrm{E}(f)_{\text{shear}}$

$$\phi(f)^{BS}_{\text{substrate}} = \left(\frac{E(f)_{\text{bulk}}}{E(f)_{\text{total}}}\right)\phi_{\text{bulk}} + \left(\frac{E(f)_{\text{shear}}}{E(f)_{\text{total}}}\right)\phi_{\text{shear}}, \tag{5.9}$$

where $\phi(f)^{BS}_{\text{substrate}}$ is the loss accounting for components ϕ_{bulk} and ϕ_{shear} of each dissipation mechanism. This statement can be used to calculate the loss of ϕ_{bulk} and ϕ_{shear} for the uncoated substrate, and for the loss of the coating, after deposition. G.Cagnoli et al. [135] found that for a cSi disk, the ratio of $E_{\text{bulk}}/E_{\text{Total}}$, also known as the dilution factor $\mathbf{D}_{\text{dil}}$ are dependent on the samples material properties with the Poisson ratio of silicon, contributing the largest change in $E_{\text{bulk}}/E_{\text{Total}}$ in their work

$$\mathbf{D}_{\text{dil}} = \frac{E_{\text{bulk}}}{E_{\text{Total}}}. \tag{5.10}$$

Frequency	$\frac{E_{\text{shear}}}{E_{\text{Total}}}$	$\frac{E_{\text{bulk}}}{E_{\text{Total}}}$	Mode Number (m,n)	$\frac{E_{\text{bulk}}}{E_{\text{shear}}}$
2950.1	93.2%	6.8%	(2,0)	0.068
2950.1	93.2%	6.8%	(2,0)	0.068
4475.7	17%	83%	(0,1)	0.83
6727.9	89%	11%	(3,0)	0.11
6727.9	89%	11%	(3,0)	0.11
10318.5	48%	52%	(1,1)	0.52
10318.5	48%	52%	(1,1)	0.52
11625.7	86%	14%	(4,0)	0.14
11626.3	86%	14%	(4,0)	0.14
17565.9	84%	16%	(5,0)	0.16
17565.9	84%	16%	(5,0)	0.16
17696.2	58%	42%	(2,1)	0.42
17697.4	58%	42%	(2,1)	0.42
19163.9	51%	49%	(0,2)	0.49
24469.9	83%	17%	(6,0)	0.17
24470.6	83%	17%	(6,0)	0.17
26314.3	62%	38%	(3,1)	0.38
26314.3	62%	38%	(3,1)	0.38
29420	55%	45%	(1,2)	0.45
29420	55%	45%	(1,2)	0.45

Table 5.1: Comparison of bulk and shear and total elastic strain energy densities for resonant frequencies of a 3"×0.1" (76.24 mm×2.69 mm) SiO_2 disk calculated with COMSOL.

This was confirmed by the author, running similar simulations in COMSOL to confirm how the distribution of energy would change with increasing mode frequency and shape. The standard form of coating loss calculation to take the following form:-

$$\phi(f)^{BS}_{\text{coating}} = \frac{E_{\text{coated}}}{E_{\text{substrate}}}\left(\phi(f)^{BS}_{\text{coated}} - \phi(f)^{BS}_{\text{substrate}}\right) \tag{5.11}$$

with the contributions of bulk and shear energies from the uncoated, and coated substrates taken into consideration. E_{coating} and $E_{\text{substrate}}$ are the calculated strain energies stored in the coating and substrate respectively. The total energy stored in each component of a mode can be calculated using equations 5.4 and 5.5 through FEA modelling. However, the values for ϕ_{shear} and ϕ_{bulk} cannot be calculated in this manner. As these values are 'entangled' in the measured values of ϕ_{mech}, they can be estimated through regression analysis. Work carried out by M. Fletcher [203] showed that with measured values $\phi(f)_{\text{mech}}$ and FEA simulated values for $E(f)_{\text{bulk}}$ and $E(f)_{\text{shear}}$, values of ϕ_{shear} and ϕ_{bulk} could be estimated using linear regression analysis. $E_{\text{bulk}}/E_{\text{shear}}$ values for thin cylinders ($t \ll \varnothing$) are invariant for geometries with the same material properties [135]. Only when the thickness of the geometry becomes comparable to the samples diameter do these values begin to change. This produces a stable model which can be used to classify the shape of a mode from FEA in further analyses.

5.1.1.1 Investigation into Anisotropy and $E_{\text{bulk}}/E_{\text{shear}}$

To better understand the how the material properties of a system affect the distribution of $E_{\text{bulk}}/E_{\text{shear}}$, this study is carried out for cylindrical substrates, and cantilever substrates which have previously been used to characterise the mechanical loss of thin film coatings [83, 93, 204, 205]. Unlike disk substrates, a cantilever sample is composed of a thin rectangular 'ribbon' ($t_{\text{ribbon}} \approx 60\,\mu\text{m}$) with a thicker region ($t_{\text{clamp}} \approx 0.5\,\text{mm}$) which is then firmly clamped between two stainless blocks. An example of a cSi cantilever and its relative dimensions is shown in Fig 5.3.

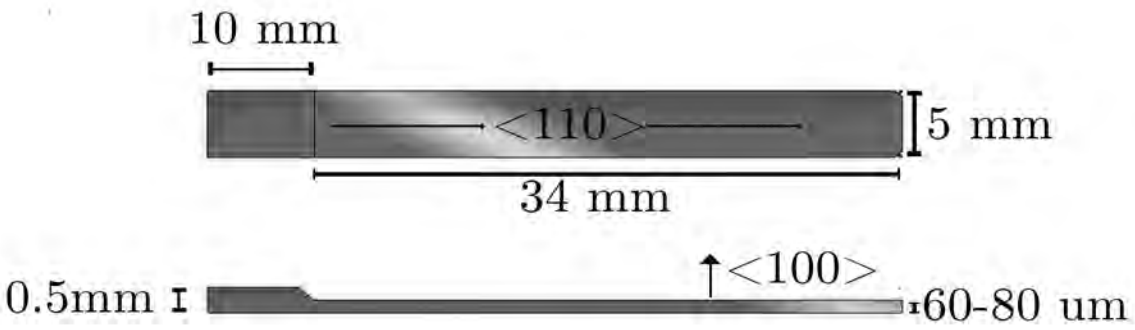

Figure 5.3: Representation of a crystalline silicon cantilever noting its relative dimensions and orientation of its crystal axes.

For each sample case, the properties of the material each sample was composed

of was varied between an isotropic material with the same properties of silicon, and an anisotropic material. This allows any changes in $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ for each material and geometry to be calculated.

Figures 5.4 and 5.5 highlight key differences in how each sample is affected by the intrinsic underlying dissipation mechanisms in cantilever and disk substrates. In each case both a disk and cantilever were modelled. Fig 5.4 shows the calculated dilution factors for each mode shape of the cantilever geometry. The dilution factor torsional mode for a cantilever, which is almost purely dominated by shear energy, shows little to no change when the material is isotropic or anisotropic, increasing by only 1% over the calculated frequency range. Bending and lateral modes, however, for an isotropic Young's modulus, a much larger portion of energy is stored in bulk motion. For each type of structure the calculated values show a large frequency dependence, increasing to a maximum level of bulk energy at 1.6 kHz, where the predicted values then begin to decrease. Fig 5.5 shows the results of the same analysis carried out on a cylindrical geometry.

Performing the same analysis on a 76.2 mm $\times$ 430 µm 'silicon' disk allows a similar trend to be produced. In this case, $E_{\text{bulk}}/E_{\text{shear}}$ has been plotted against mode number, as this clearly separates degenerate mode pairs which have the same mode frequency. For disks 'butterfly modes' ('$\star$'), on isotropic substrates, again have a higher amount of energy in shear than bulk. However, the difference between the two materials in this geometry has less profound effect on its energy distribution. For Second and Third Family modes changing the structure of the disk from isotropic to anisotropic shows a $\sim$30% difference in $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$. Introducing anisotropy into the material also causes the frequencies of previously degenerate modes to shift to higher values, removing the degeneracy. Cylindrical geometries show a maximum difference in bulk and shear energies for lower frequency modes, converging on a value $\sim$20% in both an isotropic and anisotropic case. For the case of the disk substrates, it is also clear that as the frequency increases, the ratio of $E_{\text{bulk}}/E_{\text{shear}}$ appears to converge[1]. This is also apparent upon inspection of the mode shape for higher frequency modes, where the distinguishing features of each

[1]Cantilever substrates at higher frequencies do not converge on value of $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ but manifest as different mode-shapes which do not follow this trend. As this effect occurs out-with the frequency range of the measuring apparatus, and is not investigated further.

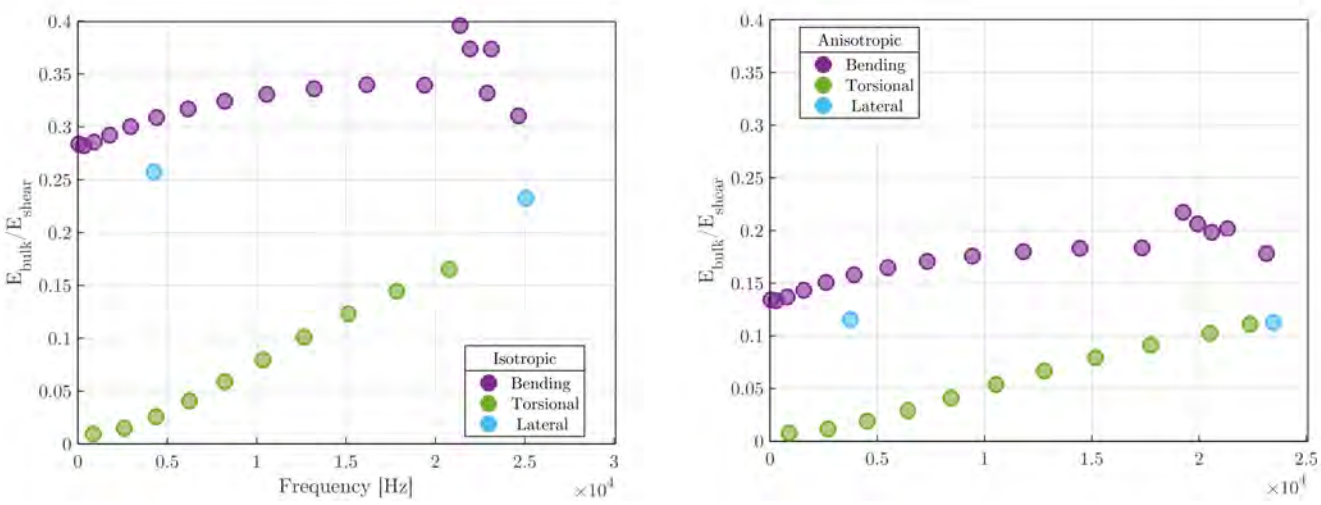

Figure 5.4: Calculated energy in bulk (E_{bulk}) and (E_{shear}) for a cantilever of the same dimensions, changing only the material structure from an isotropic (LHS) to anisotropic silicon (RHS).

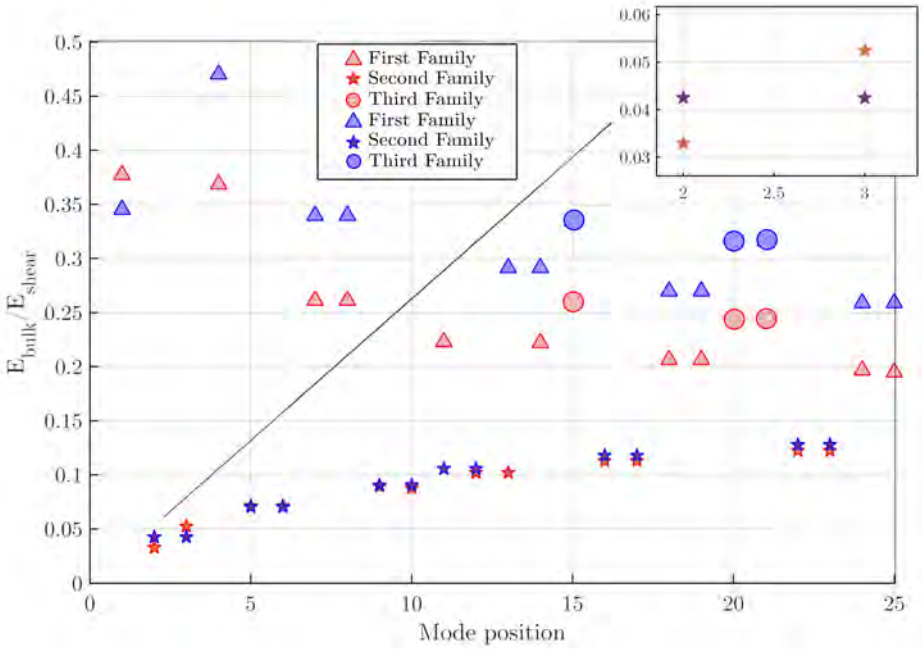

Figure 5.5: Calculated energy in bulk (E_{bulk}) and (E_{shear}) for a 3" disk of the same dimensions, changing only the material structure from an isotropic (blue) to anisotropic silicon (red).

type of motion are no longer apparent, as $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ converges.

Viewing each mode in terms of $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ allows the identification of mode-pairs on cylindrical, isotropic geometries. An example of this is shown in Fig 5.1 for both (2,0) modes in positions (b) and (c). Each mode has the same magnitude of bulk and shear energy in an isotropic material, both modelled to resonate at 656.2 Hz. In the case of an anisotropic disk, the energy ratio for the same modes changes, as do their resonant frequencies, predicted at 575.1 Hz and 672.0 Hz respectively. A change in structure has been shown to have a greater effect on modes which have larger relative deformations such as ($m \geq 2$,0) modes, as the resultant shape of the mode is then defined by Young's moduli and Poisson ratios in different directions. At higher frequencies, the relative deformation of each mode decreases, reducing this effect.

5.1.2 Measuring Mechanical Loss With Wire Suspensions

To repeatably measure low mechanical losses, any external interactions with the sample will introduce a source of damping, which can cause a higher loss to be measured. Cantilever substrates are manufactured such that the thicker clamping block reduces interactions between the resonating flexure and the clamp. However, it has been shown [83] that force applied across this region both depends on the position and angle of the sample inside the clamping block, which can reduce the repeatability of consecutive loss measurements on the same sample. Building on the work of Numata [206] and Cumming [96] it was demonstrated that very low mechanical loss could be measured from a 3" silica disk suspended by thin wires.

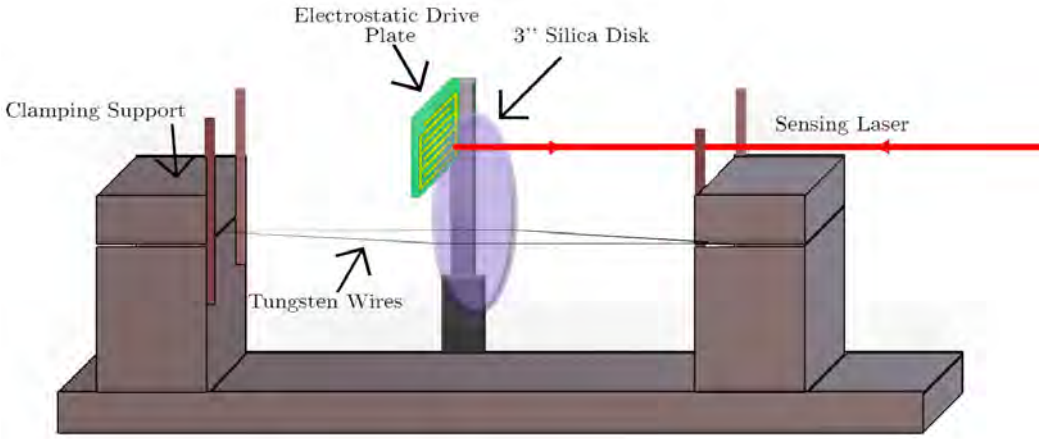

Figure 5.6: Drawing of the wire suspension method, detailing the positions of clamping blocks, exciter plate and sample.

The sample is held in place using two tensioned pieces of polished tungsten wire such that the centre of mass of the sample is on the line between the two wires ($\varnothing_{\text{wire}} = 50\,\mu\text{m}$). The tension of the wires can be altered by moving the position of the clamping blocks with respect to one and other. The suspending wires produce two points of contact at which the sample is held. Any interactions between the sample and its suspension are confined to these regions. As the deformation of the tungsten is minimal, the contact regions are assumed as two point-contacts on diagonally opposite points of the barrel of the disk [96]. This method allows for modes which are have a higher percentage of volumetric (bulk) deformation and shape-deformation (shear) modes to be sampled. The position of a sample was monitored during translation using a SIOS GmBH vibrometer [207]. This allows any motion in the probed region on the sample to be tracked as a function of time and frequency during mechanical loss measurements.

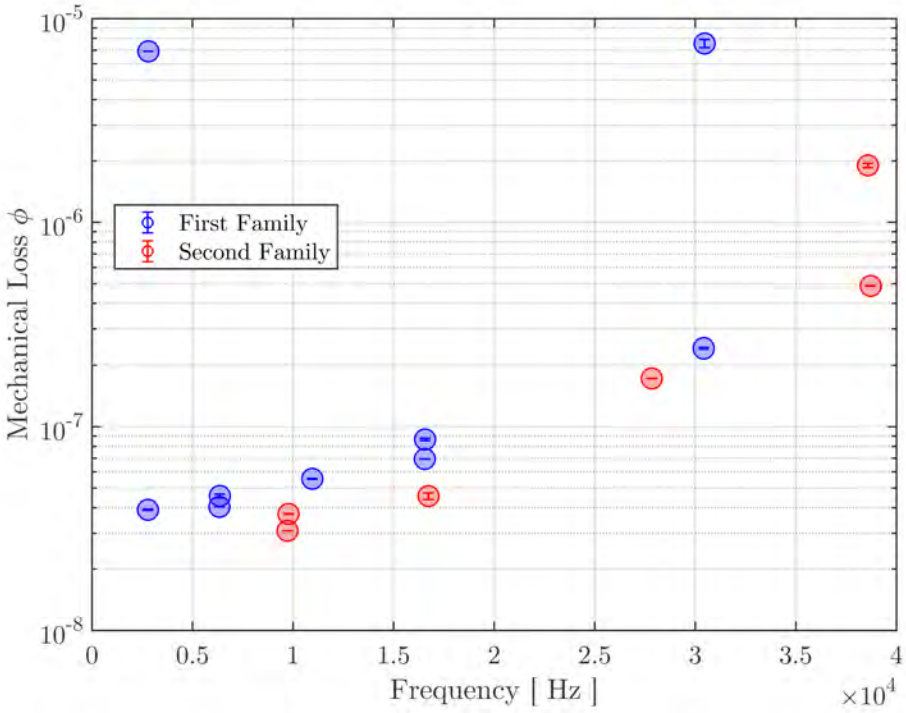

Figure 5.7: The lowest of three consecutive measurements of an uncoated 3"× ≈0.1" (76.24 mm ×2.69 mm) SiO_2 disk (Disk 5) on the 2-POC wire suspension set up. Error-bars are the standard deviation of the loss values from which the suspension on which it was measured. Mechanical modes which have interactions with the suspension are shown in blue circles, modes which do not interact are shown in red squares.

Figure 5.7 shows the mechanical loss of a 3" Corning 7980 SiO_2 disk measured in a wire suspension where 9 mode-shapes were sampled from 2.8 kHz to 38.3 kHz. The same sample was re-suspended a further two times to test the repeatability of the initial measurement. It is clear that there is a large difference in the measured loss values at similar resonant frequencies for the majority of the measured modes. However, the mechanical loss of modes measured at 4.2 kHz, 9.7 kHz and 28 kHz (identified as $(m \geq 0,1)$ modes) have a much smaller error through consecutive measurements. These modes have very little motion close to the circumference of the disk, producing very little / no interaction with the wire suspension (See Fig 5.1). In these cases, the measured ring-downs are purely exponential and can be fit using a single exponential decay. However, modes which in contrast have the majority of their motion around the circumference of the disk have a much higher probability of interacting with the wire suspension. This interaction 'couples' into the measurement, damping the mode.

Two degenerate modes which are sufficiently close in frequency can both be excited as the drive frequency is increased, and can decay simultaneously. While these modes have the same shape, their relative rotation can be such that one is more likely to be damped by interactions with the suspension, increasing the loss measured for this mode. Even though the first mode in the pair has no physical contact with the suspension, the spacing between the modes in frequency allows the now higher loss of the second mode to couple into the first mode. This coupling effect causes a 'splitting' in the measured losses for each mode pair, producing the large spread in measured losses as shown in Fig 5.7. Fig 5.8 shows the measured decay of such a mode pair. Interactions between the sample and the suspension can be seen. After 800 seconds the rate of amplitude decay changes sharply. A single exponential decay can no longer approximate this ringdown due to this coupling and care must be taken to approximate its true loss.

The fact that disks are suspended between the wires by hand can introduce inconsistencies in measured loss between repeated measurements. Changes in the wires tension and their relative position on the sample will change the loss measured and can change the measured frequency of the mode by more than $\pm$ 10 Hz.

In the case of a disk freely floating in space it would be expected that mechanical mode pairs could also periodically transfer energy between one and other, determined by their spacing in frequency with respect to one and other. This effect is discussed a single mechanical mode pair in section 5.1.4.1. This mechanical coupling between modes is also thought to occur between multiple mode pairs which could effect the predicted level of mechanical loss, however characterising this effect goes beyond the scope of this thesis and is a subject for later work.

In order to quantify the loss of a coating applied to a sample, the ϕ_{mech} of the same modes must be measured on the uncoated and coated disk. Due to the inconsistencies between consecutive suspensions of the same sample and the shift in mode position after heat treatment or coating deposition, identification of mode-shapes from measured ringdowns can be difficult.

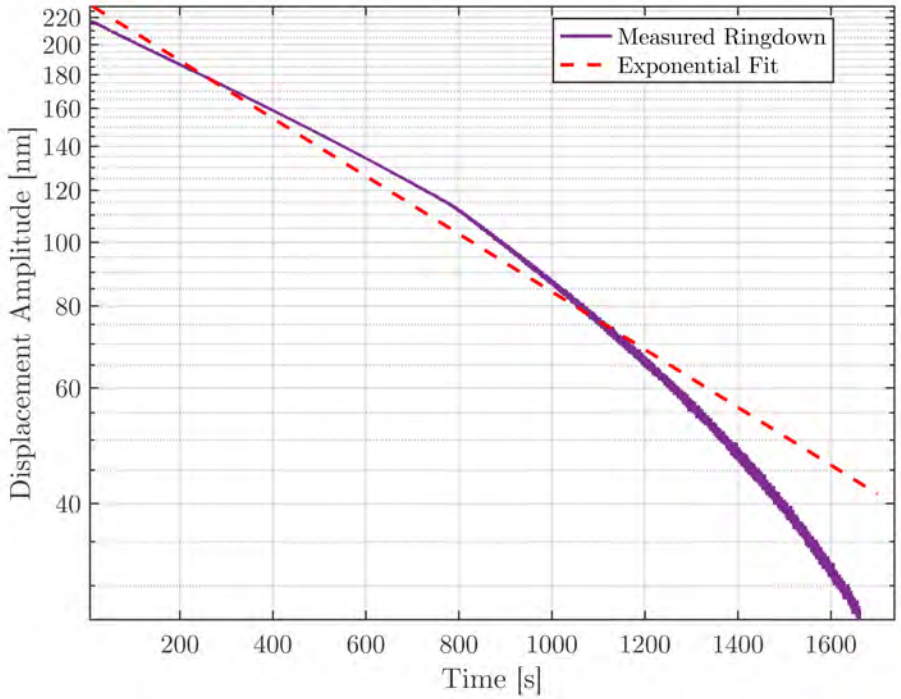

Figure 5.8: Example measured ringdown of a 2.8 kHz mode on a 3"× ≈0.1" (76.24 mm×2.69 mm) Corning 7980 SiO_2 disk (Disk 5) measured on a wire suspension. The purple line denotes the measured decay and the red line an exponential fit.

As the amount of damping of the sample from the suspending wires also affects the measured loss of mode pairs, the comparison of the same sample before and after coating can give losses lower than that of the uncoated substrate which, in turn, leads to nonphysical, negative, calculated coating losses. It has been shown by [96] that this effect can be reduced by reducing the contact area between the wire and sample. However, as the level of suspension damping is not constant between measurements, this systematic effect cannot be easily fully corrected. This makes it difficult to discern if the measured loss values are the intrinsic properties of the sample or if under different conditions, a lower loss value could be measured.

5.1.3 Gentle Nodal Support (GeNS)

A significant limitation of the wire suspension method is the 2 points of contact on the barrel of the sample, potentially damping down nodes of deformation during oscillation. To reduce these effects, the position on which the disk is contacted must be moved to a region with consistently lower amounts of displacement during

excitation, for all potential mode-shapes to be measured. E. Cesarini et al. [176] showed that a cylindrical sample of thickness t could be balanced on a spherical lens with a radius of curvature r_c and diameter D. This creates a method from which mechanical losses can be measured with a single point of contact system.

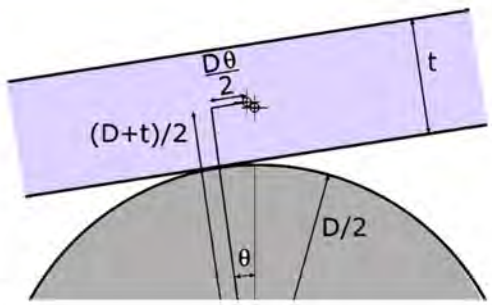

Figure 5.9: Geometric representation of a cylindrical sample of thickness t balanced on a lens of known radius of curvature ($r_c = 60.44\,\text{mm}$) and diameter D [176].

To provide such a surface to balance a 3" sample on, a 1" plano-convex silicon lens ($r_c = 60.44\,\text{mm}$) was used. The potential energy $V(\theta)$ of the sample must be nominally constant so that the sample will not overcome the coefficient of static friction μ_s and fall from the lens. If the form factor of the sample and balancing lens are known, the potential energy of the sample at different values of θ can be calculated [176].

It was shown that for values of $t < D$, the potential energy $V(\theta)$ of the disk can be considered flat when the tilt is ± 2 degrees. Thus any small deviations in the angle of the sample relative to the horizontal plane will not cause the sample to fall. The maximum tilt θ at which a sample will remain balanced is a function of the samples thickness t, diameter D

$$\theta \sim \pm\sqrt{3\frac{(D-t)}{t}}. \tag{5.12}$$

For the plano-convex lens with r_c used by the author, a $\varnothing = 3"(76.24\,\text{mm})$, $t \approx 0.1"$ $(2.7\,\text{mm})$ SiO_2 disk will be stable between $\pm\ 15\,°$. For tilts $-15° < \theta < 15°$ the potential energy of the sample increases, and if the sample tilts it will

'self-right', moving back to a balanced position. In regions greater than $\pm$ 5° the potential energy sharply decreases, indicating regions where sample balance is no longer possible.

Due to the coincidence between the point of contact between the curved lens and centre of mass of the disk, any mode shape which has a significant displacement ($m \geq 0,1$) modes in this region will experience frictional damping.

5.1.3.1 System Prototyping and Development

To test this suspension method, a prototype lens holder and base, were fashioned from acrylic plastic. To reduce the effects of excess vibrational damping on the sample, the piece could be screwed into a horizontal plate on the inside of the vacuum chamber.

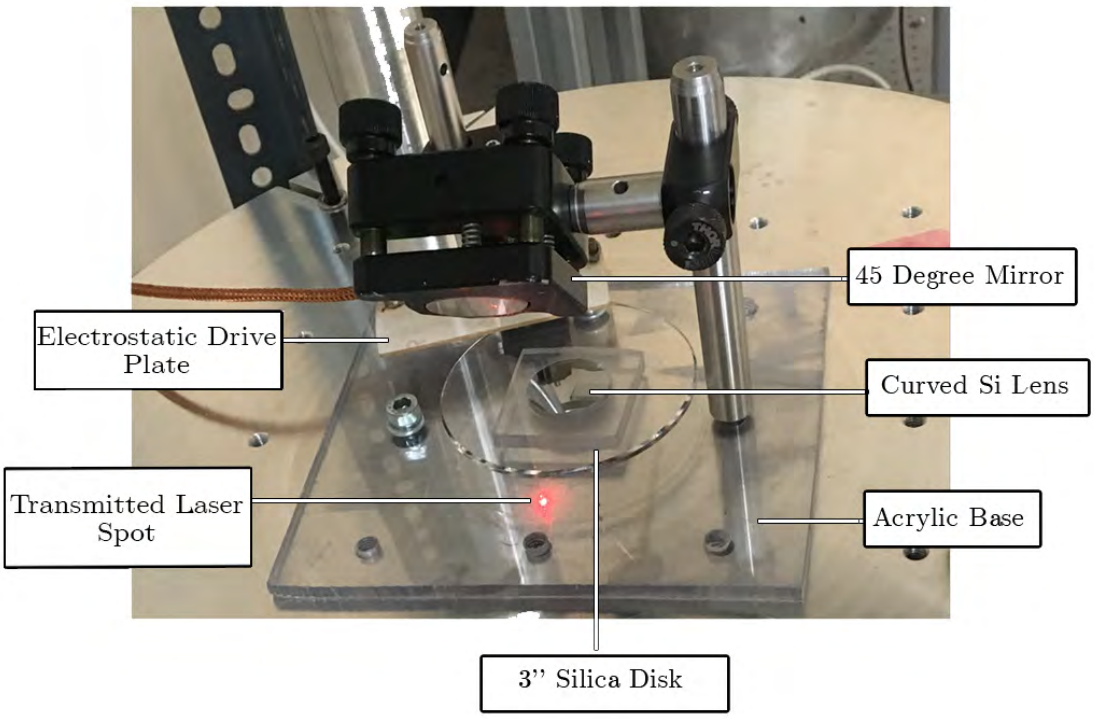

Figure 5.10: Image of a $\varnothing = 3$" (76.24 mm), $t = \approx 0.1$" (2.7 mm) silica disk successfully balanced in the prototype static GeNS. Acrylic base plate, exciter plate and 45° mirror also pictured.

Fig 5.10 shows a balanced $\varnothing \approx 3$" (76.24 mm), $t \approx 1$" (2.7 mm) silica disk in the prototype apparatus. A commercial interferometer was used to measure the ringdowns of the sample, with a 45 degree mirror used to direct the laser beam onto

the sample. The laser beam is aligned to the top surface of the upper face of the sample which is shown in Fig 5.11. As the exact plane in which the sample lies cannot be easily seen by eye, the interferometer beam is calibrated and locked to a horizontal plane defined by a vessel of water. When the water is removed and the disk installed, small changes were made to its position on the lens until the interferometer signal equalled that obtained from the water, indicating that the disk was parallel to the base of the vacuum system. An ESD plate is then lowered, parallel to the upper face of the sample.

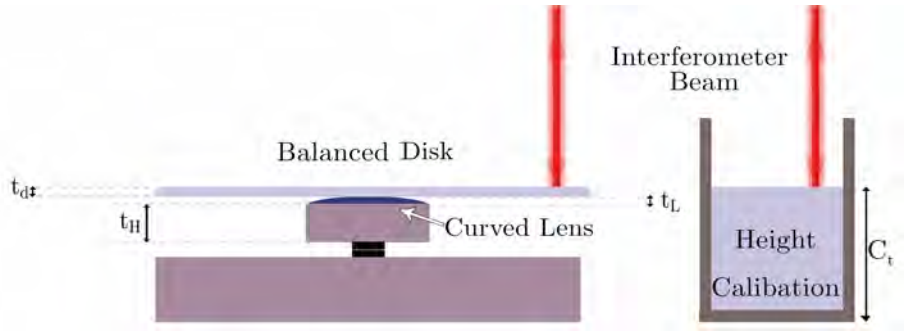

Figure 5.11: Illustrated representation of the calibration process of a SIOS interferometer using a water container.

5.1.3.2 Loss Comparison Between Suspension Methods

To test the apparatus, six ($m \geq 2,0$) and two ($m \geq 0,1$) mode pairs (between 2.8 kHz and 31 kHz) were measured on both the wire suspension and the GeNS. A comparison of these measurements is shown in Fig 5.12. Mechanical modes which previously showed signs of loss 'splitting' on the wire suspension do not exhibit this behaviour when measured using a GeNS. The degenerate mode pairs, are much closer in frequency, with nominally identical loss parameters in most cases, highlighting key differences in loss measured using each suspension method. It is also interesting to note that some mode pairs when measured on the wire suspension produce lower losses than the same modes measured on the GeNS. This may be an example of energy transfer between different modes which is facilitated by the wire suspension, leading to an artificially lower mechanical loss. This could also suggest that there are other sources of loss which have not been minimised. Individual ringdowns measured using the GeNS no longer exhibit the characteristics of a single component decay shown on the wire suspension.

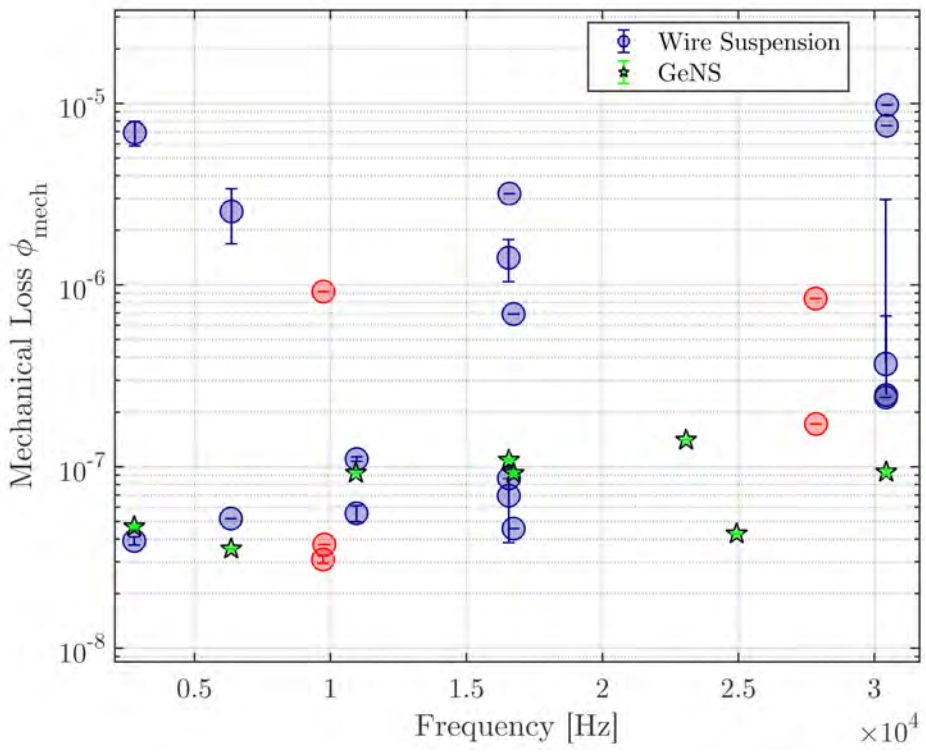

Figure 5.12: Comparison of mechanical losses measured on the same 3"(76.24 mm), $t \approx 0.1$" (2.7 mm) SiO_2 disk using a wire and GeNS suspensions. Error bars denote the spread in ϕ_{mech} from the same sample, over multiple suspensions, on each apparatus. Frequencies highlighted in red cannot be measured using the GeNS suspension.

It is also noted that low-frequency ($m \geq 0,1$) modes at 4.2 kHz and 9.7 kHz (highlighted in red) can not be measured on the GeNS. With increasing mode frequency, the amount of relative displacement at the disks' centre decreases, and the next mode with a high proportion of motion near the centre of the sample, at 16.8 kHz, can be successfully measured. This relative shift in mode shape can be described in terms of the amounts of bulk and shear energy stored in each mode shape, which is discussed in Section 5.1.1.

Fig 5.13 shows a comparison of ringdowns of the same mode measured with the two support systems. The loss measured on the wire suspension is 1.2× higher than the same frequency measured on a GeNS support.

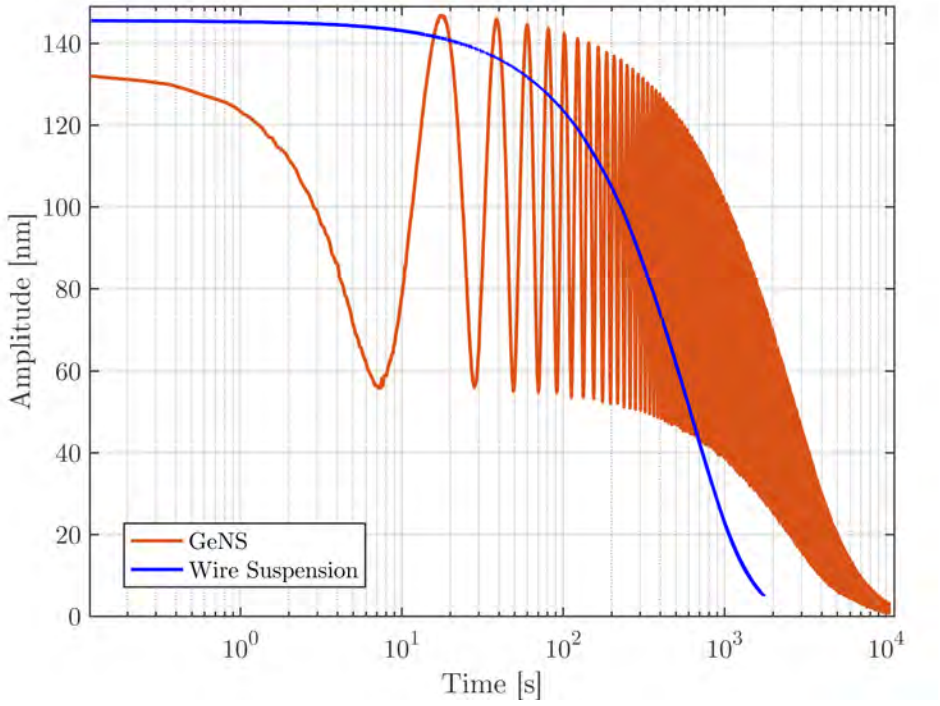

Figure 5.13: A ringdown comparison of the 2.8 kHz (2,0) mode of a 3"(76.24 mm), $t \approx$0.1" (2.7 mm) SiO_2 disk, measured on a wire suspension and on a GeNS support. Time has been plotted a log scale such that the effects of beating are clear.

The measured ringdown measured on a GeNS contains the ringdown of both component mode pairs. In order to reproduce a ringdown with the same functional form using a wire suspension the position of the wires would need to be far from any anti-nodes of displacement on the sample, making this difficult to achieve. As on the GeNS, both modes are able to oscillate uninhibited, the ringdown can no longer be approximated as a single exponent, but must be considered as two component oscillations with different paramaters.

5.1.3.3 Non-Linear Minimisation of Beating Decays

We will assume that the measured signal arises from a pair of exponential decays from two modes in a degenerate pair, A and B:

$$A = A_1 e^{i\varphi_1 - t/\tau_1 + it\omega_1}, \quad B = A_2 e^{i\varphi_2 - t/\tau_2 + it\omega_2}. \tag{5.13}$$

where φ_1, φ_2, are the respective phase of each wave where the time dependent phase $\varphi_2 = \Delta\varphi + \varphi_1$ and decay constants decay constants τ_1 , τ_2. The amplitude A_1 and

A_2 of each component are assumed to be within the bounds of the upper and lower envelope of the curve. Assuming that each decaying wave contains information about the other, the sum of both exponential terms yields the following expression by factoring out the average frequency

$$e^{-it(\omega_1+\omega_2)/2}\left[A_1 e^{\varphi_1+it\omega_1+t/\tau_1} + A_2 e^{\varphi_2+it\Delta\omega+t/\tau_2}\right] \tag{5.14}$$

The period of the oscillation $\Delta\omega$ is expressed in terms of the beat frequency f_{beat} where $f_{\text{beat}} = f_2 - f_1$:

$$\Delta\omega = \frac{2\pi(f_{\text{beat}})}{2}. \tag{5.15}$$

These equations then give two components, a high and a low frequency oscillating component. Adding both of these components in quadrature gives the intensity of the decaying waveform, allowing a simple formulation of the interference between both waves to be obtained:

$$I^2 = \sqrt{|A^2 + B^2 + 2AB\cos(\Delta\omega t)|}. \tag{5.16}$$

This representation of a ringdown arising from two decaying wave-forms was tested against mechanical loss measurements and found to produce a good fit to the measured data as shown in Figures 5.15 and 5.16. Non-linear least squares fitting was used to fit Eq. 5.16 to the data, resulting in an R^2 of 0.998. However, a large sum of squares error value of 1×10^4 was found, and the fit was highly dependent on the initial estimates of the fit parameters. This requires a more robust method to estimate the magnitude of each initial parameter, to ensure the best possible fit.

Fig 5.14 shows a flowchart representation of each section of the fitting code. Before analysis is carried out, the mechanical loss measurements are first inspected, allowing regions which contain over excitation or noise to be removed. Following the method produced by G. Vajente [177], several steps are taken to produce the final fit parameters for each component decay. Initial estimate values of A_1 and τ_1 can be produced by fitting the raw data with a single exponential decay allowing this function allows upper and lower bounds to be placed on A_1 and A_2.

$\Delta\omega$ can be estimated by subtracting a mean exponential trend from the

measured data to produce a flat oscillation, a process known as 'de-trending'. The distance between common points on the decay, such as peaks or troughs is used as the initial guess for the oscillation period. Taking the power spectral density of this data produces two distinct peaks, representing the two decaying terms in the ringdown. The peak separation frequency is used to refine the estimate of $\Delta\omega$. In the case of a purely exponential ringdown, only one peak is produced by the power spectral density, setting $\Delta\omega$ to zero.

To estimate the difference in phase, $\Delta\varphi$ between the two components, the de-trended data is fitted with a cosine function, evaluating the best fit at 0, $\frac{\pi}{2}$,π, and $\frac{3\pi}{2}$. This fit produces an estimate $\Delta\varphi$, while also placing upper and lower limits on φ due to its periodic behaviour. At this point, initial guesses, or bounds have been placed on each parameter of Eq. 5.16, allowing the optimisation of the non-linear model to produce estimates of all six parameters. During this process, two parallel minimisations are performed. The first considers the case where $\tau_1 = \tau_2$, and the second where $\tau_1 \neq \tau_2$, minimising the residuals produced in each case. As there are only very specific cases where ringdowns on cylindrical samples can produce a ringdown which is purely exponential ($\tau_1 = \tau_2$) these results are weighted against the power spectral density result, and are only considered where the beating frequency is zero.

Fig 5.15 is an example of a two component ringdown which has been analysed using this method. The amplitude of the beating oscillation is low and the ringdown is close to displaying a single exponential decay, it would be previously expected that approximating this data as such would produce a good estimate of its mechanical loss. However, when the same data was analysed using this method, it is possible to reproduce the high frequency beat and obtain a loss for both of the degenerate modes. Also shown are the residual points, the difference between the optimised fit and the measured data. Tab 5.2 shows the minimised parameters produced using a single exponent and two component analysis.

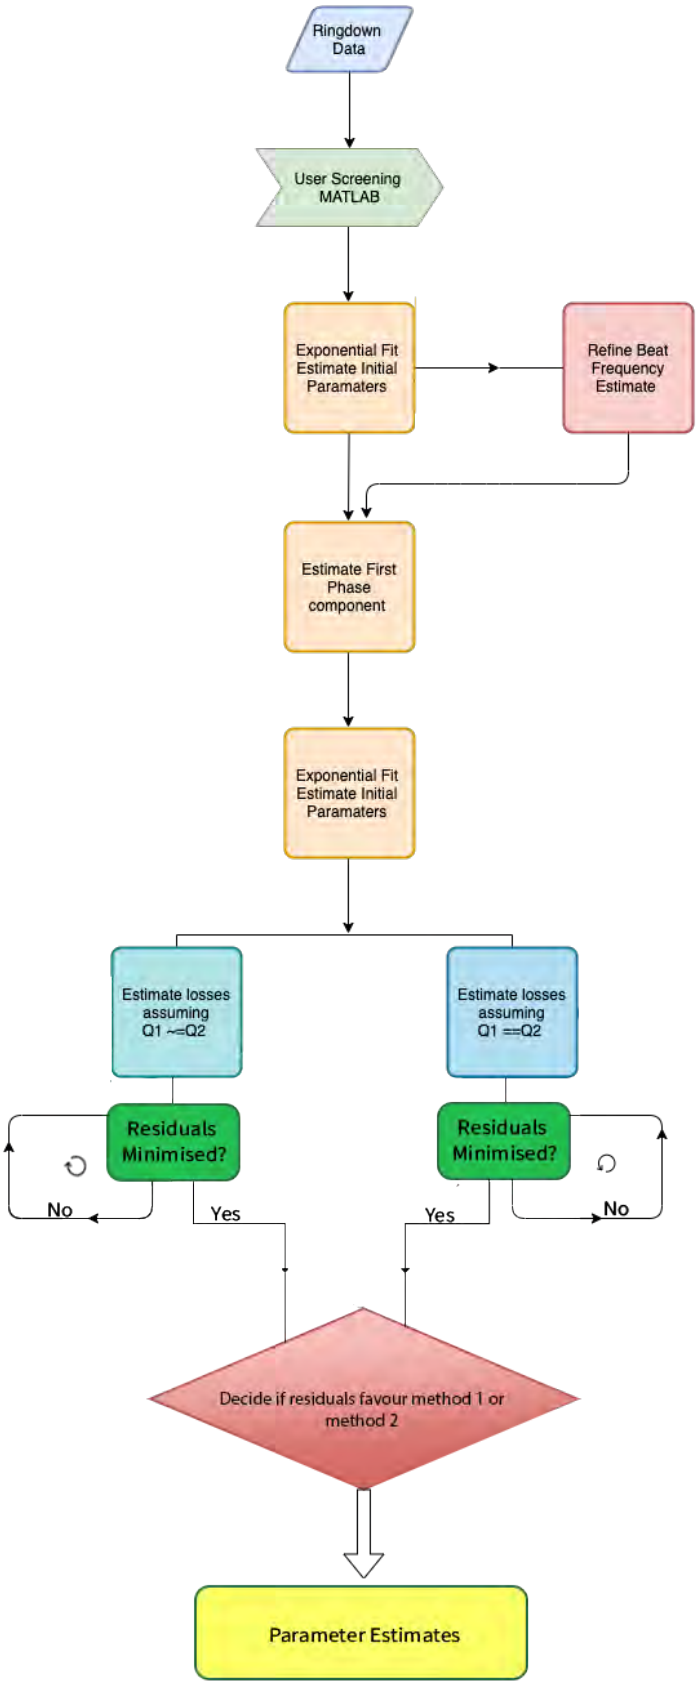

Figure 5.14: Flow diagram showing the major processes which are used to perform the optimised fit of beating ringdown data to Eq 5.16.

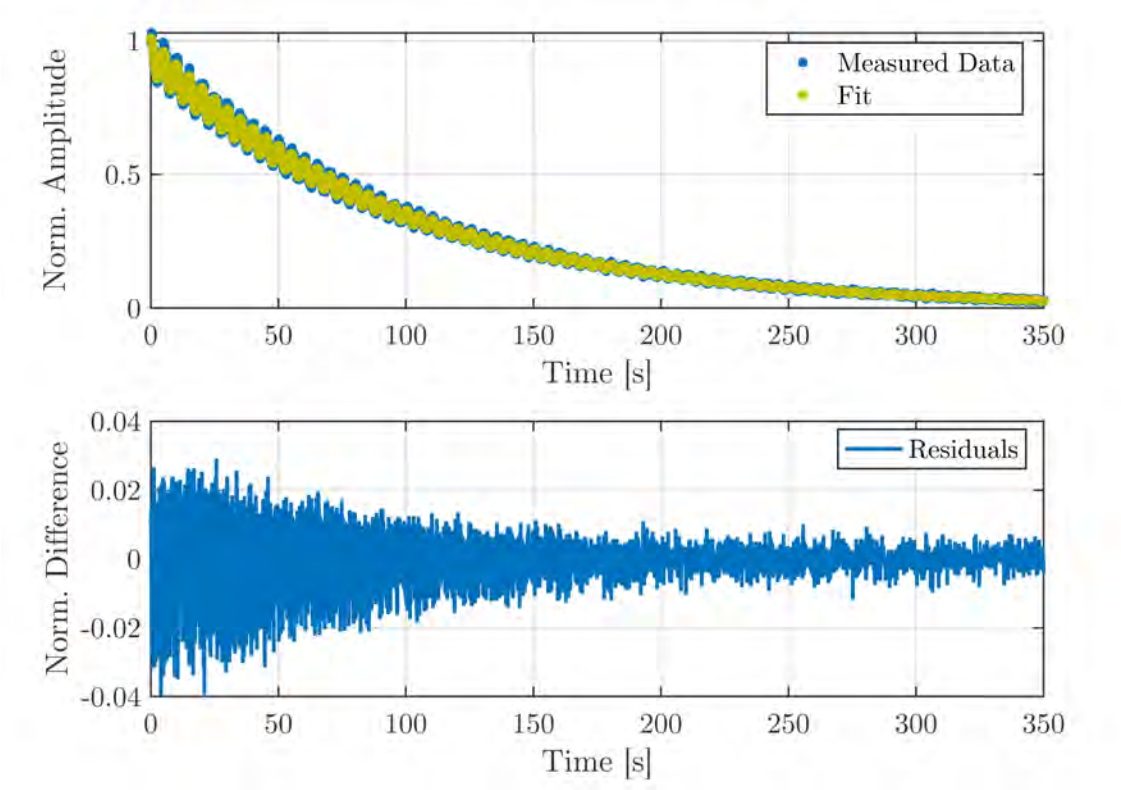

Figure 5.15: Upper plot shows a comparison between the measured ringdown of a 10.9 kHz mode on blank Disk 5. Lower plot shows the residual data from the fit where this is directly the difference between measured and fitted data.

In each case the calculated amplitudes have been re-scaled to reflect the input data, and τ_n values have been converted to mechanical loss ϕ_n, where $\phi_n = \frac{1}{f\tau_n\pi}$. For comparison, the parameters which can be extracted from a single exponential fit are also detailed. The level of mechanical loss estimated by the single exponential fit and two component method are very similar for this data set, with the estimated values of ϕ_1 agreeing within their respective error. The second loss component ϕ_2 is estimated close to ϕ_1, with a higher uncertainty on its value.

Parameter	Non-Linear Minimisation		Single Term Exponential	
	Value	(Lower Bound: Upper Bound)	Value	(Lower Bound: Upper Bound)
A_1	26.9	-	24.74	(24.66 :24.78)
A_2	1.9	-	-	-
ϕ_1	2.93×10^{-7}	$(2.93\times10^{-7}: 2.94\times10^{-7})$	2.95×10^{-7}	$(2.94\times10^{-7}:2.96\times10^{-7})$
ϕ_2	2.95×10^{-7}	$(2.85\times10^{-7}:3.06\times10^{-7})$	-	-
f_{beat}	0.2 Hz	(0.2:0.2)	-	-

Table 5.2: Calculated parameters using non-linear minimisation compared to single exponent fitting of a 10.9 kHz beating decay measured on a $\varnothing = 3$",$t = 2.7$ mm silica disk.

At the start of the ringdown, the fit shows a good approximation of the oscillation, as there is no visible periodicity present in these points. The residual points detail that the calculated amplitudes of the decays differ by ±0.03 from the measured data at the start of the ringdown, decreasing to ≈ 0 after 200 s with a total sum of squares error (SSE) of 0.242. This is clearly a much better fit than obtained through the single exponential model (SSE 949.8).

Fig 5.16 shows the same ringdown of the 2.8 kHz ringdown from Disk 5, presented in Fig 5.13 after being analysed using this technique. The estimated values for each parameter, and their respective errors are detailed below in Tab 5.3.

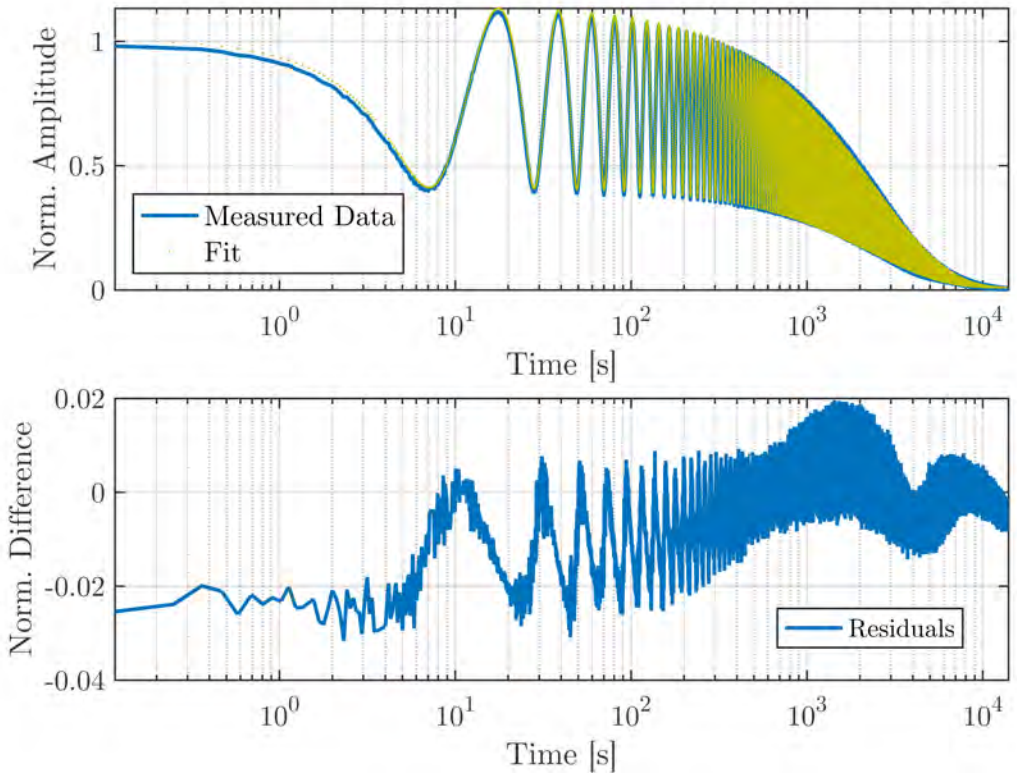

Figure 5.16: Upper plot shows a comparison between the measured ringdown of a 2.807 kHz mode on blank Disk 5. Lower plot shows the residual data from the fit where this is directly the difference between measured and fitted data.

This results values of ϕ_1 and ϕ_2 which differ by only $4{\times}10^{-9}$. Beating decays which do not decay symmetrically, show signs of excess damping from external sources. Pairs of modes which are close to one and other in frequency can transfer energy between each vibration, known as 'coupling' [118]. This can cause an artificially higher loss to be measured for these frequencies, due to damping from higher-order modes. This is discussed further in section 5.1.4.

Parameter	Non-Linear Minimisation		Single Term Exponential	
	Value	(Lower Bound : Upper Bound)	Value	(Lower Bound : Upper Bound)
A_1	83.3570	-	83.3570	(inf : inf)
A_2	36.5827	-	-	-
ϕ_1	4.627×10^{-8}	$(4.624\times10^{-8}:\ 4.629\times10^{-8})$	4.839×10^{-8}	$(4.74\times10^{-8} : 5.03\times10^{-8})$
ϕ_2	4.229×10^{-8}	$(4.423\times10^{-8} : 4.435\times10^{-8})$	-	-
f_{beat}	0.048 Hz	(0.048 : 0.048)	-	-

Table 5.3: Comparison of calculated parameters using non-linear minimisation and single exponent fitting of a 2.807 kHz beating decay.

In a single point suspension, such as GeNS, if there are no interactions between the suspension and the sample, it would be expected that both modes should oscillate at the same frequency. The calculated separation frequency is suggestive of small deformations on the sample from a perfect cylinder. As the sample has effectively no motion in the interface region for this mode, it seems unlikely that interaction between the suspension and sample is the cause of the frequency splitting.

The residual of the fit shown in Fig 5.16 shows a degree of periodicity matching the period of the original measured data. This indicates that there is something in the measured data which is not being taken into account for by Eq. 5.16. Analysing the residuals using the same method detailed above shows that there are two oscillating components contained within the residual data: a faster oscillation, whose period matches the residual data, and a slower oscillation which lasts on the order of hours - possibly the result of temperature fluctuations in the laboratory. [2]

In the following sections any mechanical resonances measured on silica disk substrates will be analysed using the method described above. Any differences in the contributions of each mode pair will be discussed.

[2] This conclusion requires further experimental testing to confirm this hypothesis. This work will be explored at a later date, beyond the bounds of this thesis.

5.1.3.4 Analysis and Repeatability of Mechanical Loss Measurements

To test the repeatability of the GeNS measurements, a different disk (Disk 3) was suspended three times and the loss of a number of modes measured. The loss of each mode was then calculated using the method described in section 5.1.3.3. The results are shown in Fig 5.17, where blue and red points denote the first and second components of degenerate mode pairs.

Discounting the higher loss measured on the second suspension at 16.5 kHz, the measured loss for all modes shows much higher repeatability when compared to measurements of the same sample on a wire suspension, with a calculated standard deviation of <1.5 Hz variation in mode frequency position and a standard deviation of 4.5×10^{-9} in the average mechanical loss. This significant change in the amount of spread between measured losses on a single sample suspension is due to the reduction in suspension damping. Without this systematic artefact of the measurement, a clear trend in measured losses emerges with frequency and mode shape. It is hypothesised that the more extensive spread in measured mechanical loss shown at 16.5 kHz can be explained when mode-coupling [118] is taken into account. As there are two mode pairs at 16.5 kHz and 16.7 kHz, the small separation in frequency allows a small percentage of vibrational energy to be exchanged during excitation.

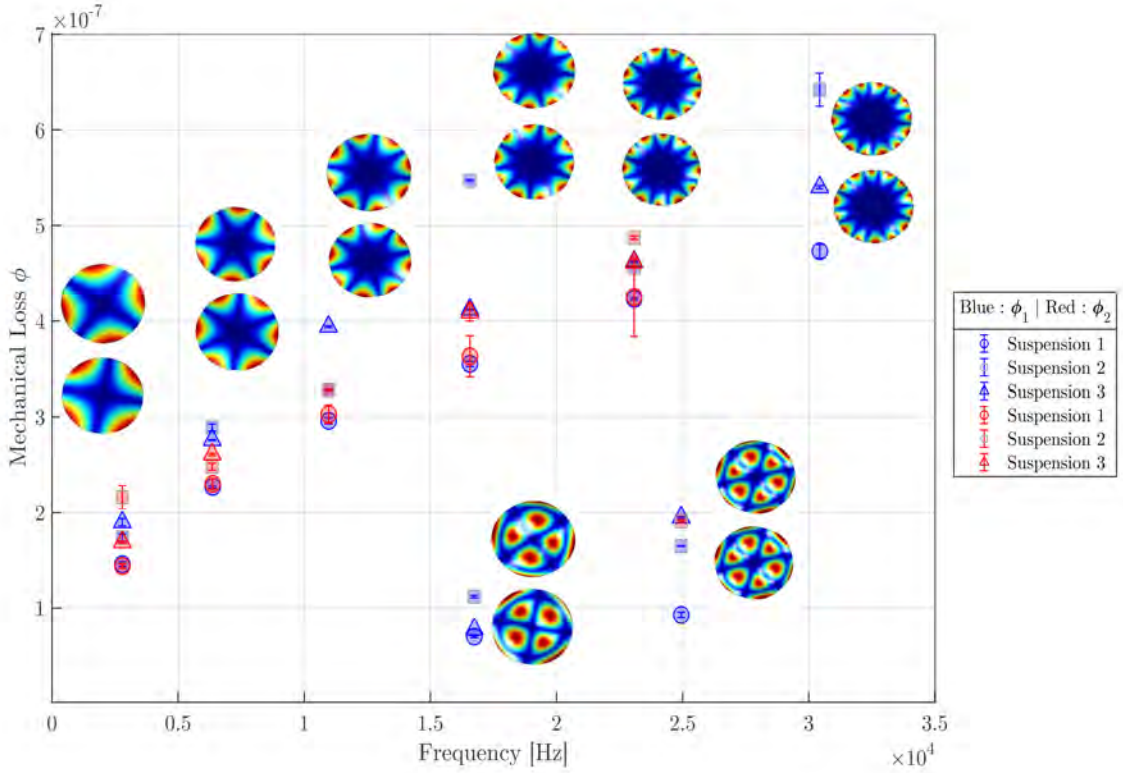

Figure 5.17: The first three suspensions of blank disk 3 on the prototype GeNS support system, showing the associated mode shape with the measured frequency. Loss components ϕ_1 and ϕ_2 are denoted in blue and red respectivley.

This could reveal some inconsistencies in the manual suspension of the sample. It is hypothesised that the measured loss of the sample could exhibit a dependence on its respective angle θ to the curved lens, changing the level of frictional damping. It would be later observed that the plastic base plate on which the apparatus is built is susceptible to environmental temperature fluctuations resulting in changes to the position of the samples over time. As the lens holder was also made of the same material, this was replaced with a steel counterpart with a much lower coefficient of thermal expansion.

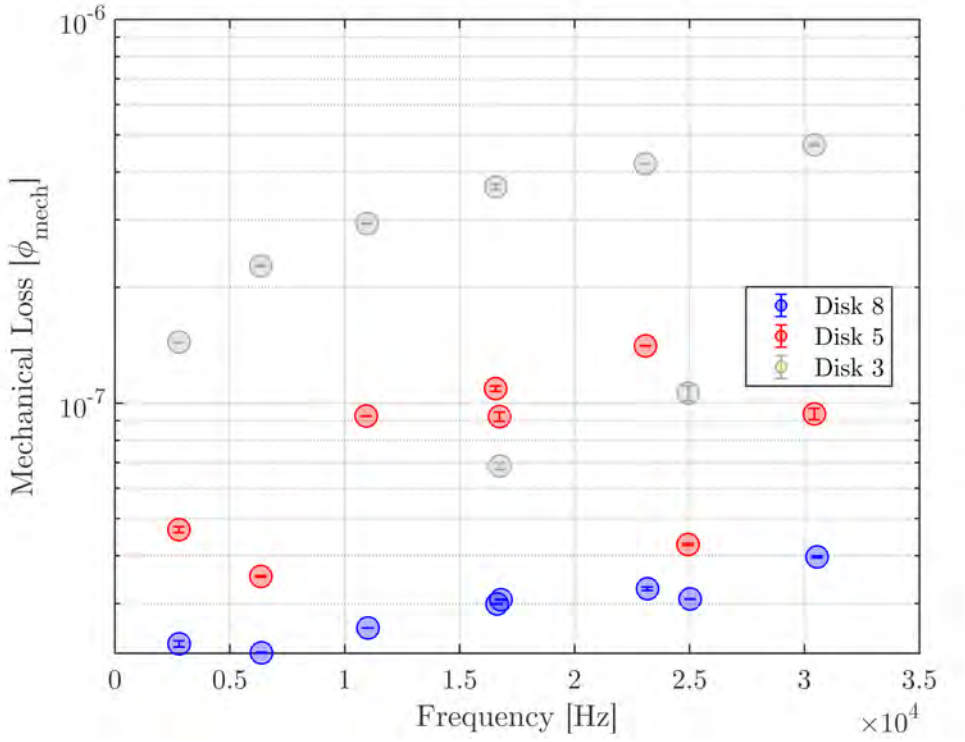

Figure 5.18: Comparison between the lowest measured losses from consecutive suspensions on Disk 3 and Disk 5. Plotted error-bars represent the standard deviation of the measured loss, from the suspension on which it was measured.

Figure 5.18 compares the lowest mechanical losses measured on Disk 3 and Disk 5 over multiple suspensions. With this comparison it is clear that the loss of Disk 5 is an order of magnitude lower than the losses measured on Disk 3.

To confirm that the changes made to the apparatus did not negatively impact the losses measured on disk samples, a further set of measurements were carried out on each Disk, revealing that the set-up change had not negatively impacted the loss, allowing the confirmation that the loss of the sample was not being overall limited by the suspension method.

After changes to the set-up were made, it does not affect the frequency at which the modes are found. However, these subsequent measurements show a smaller average standard deviation compared to the measurements taken on the acrylic apparatus. As the measured frequencies and losses do not change after altering the setup, it can be assumed that the increase in loss repeatability can be attributed to

the removal of acrylic parts from the system, increasing the rigidity of the system.

A comparison of two uncoated disks from the same manufacturer shows a large difference in $\phi(f)_{\text{mech}}$ in each sample. Disk 5 shows a general increasing trend over the measurement range approximated by $\phi(f)_{\text{mech}} = 1.5 \times 10^{-8} f + 8 \times 10^{-8}$, matching with the expected trend in loss shown on similar measurements [142, 177]. However, loss measurements of Disk 5 on the same apparatus, shows a much larger frequency dependence for ($m \geq$2,0) modes with large deformations near the edge of the disk. Isolating the losses just from ($m \geq$2,0) shows a trend of $\phi(f)_{\text{mech}} = 1.1 \times 10^{-8} f + 2.1 \times 10^{-8}$. This trend is much more comparable with the losses measured on Disk 8. As the average loss of Disk 8 was an order of magnitude lower than on Disk 5, this implies that there must be some physical difference between the two samples.

For modes which are dominated by motion near the edges of the sample, any defects around the circumference of the sample will result in a larger measured loss. Studies conducted on similar disks with mechanically ground edged by A. Cumming [96] and M. Granata et al. [142] show the effects of barrel surface quality relating to the mechanical loss, showing that the frequency-dependent loss associated with these modes is directly related to the quality of the barrel surface. Samples which have been mechanically ground or more etched edges have a higher loss compared with samples which have been mechanically or laser polished [96].

Disk 3 Disk 5 and a third sample (Disk 8) were then reviewed under an optical microscope.This revealed that the surface quality of the barrel of both disks was not finished to the same standard. The barrel of the comparatively higher loss sample (Disk 3) showed signs of mechanical grinding in the form of dark lines across the barrel's surface. Comparing this with the quality of the lower loss sample (Disk 5), these darker regions are not found in the sample, indicating that the sample was polished using a different method. Disk 3 which exhibits the highest loss contributions from ϕ_{barrel} shows signs of mechanical grinding on the top and bottom chamfered edges (Fig 5.19). This is seen to a lesser extent on Disk 5 (Fig 5.20). Each of the chamfered faces on Disk 8 show little/no signs of mechanical grinding and are thought to have been optically polished. It is thought upon review of the work by A. Cumming [96] that this is responsible for the measured changes in mechanical loss.

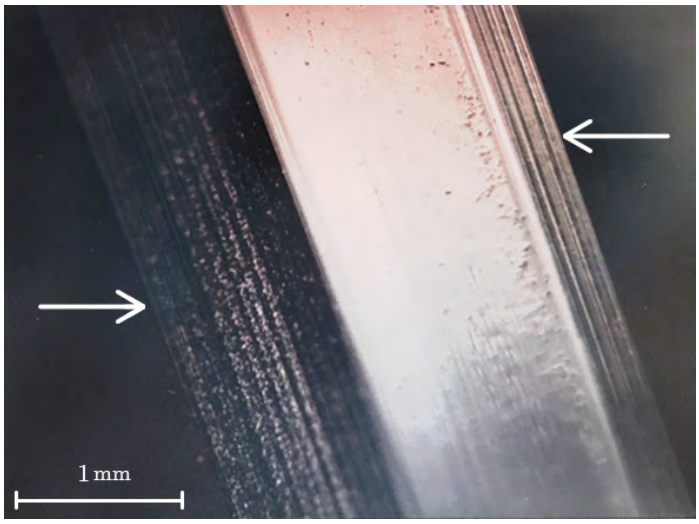

Figure 5.19: Optical microscope image of the barrel of Disk 3. Each edge of the disk is highlighted.

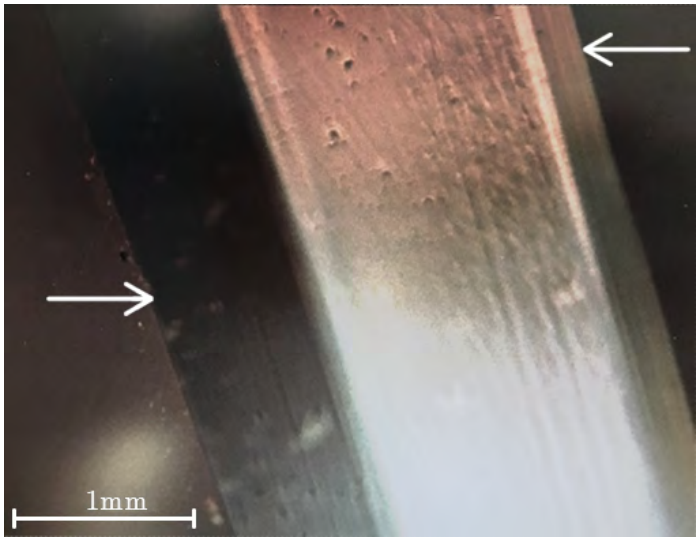

Figure 5.20: Optical microscope image of the barrel of Disk 5. Each edge of the disk is highlighted.

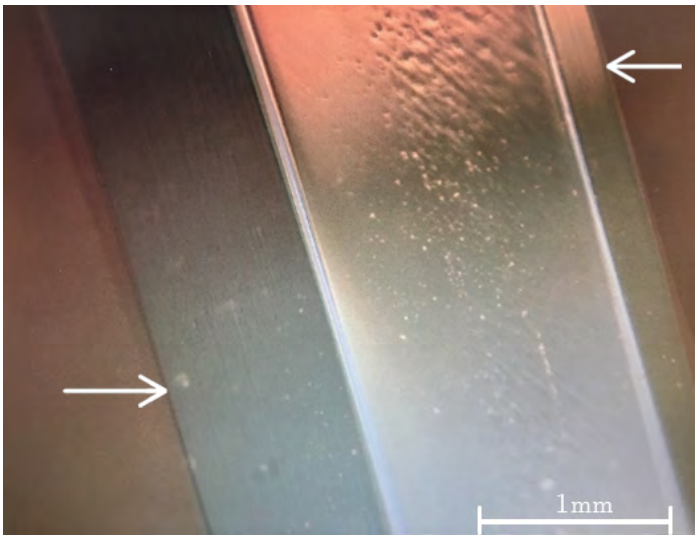

Figure 5.21: Optical microscope image of the barrel of Disk 8. Each edge of the disk is highlighted.

As the loss of modes such as ($m \geq 2$,0) modes have anti-nodes of deformation around the circumference of the sample, there is an interaction between the edge of the disk and these modes. The smoother polish is correlated with low mechanical loss, reducing the surface loss contributions from these edges [96, 208].

This hypothesis was then confirmed upon contacting Gooch and Housego$^{\text{TM}}$, that the samples were in fact from different batches and different polishing methods were used. It then seems likely that the surface quality of the disks dominates the losses of the sample in these regions. Work carried out by M. Granata et al. [142] has shown that the trend in losses measured on mechanically ground disk samples, can be quantified when the loss of the samples barrel ϕ_β is taken into consideration. His work showed that modes with anti-nodes of deformation near the disks circumference will be more sensitive to the surface quality of the sample. If the level of loss measured on modes with different amounts of surface deformation in these regions is compared, the measured loss is lower than for any other mechanical-mode and does not exhibit the same frequency dependence. It can be shown that the loss contribution of the barrel to ϕ_{mech} on cylindrical samples, can be calculated for modes dominated by ϕ_β using a frequency-dependent power law :

$$\phi_{\text{mech}}(f) = af^b + \epsilon d\phi_\beta. \tag{5.17}$$

a, b, ϵ and d can be determined by fitting the frequency-dependent points with

Eq. 5.17. From this analysis allows the contributions of $\epsilon d\phi_\beta$ to the total measured loss be estimated as the difference between the losses of the two mode-shapes. Applying this method to the measured mechanical loss for Disk 5, estimates a difference of $\epsilon d\phi_\beta = 1.5\times10^{-7}$.

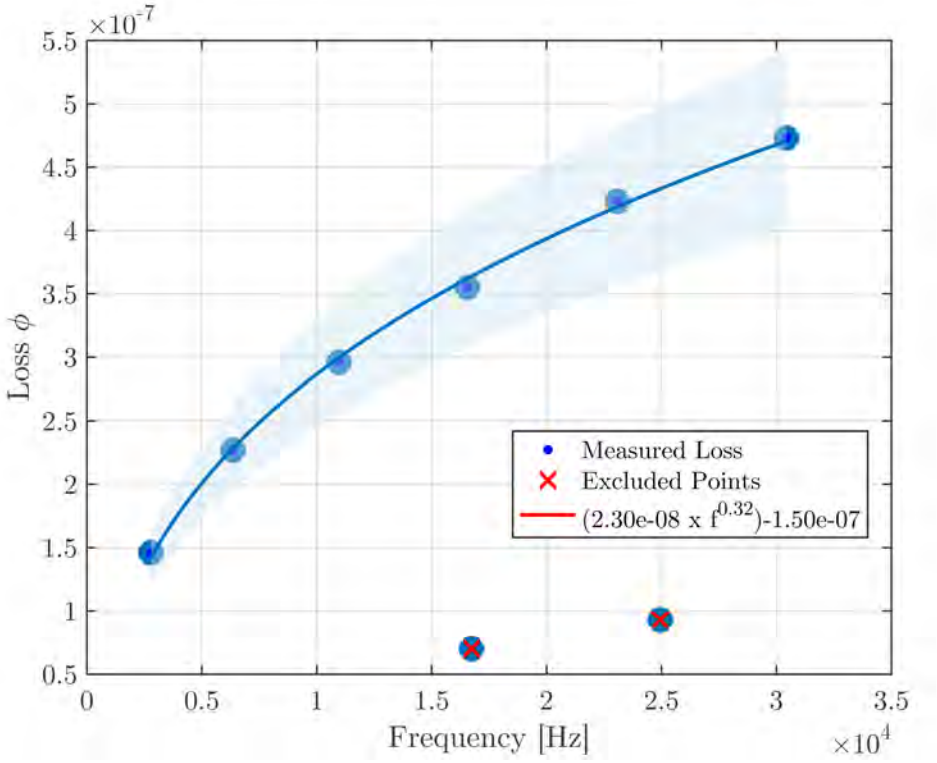

Figure 5.22: Demonstration of power law approximation for lowest losses measured on uncoated Disk 3. Shaded area around fit line follows an 15% error on the fitted curve.

In this fit the mechanical losses for modes at 16.6 kHz and 25 kHz are excluded as they do not exhibit the same frequency dependence as ($m \geq 2,0$) modes. Comparing this with the difference in loss measured between the ($m \geq 2,0$) and ($m \geq 2,1$) modes of 2.09×10^{-7} is largely in agreement with the values fitted values for this data-set. As the ($m \geq 2,0$) modes for this sample show a much larger frequency-dependent loss, this sample is therefore dominated by ϕ_B for these modes. Disk 8 does not exhibit the same level of frequency dependence on the measured loss values, which can now be fully explained by the surface finish on the barrel of the two samples.

5.1.4 Automated GeNS

5.1.4.1 System Design

Although the use of the GeNS improved the repeatability of loss measurements between suspensions compared to the wire suspension, Fig 5.17 shows that ϕ_1 and ϕ_2 can differ between consecutive suspensions. This difference in loss is expected to depend on the horizontal alignment of the sample. To remove the variation arising from hand-balancing disks, an apparatus was developed to raise and lower the sample without changing its horizontal orientation. This built on the work of G. Vejente [177] at Caltech.

Fig 5.23 shows a CAD representation of the system designed by R. Jones (University of Glasgow) and the author, which consists of a fixed supporting structure or 'table' and a translating stage which holds the sample and balances it on a mounted lens. A driving motor for raising and lowering the stage is fed through the bottom plate of the table and affixed to the system in the position marked in Fig 5.23. The apparatus must be able to place a disk sample on a curved lens, such that the contact radius between the sample and lens are concentric with the samples centre of mass.

The upper plate of the 'table' also contains a countersunk recess directly above the motor which holds the curved silicon lens on which disks are balanced. The lens is friction fit into a recess in the lens holder such that the final lens position is higher than the rest of the 'table'. The translating stage is composed of three thin rods with aluminium plates at either end. The upper plate, on which the sample rests prior to being balanced on the lens, has a central bore through the plate, which allows the plate to be easily raised and lowered around the lens holder, without risk of contact.

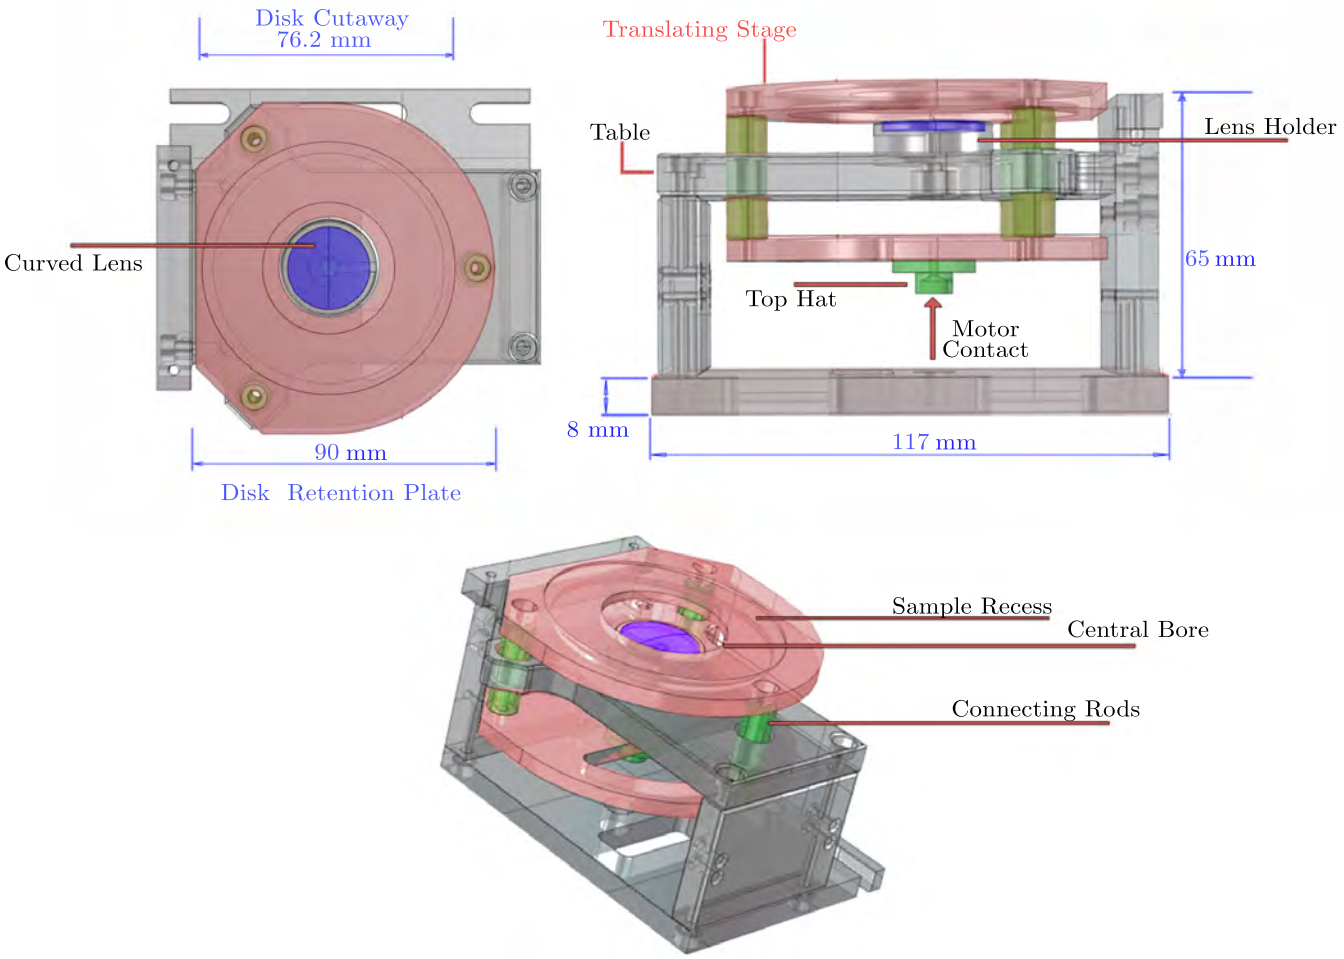

Figure 5.23: A SolidWorks rendering of the Automated GeNS disk measurement set up. Parts which are rendered in grey are aluminium, green rendered parts are made from a bearing material and blue, the silicon lens. Parts highlighted in red are capable of vertical translation. Central motor not pictured.

A second broader circular region $\varnothing = 3$" was counter-bored, to a depth equal to half the thickness of the disk samples measured previously. The lower plate is mainly flat, with a small region of material removed such that the centre of mass of the translational stage will be aligned with the centre of the lens holder when assembled.

The three connecting rods between the upper and lower plates are arranged in a triangular formation, which prevents the stage from pivoting out of the horizontal plane during translation. To reduce the slip/stick coefficient of the rods, they must

be machined to within a tolerance of a few 100's of nano-meters to the holes in which they are mounted such that movement of the stage is restricted in all axes, bar the vertical. A $\varnothing = 3" \times \approx 0.1"$ (76.24 mm $\times$ ~2.54 mm) cylindrical sample is then placed in the matched recess of the upper plate, where it can be balanced on the highest point of the curved silicon lens by lowering the translating table stage. The system was designed such that the only sources of vibrational energy are from the driving motor, while minimising external vibrations. For the stage to translate smoothly the three rods of the translation stage must have a low coefficient of static friction with the aluminium 'table' plate, allowing for the stage to move smoothly raise and lower, while minimising the force acting on the rods from the surrounding metal. As all materials adhere to LIGO vacuum compatibility standards [2], the choices of materials are limited.

5.1.4.2 Characterisation of Bearing Materials

Connecting rods made from polyether ether keytone (PEEK), polytetrafluoroethylene (PTFE) and phosphor bronze were tested with two different 'table' plates made from aluminium and stainless steel. The relevant coefficients of static friction μ_s are given in Tab. 5.4. All parts were cleaned in an ultrasonic bath to remove residual oils from the machining process and then baked for 48 hours at low temperature ($< 100\,^{\circ}$C) to reduce out-gassing. During this process, each material will expand slightly in accordance with its material properties. This expansion is taken into account in the machining tolerances of each of the metallic components.

A vacuum compatible picomotor[3] providing <30 nm positioning accuracy was used to move the translating stage. To increase the likelihood that the motor is driving on the centre of mass of the translating stage, an additional part known as the 'top-hat', shown in Figure 5.23, was manufactured from the same bearing materials listed in Table 5.4. The top-hat increases the driving surface area of the motor, such that driving force is spread equally and symmetrically around the centre of mass of the translating stage.

[3](Newport - 830X-UHV-KAP)

Material	Coefficient of Static Friction μ_s^{Al}	Coefficient of Static Friction μ_s^{Steel}	Thermal Expansion Coefficient ($\times 10^{-6}$ K^{-1})
Stainless Steel	0.21 [209]	0.78 [209]	9.5–12 [210]
Aluminium Polished	1.05 [209]	0.61 [209]	13.1 [211]
Polyether ether Keytone (PEEK)	0.1-0.24 [212]	0.04-0.35 [213]	46.8 [214] [118]
Polytetrafluoroethylene (PTFE)		0.04 [209]	0.69 [215]
Phosphor Bronze	0.35		17.8-18.4 [211]

Table 5.4: Thermal expansion and friction coefficients[4]of vacuum compatible bearing materials listed in LIGO vacuum compatible materials documentation [2].

Each combination of bearing and table materials was tested to find which provided the smoothest vertical translation. This process also allows the expansion of materials during the cleaning process to be tested. Any parts which expanded more than the tolerances accounted for during machining would no longer move smoothly through the system. The translation of the stage was judged on three criteria:

1. The stages ability to translate purely in the vertical plane, where any deviations from this could result in a sample being placed at a larger angle than tolerable.

2. Evidence of material expansion beyond tolerances which can cause the stage to stick or pivot during translation.

3. The resultant position of a cylindrical test sample after the stage had translated.

The position of a sample was again monitored during translation using a SIOS GmBH vibrometer [207]. Any deviations in the angle of the sample relative to an initial calibration point would result in a change in signal strength; thus, the resultant position of the sample could be determined. Two external cameras were mounted at 90 degrees to the GeNS system inside the tank, which allowed the stage's motion to be tracked as a function of time where any non-smooth translation could be reviewed.

[4]Values for the coefficient of static friction for material/material interfaces vary largely depending on the manufacturing process. Should be used for reference only.

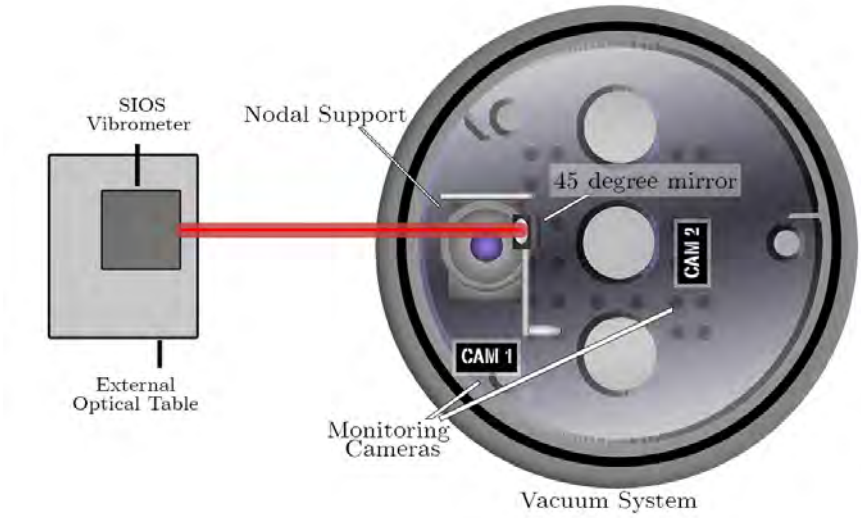

Figure 5.24: Illustrated representation of main experimental space inside the vacuum chamber. Shown from above, the GeNS system is mounted onto the tank where two cameras are placed at 90° angles to the system. A mirror mounted at 45° to the plane of the balanced sample is also shown.

Initial tests were carried out with stainless steel connecting rods and without a top-hat part on top of the driving motor. The difference in μ_s of the steel and aluminium was too large, resulting in the stage sticking, and failing to translate smoothly for prolonged periods. Without the 'top hat' part atop the motor, the rotation of the motor caused it to pivot across the bottom plate resulting in a precession of the stage as it translated. As a result, the sample failed to balance within a tolerable angle of θ.

Further translation tests were also carried out using connecting rods made of PEEK, a vacuum safe thermoplastic, and phosphor bronze alloy. Both of these showed an improvement in stage translation compared to stainless steel. Due to the machining quality of the phosphor bronze rods, there was not ample spacing between the rods and table plate to allow for a minimal procession during translation, which would cause the stage to become stuck. To allow the PEEK translation rods to be tested under vacuum, this material would first have to be baked to reduce the risk of outgassing [2]. To confirm that the parts were manufactured to the described tolerances, translation testing was carried out before and after the cleaning process. It was noted that before cleaning the stage

translated smoothly, without sticking. However, after the cleaning process, the parts expanded non-uniformly across its length and were no longer fit for purpose.

The last material tested was PTFE which has the lowest value of μ_s with steel and aluminium out of all materials tested (see Table: 5.4). This material pair proved optimum during translation testing with the steel table plate with minimal resistance between the material interfaces. Testing the material parts after the cleaning process showed a negligible change in the dimensions of the material and the PTFE rods were integrated into the system.

5.1.4.3 Motor Vibration Characterisation

One systematic feature which was apparent during translation testing were the vibrations produced when driving the picomotor. Higher translating speeds produce a higher frequency, larger amplitude vibration. It was observed that these vibrations could cause the sample to oscillate, rotate about its centre of mass and ultimately fall from the curved lens. The power spectral density (PSD) of the vibration was measured using a spectrum analyser, as shown in black in Figure 5.25 reveals a large amplitude signal generated at $\approx 2\,\mathrm{kHz}$ with periodic side-bands at $2\,\mathrm{n}\times f_o$.

Measuring the power spectral density (PSD) of the produced signal using a spectrum analyser, without a top-hat, the motor is in direct contact with the translating stage of the GeNS and produced vibrations and could then directly travel through the material of the table and into the sample.

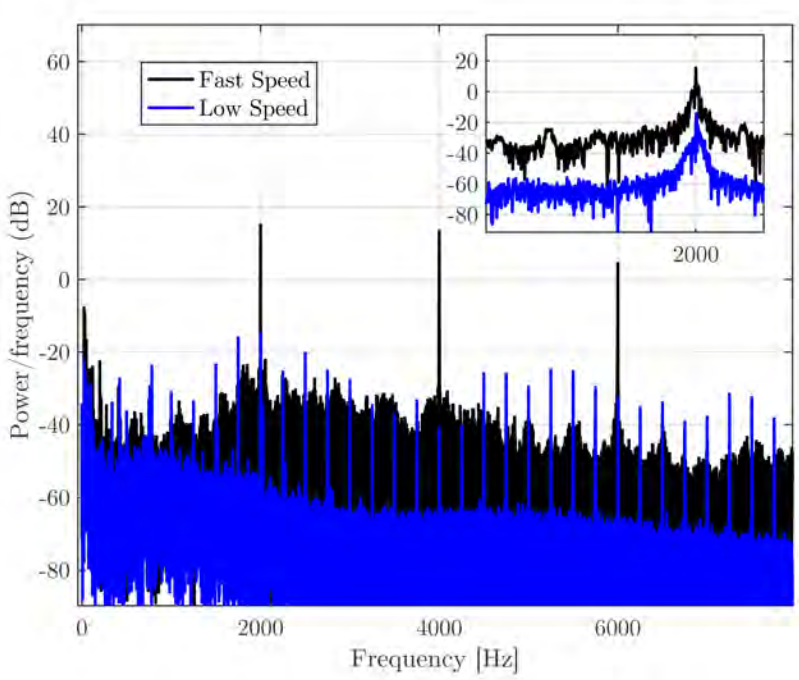

Figure 5.25: Power spectral density of the signals produced translating the picomotor at high and low speeds. The measured frequency of both signals is normalised between 0 and 1.

Given the proximity of the generated vibration to the fundamental butterfly (2,0) mode of a ⌀ 3" (76.2 mm) sample at 2.8 kHz, it seems likely that the produced vibrations are directly driving this resonance of the sample causing balancing failure.

It was observed that lowering the translation speed of the motor considerably reduces the frequency which is generated as shown in blue in Fig 5.25, where the large peak is now at 2.6 kHz and repeats at 200 Hz intervals. The peak amplitude is 29.76 dB lower than for the higher speed translation. Using the lower speed improves overall stability of the sample eliminating the sample rotation when lowering the stage. However, the vibrational coupling could still affect the position of the sample. A layer of chloroprene rubber was fitted below each table leg to damp the vibrational motion [216]. This allowed the stage to be translated at low speeds without influence on the sample.

5.1.4.4 Mechanical Losses of Automated Sample Suspensions

The automated suspension technique was tested by measuring the loss of number of modes of a coated silica disk, 3" (76.2 mm) in diameter by ∼0.1" (2.6 mm) thick. Measurements of the samples loss when balanced by hand are used as a control. To make a direct comparison on the repeatability of the measured loss using the automated system, the coated sample placed inside the system and $\phi(f)$ measured over three consecutive suspensions. It would be expected that any changes in $\phi(f)_{\text{mech}}$ between each automated suspension could be attributed to any deviations in the resultant position and balance of the sample.

Using the analysis methods described in Subsection 5.1.3.3, the losses of degenerate mode pairs shown in Fig 5.26 were evaluated for three successive automated suspensions.

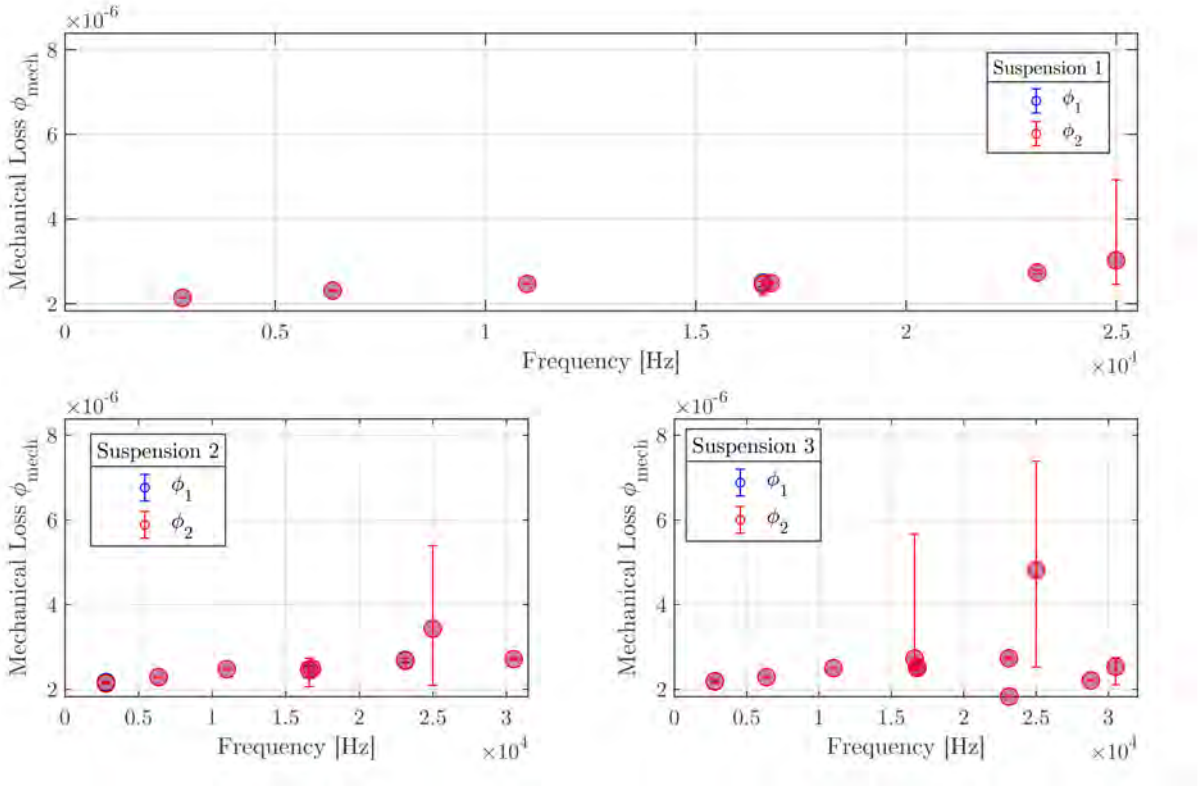

Figure 5.26: Comparison of the average mechanical losses ϕ_1 and ϕ_2 measured from three successive automated suspensions of a coated 3" (76.2 mm) in diameter by ∼0.1" (2.6 mm) thick silica disk. Error-bars denote the total spread of ϕ_1 and ϕ_2 values for each mechanical mode.

Each suspension of the coated disk[5] shows that the losses of mechanical modes below 15 kHz can be repeatably measured through multiple automated suspensions, varying by less than 3%. This repeatability is also observed for mechanical modes which do not show evidence of multi-mode interactions (MMI). Pairs of modes which are close together in frequency (Δf <1000 Hz) such as those occurring at 16.6 kHz and 16.8 kHz are measured to have a higher loss, due to the contributions of its degenerate mode pair. It is thought that this artificially higher measured loss can transfer energy to mechanical mode pairs which are sufficiently close in frequency, where the mechanical loss of each set of degenerate mode pairs can affect the other, resulting in a larger spread in measured mechanical loss values.

The effect of mode coupling is more prevalent on the second and third suspensions of the coated disk, with the spread in losses measured at 16.6 kHz and 25 kHz increasing with each suspension. Fitting the measured loss of each suspension using Eq. 5.17 allows for the contribution of ϕ_{barrel} from the barrel of the sample to be calculated. From this data is can be seen that the measured loss, for the majority, of measured frequencies is dominated by the loss contribution of the sample barrel (see section 5.1.3.4).

Suspension	Loss Component	$\epsilon\phi_{\text{barrel}}$ ($\times 10^{-6}$)
Suspension 1	ϕ_1	1.5
	ϕ_2	1.5
Suspension 2	ϕ_1	1.6
	ϕ_2	1.7
Suspension 3	ϕ_1	2.0
	ϕ_2	1.8

Table 5.5: Calculated loss contributions of $\epsilon\phi_{\text{barrel}}$ from three successive suspensions of a coated 3" (50.8 mm) in diameter by $\sim$0.1" (2.6 mm) thick silica disk.

Comparing the estimated contributions of $\epsilon\phi_{\text{barrel}}$ shown in Tab. 5.5, for each suspension, an increasing trend with the number of suspensions is observed. The

[5]Prototype HR coating material is composed of $aSi:SiO_2:Ta_2O_5$ and is discussed more in Chapter 6.0.1.

increasing spread in each losses measured from mechanical losses which exhibit MMI and the increasing predictions of $\epsilon\phi_{\text{barrel}}$ with successive suspensions are both indications that the placement of the sample during automated suspensions is not constant. To understand the impact that successive automated suspensions has on the estimated mechanical loss a given sample, the total spread in $\phi(f)_{\text{mech}}$ is compared to values of $\phi(f)_{\text{mech}}$ when the sample is manually suspended.

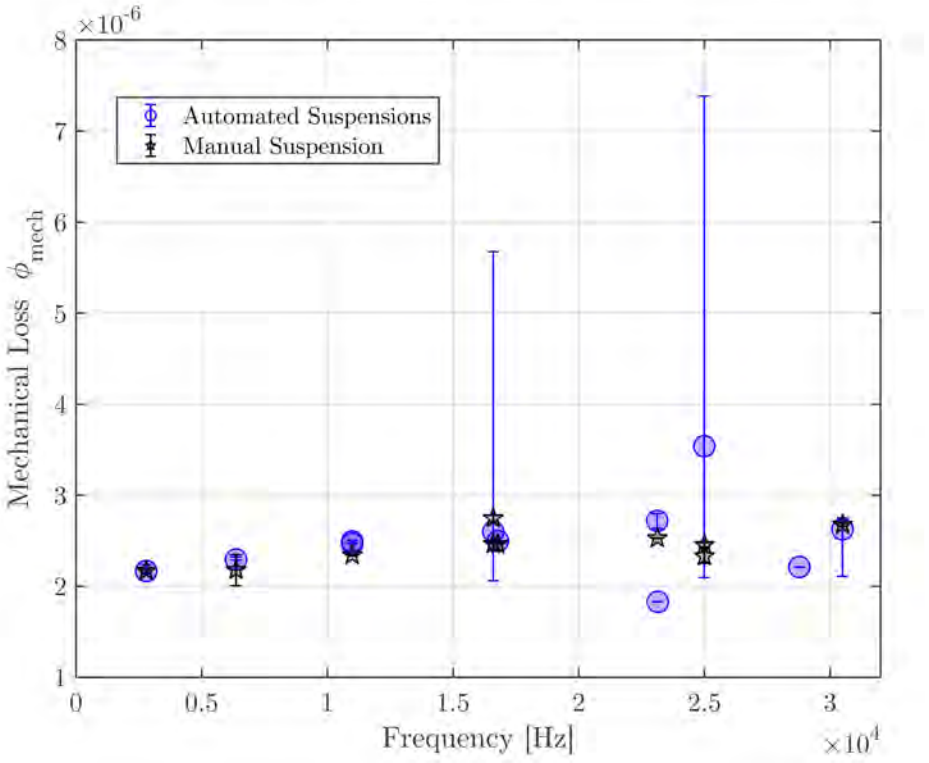

Figure 5.27: Comparison of multiple consecutive suspensions of a prototype highly reflecting coating material[6]deposited on SiO_2 disk measured with automated suspension. Manual suspension of the same sample used as a control.

Fig 5.27 compares the loss of the same sample for manual and automatic suspensions. The average loss at low frequencies (f<10 kHz) show considerable agreement within experimental error of the successive automated and manual suspensions. This shows that the lower limit of the measured loss is not compromised by the automated suspension method. Modes which show multiple mode interactions (MMIs) show larger disagreement between both methods, with the losses measured via manual suspension showing a slightly less spread at 16.8 kHz. The average losses measured at 16.6 kHz shows good agreement between

both suspension methods, with a much larger spread observed from the automated suspension method. Similar points are also found for MMI modes between 25 kHz with a higher average loss measured in the automatic loading case. The modes measured at $\sim$30.4 kHz show similar agreement between both suspension methods.

As it is conventional to report the lowest mechanical losses measured over multiple suspensions (with their respective errors) it is then pertinent to compare the lowest values of $\phi(f)_{\mathrm{mech}}$ from each automated suspension to the values produced by manual sample suspension.

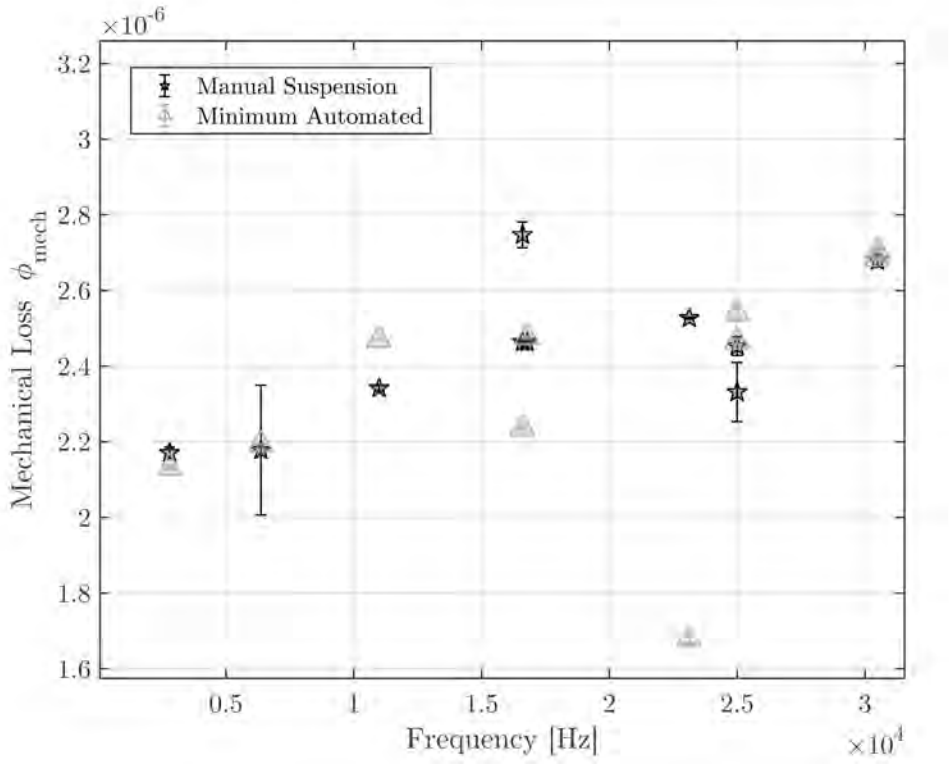

Figure 5.28: Lowest losses measured over consecutive suspensions of a prototype highly reflecting coating material deposited on SiO_2 disk measured with automated suspension. Manual suspension of the same sample used as a control.

Fig 5.28 shows the lowest mechanical losses measured over the three automated suspensions shown in Fig 5.26. Comparing these values to the same manual suspension losses shown in Fig 5.27 is a true comparison of the mechanical losses measured via both suspension methods. The spread in ϕ_{mech} measured from MMI modes at 16.6 kHz and 25 kHz is now of less consequence. Comparison of both

measured data-sets shows agreement between measured values of ϕ_{mech} produced through automated and manual suspensions of the same sample. In comparison to the single set of control data, the anomalously high loss measured from MMI modes using the automated suspension method are not a true reflection of the samples loss.

In order to correctly count for the contributions of MMI's, it is suggested that the systems can be modelled via an FEA harmonic analysis, to simulate the systems response to being driven at different frequencies. However this effect has not been simulated by the author at the time of writing and will be carried out at a later date.

5.1.5 Conclusions

In this chapter an automated gentle nodal support, adapted from the designs of. G. Cagnoli *et al.* and G. Vajente was designed by the author. Implementing this method of mechanical loss categorisation for disk substrates has decreased variations in mechanical losses measured over multiple suspensions and can increase sample throughput. This apparatus, through a single point of contact has minimal interaction with the sample of interest for certain modes and therefore reduces energy transfer from the sample producing a more accurate measure the loss of the sample. As there is minimal damping of resonant modes, pairs of modes which are close in frequency can be excited, entangling their ring-downs. In order to calculate the mechanical loss of each frequency component, a non-linear minimisation method developed by G. Vajente was used.

This experiment has demonstrated that mechanical loss measurements on an automated GeNS system can reduce the effects of external damping effects due to alternative suspension methods. This measurement technique produces a much more accurate reflection of the true loss of the sample. Measurements on an automated GeNS system, allows for consecutive measurements of the same sample, without having to open the vacuum system, thus reducing the overall time taken to measure $\phi(f)_{\text{mech}}$ without compromising on measurement accuracy.

Chapter 6

Room Temperature Mechanical Loss of Multi-material Coatings

Current gravitational wave detector end test mass (ETM) coatings consist of 19 bi-layers of silica and titania doped tantala (TiO_2-Ta_2O_5) deposited by ion beam sputtering (IBS) by Laboratoire des Matériaux Avancés in France [217] (see section 4.3.2), [131, 142]. By optimising the layer thickness to the wavelength of the main interferometer laser ($\lambda_{\mathrm{IFO}} = 1064\,\mathrm{nm}$[1]) alternating layers of high refractive index $n_{\mathrm{high}} = 2.12$ (tantala) and low $n_{\mathrm{low}} = 1.44$ (silica) can produce a highly reflective mirror coating with $R \approx 99.9999\%$.

The combination of both of these materials produces a final coating stack which absorbs 1% of the circulating power in the arm cavities which is crucial for detector duty cycle [94]. By doping Ta_2O_5 with TiO_2, it was found [145] that the crystallisation temperature of Ta_2O_5 could be increased. It was also found that this reduced the mechanical loss by $\approx 40\,\%$ producing a coating with $\phi \approx 2\times10^{-4}$ [145]. As coating absorption, α, and mechanical loss contributions of the coating layer produce sizeable contributions to the duty cycle and the level of thermal noise in GWD respectively (see Chapter 2), detector sensitivity is both facilitated and limited by these coatings.

[1] in practice the coating thickness is optimised for reflectivity at $\lambda_{\mathrm{IFO}} = 1064\,\mathrm{nm}$ and a secondary laser $\lambda = 533\,\mathrm{nm}$ which is used for lock acquisition.

In this section, measurements of a novel coating design which utilises layers of amorphous silicon aSi ($\phi \approx 2\times10^{-5}$ when heat treated at 500 °C [149]) to reduce the mechanical loss of an ETM coating stack while maintaining low optical absorption at $\lambda_{\mathrm{IFO}} = 2000\,\mathrm{nm}$. Companion measurements of the coatings optical absorption can be found in Chapter 4.

6.0.1 Coating Design

From the work of Hong *et al.* [132], Yam [134] and others [218, 219] it has been shown that the mechanical loss of a HR coating stack can be reduced by adding layers of low mechanical loss material such as amorphous silicon (aSi). Applying this concept to an HR mirror coating, such as is currently used in aLIGO / AdV could significantly reduce the level of thermal noise the coating produces. Another important consequence of these works [132, 134, 218, 219] show that the thermal noise contribution of coating layers increases closer to the substrate. Therefore, by replacing these layers with lower loss materials such as aSi, it is expected to further reduce the level of thermal noise produced by the HR stack.

The multi-material coating design studied in this chapter aims to take advantage of the low absorption of some SiO_2-Ta_2O_5 layers, to reduce the transmission of laser light to the more highly absorbing lower layers consisting of SiO_2-aSi [149]. This then reduces the total light intensity in the layers of highly absorbing aSi, allowing the coating to benefit from its low mechanical loss, but to not be constrained by its high optical absorption. This type of design allows for a reduction in the total coating thickness by reducing the number of layers (due to the high index of aSi) and a lower thermal noise due to the use of aSi instead of Ta_2O_5. The key results from Chapter 4 and this chapter have been published in Physics Review Letters [192].

The '*Full Stack*' coating is composed of 5 bi-layers of SiO_2 and Ta_2O_5 followed by 5 more bi-layers in which the Ta_2O_5 is replaced by aSi. Throughout this section, the five bi-layers of SiO_2-Ta_2O_5 will be referred to as the '*Upper Stack*', the five bi-layers of SiO_2-aSi, the '*Lower Stack*' and the total coating as the '*Full Stack*'. Each coating stack was produced by RLVIP deposition by Tafelmaier $_{\mathrm{GmbH}}$ [6] which was discussed earlier in more detail in Chapter 4. The RLVIP deposition

process is thought to increase the packing density of deposited coating layers, which can be linked to changes in two level systems in coating materials, thought to be responsible for mechanical losses [152]. During the coating deposition it was recorded that the internal chamber temperature reaches ∼250 °C, effectively 'pre-annealing' the coating materials to this temperature.

Fig 6.1 shows the calculated normalised electric field intensity (EFI) at λ =2000 nm of the *Full Stack* coating, allowing the working principle of the coating stack to be visualised. From Fig 6.1, the total electric field intensity at 2000 nm inside the first bi-layer of aSi-SiO_2 has reduced by more than a factor four, due to the light reflected by the *Upper Stack* layers. Therefore in these layers, less laser light circulates and optical absorption is of lower consequence. These materials have been previously investigated as single layer coatings, but the use of aSi layers in a multimaterial coating stack is a novel concept. It is unknown if thermally induced stresses due to mismatches in thermal expansion coefficients will alter the mechanical properties of each layer. As such, the expected mechanical properties of each material may differ when they are deposited in such a configuration.

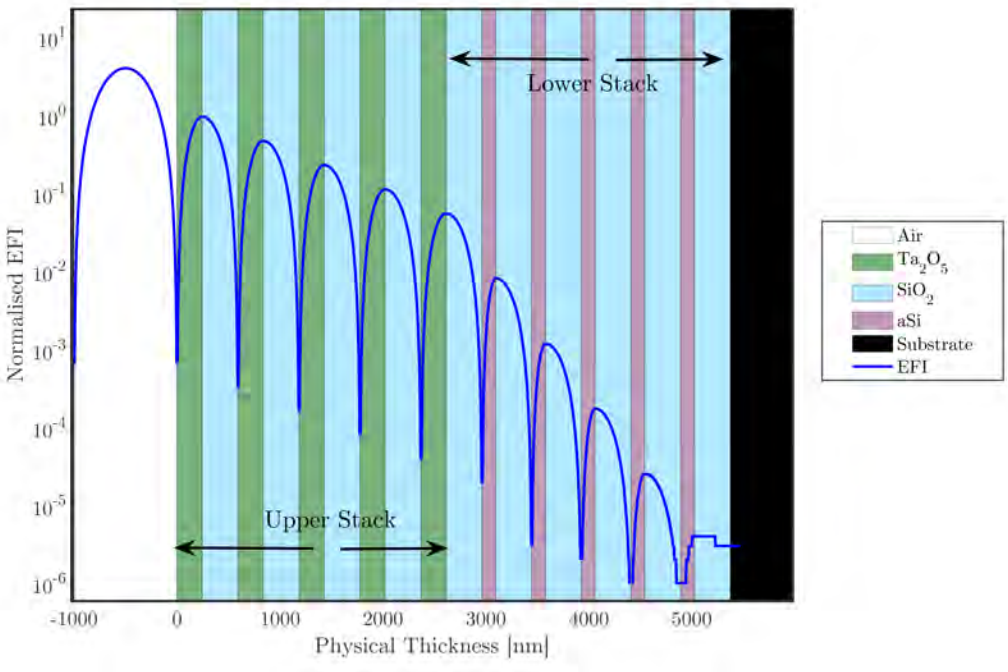

Figure 6.1: Calculation the of electric field intensity inside the *Full Stack* coating design. Light propagating from left (air) to right (substrate).

Measurements of optical absorption of each coating stack (see Chapter 4) suggest that incremental heat treatments can produce variations compared to coatings when directly heat treated, which could stem from this hypothesis. In order to provide direct comparisons between samples used for loss and optical absorption, each loss sample was incrementally heat treated in 100 °C steps.

To be able to relate any physical or chemical changes in the *Full Stack* coating during experimental measurements, the sample set was divided into three components. A variety of substrates for optical and mechanical loss were placed into the chamber to be deposited with initial layers of aSi-SiO_2. The deposition chamber was then opened and half of the samples inside removed, creating the *Lower Stack* sample set. These samples were then replaced by comparable blank samples, and the layers of SiO_2-Ta_2O_5 deposited producing samples coated in the *Upper Stack* coating alongside samples coated with the *Full Stack* coating. Pausing the deposition of coating layers in this manner produced three subsets of samples with each component of the coating stack. This process also ensured that each coating stack was produced in the same chamber under nominally identical conditions as all other samples in each subset, reducing possible variations between samples.

	aSi	SiO_2	Ta_2O_5
Youngs Modulus (Pa)	1.47×10^{11} [220]	7.20×10^{10} [178]	1.40×10^{11} [178]
Refractive Index (2 µm)	3.73	1.44	2.12
QW Thickness (m)	1.34×10^{-7}	3.47×10^{-7}	2.35×10^{-7}
Poisson ratio:	0.23 [178]	0.17 [178]	0.23 [178]
Thermal Expansion(10^{-6}/K^{-1})	0.01 [216]	(0.55 -0.99) [221]	(-44.3 : 3.6) [221]
Density(kg/m^3)	2330	2202	6850 [178]
Optimum Annealing Loss	2.0×10^{-5} [149]	3.20×10^{-5} [93]	5.66×10^{-4}
Optimum Annealing Temperature (°C)	400-500 [149]	800-900	500

Table 6.1: Single layer material properties of RLVIP thin film coatings used in the multi-material coating study. All optical properties are quoted for a laser wavelength of 2000 nm.

Tab 6.1 lists the single layer material properties of each coating material used in this study. Values used to predict the loss of the component stack are the loss values measured at the optimum annealing temperature for that material.

Performing a thickness weighted sum of the optimum annealing losses of each material should allow the loss of each component stack to be estimated [222]. However, as each material has different stiffness and thermal expansion properties, this could make it difficult to predict how each material will respond to post-deposition heat treatment. Previous measurements of mechanical and optical properties of single layer coatings with heat treatment show that the lowest energy state for each material differs, and as such it is difficult to predict how each material will behave in a coating stack. Each coating stack was studied in an 'as-deposited' state to provide baseline measurements of optical absorption and mechanical loss. Following this, each coating stack was subjected to the same heat treatment condition.

The following sections describe mechanical loss measurements of the *Full*, *Upper* and *Lower* stacks. These measurements are then used to calculate the thermal noise implications on a GWD.

To investigate the mechanical loss of each component stack, three ⌀=3" (76.2 mm) t≈2.6 mm, fused Corning 7980-0A grade SiO_2 disks manufactured by Gooch & Housego [223], were also present in the coating run described in Section 6.0.1. Prior to coating, each disk was heat-treated to 950 °C for four hours, and their mechanical losses measured using a GeNS system in section 5.1.4.1. Each sample was measured over three separate suspensions, with the minimum recorded loss plotted for each mode of each sample. A summary of the lowest losses measured on each blank sample as a function of frequency is shown below in Fig 6.2. Although the disks came from the same batch made by the same manufacturer, the loss of each was significantly different.

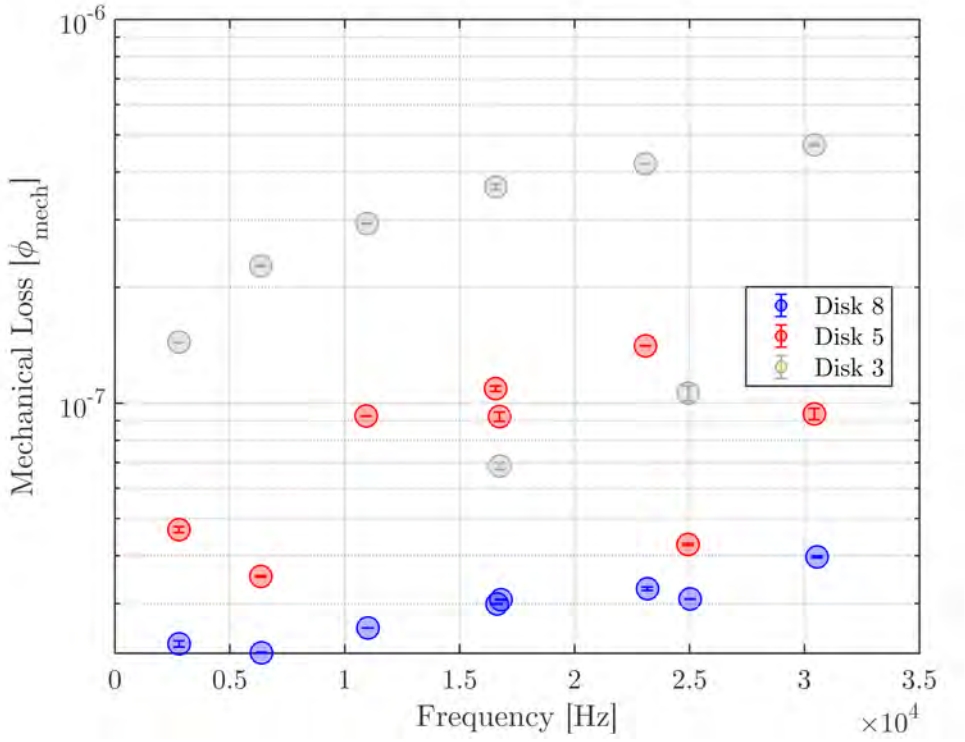

Figure 6.2: The minimum mechanical losses measured as a function of frequency for three uncoated Gooch & Housego ⌀=3" (76.2 mm) t≈2.6 mm, fused Corning 7980-0A SiO_2 disks.

The loss of Disk 3 shows a large frequency dependence reminiscent of those shown by M. Granata *et al.* [142] (see section 5.1.3.4). The loss of Disk 5 is lower than that of Disk 3, with less loss contribution from the disks' barrel (see Chapter 5.1.3.4). The lowest loss sample (Disk 8) has a minimum mechanical loss at 2.8 kHz ∼50% lower than the minimum value recorded for Disk 5. The change in level of loss, and clear separation of losses measured on Disk 3 between different mode families indicates that this sample is strongly influenced by the loss of the disk barrel ϕ_{barrel}.

On examination under an optical microscope, (see section 5.1.3.4) a correlation between the polishing quality of the samples' barrel and its mechanical loss was found.

After coating, each coated disk was measured in the as-deposited state before being heat-treated in 100 °C steps, starting at 300 °C, 50 °C higher than temperatures achieved during deposition. All coated samples were heat treated at

the same time in Carbolite large silica tube oven for 3 hrs and left to naturally cool inside the oven before being removed. After each heat treatment the mechanical loss of mechanical modes between 2 kHz and 31 kHz was measured over a minimum of three suspensions, categorising resonant frequencies between 2 kHz and 31 kHz for each disk. The losses of degenerate mode pairs were calculated using the method described in section 5.1.3.3. Once the mechanical loss of each coated sample is measured, the loss of the coating can be calculated for each mode using Eq 5.11, with COMSOL used to calculate the strain energies in the coating and substrate. Due to the complexities associated with modelling multiple concentric thin interfaces in FEA, the *Full Stack* coating was modelled as a thickness weighted average of the *Upper* and *Lower* stack coatings. If the material properties of each coating layer are known, the mean Young's Modulus of a coating layer can be expressed as the weighted sum of its parts: -

$$Y_{\text{coating}} = \sum_{i=1}^{n} \frac{Y_i t_i}{t_i}, \tag{6.1}$$

this equation states that the Y_{coating} is the average Youngs' modulus of the coating layer where Y_m and t are the Youngs' modulus and thickness of a single layer summed over n layers. Likewise, the composite Poisson ratio can be calculated by the following [224]:

$$\nu_{\text{coating}} = \frac{(t_i Y_i \nu_i (1-\nu_i^2)) + (t_j Y_j \nu_j (1-\nu_j^2))}{(t_i Y_i (1-\nu_j^2)) + (t_j Y_j (1-\nu_i^2))}. \tag{6.2}$$

Subscripts i and j denote the average values of the *Upper Stack* and *Lower stack* coatings respectively. Y_{coating} and ν_{coating} can then be used with the total coating thickness to produce a model of the *Full Stack* coating, which can be solved in FEA. The composite thickness of the each stack was calculated using the following formulation:

$$t_{\text{coating}} = \sum_{i=1}^{n} \frac{t_i}{n_i} N_i, \tag{6.3}$$

where t_i is the physical thickness of each layer, n_i is the refractive index at 2000 nm and N_i is the total number of layers of material i. This simplification vastly decreases the computational power required to calculate the energy stored in the coating layers. However it should be noted that this simplification does not allow for

interface dependent interactions between coating layers. Using this model, the respective elastic strain energy density in the average coating layer and substrate can be simulated using FEA. The values used in these models are detailed in Tab. 6.2.

	Full Stack	*Lower Stack*	*Upper Stack*
Composite Young's Modulus (Pa)	9.6×10^{10}	9.3×10^{10}	1.0×10^{10}
Composite Poisson ratio	0.2006	0.1967	0.2038
Composite thickness (nm)	5321.71	2406.35	2915.36

Table 6.2: Average properties of the *Full Stack*, *Lower Stack*, and *Upper Stack*, coatings calculated from values detailed in Tab 6.1 using equations 6.1 and 6.2.

6.0.2 Upper Stack Coating - RLVIP Ta_2O_5-SiO_2

Fig 6.3 shows the measured mechanical losses as a function of frequency of silica Disk 5 before and after coating with the *Upper Stack*. Disk 5 was used as its level of loss provides adequate resolution to the expected level of loss produced by the layers of Ta_2O_5-SiO_2. Each point is produced from three consecutive suspensions of the sample, where the lowest mechanical loss and its respective error are then shown in Fig 6.3. By calculating the ratio of elastic strain energy density stored in the substrate and coating layers ($\frac{E_{\text{coat}}}{E_{\text{sub}}}$) the loss of the coating layers can be extracted using Eq. 5.11. $\frac{E_{\text{coat}}}{E_{\text{sub}}}$ for each frequency is detailed in Tab. 6.3.

In the as-deposited state, the *Upper Stack* loss was between 3.9×10^{-4} and 4.7×10^{-4}. The minimum recorded value found at $f_o = 2.7\,\text{kHz}$. As the measurement frequency increases, the calculated coating loss also increases following the linear function ($\phi(f) = \phi_0 + mf$), $\phi(f) = 3.89 \times 10^{-4} + f(2.27 \times 10^{-9})$. These values are 60% lower than single layer RLVIP Ta_2O_5 measured on SiO_2 cantilevers measured by the author, it is expected that the additional layers of SiO_2 are responsible for this change. It can also be speculated that the intra-chamber temperature during deposition could have had an effect on the coating material, further reducing its loss.

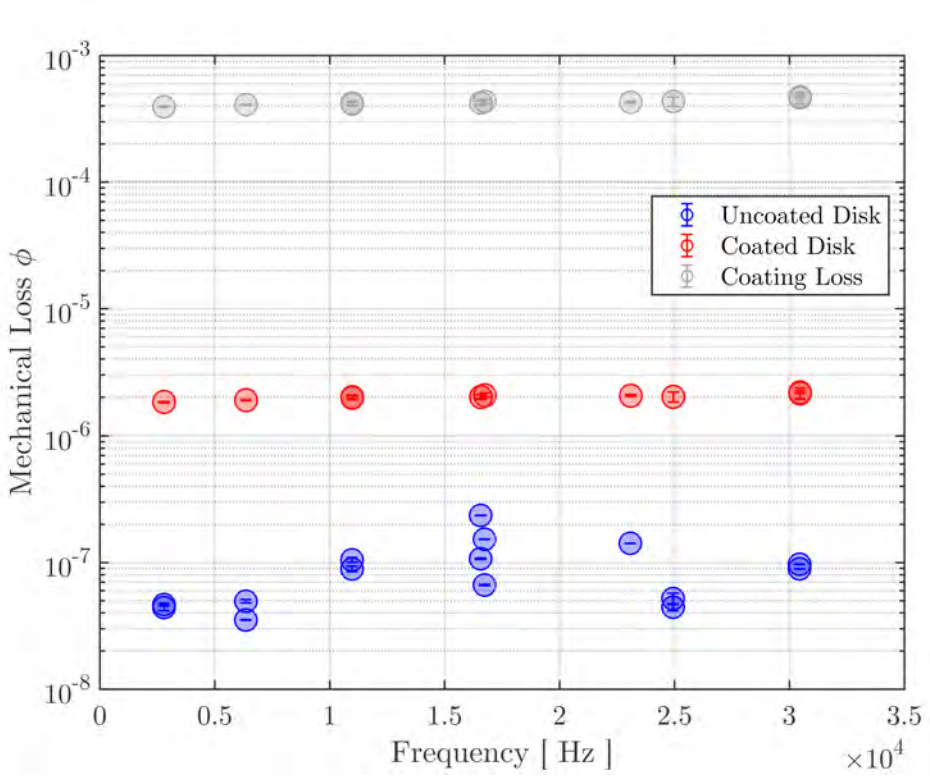

Figure 6.3: Comparison of mechanical losses measured on the uncoated silica disk before and after deposition of the *Upper Stack* coating with the calculated coating loss contributions.

It is assumed that the thickness of each coating layer does not change with heat treatment temperature allowing ϕ_{coating} to be calculated for each heat treated state of the disk using the same values of $\frac{E_{textsub}}{E_{textcoat}}$.

Conforming with the heat treatments of optical absorption samples (see Chapter 4) the *Upper Stack* sample was heat treated in 100 °C increments, starting at 300 °C. This temperature was chosen as this is above the reported 'pre-annealing' temperature of the coatings during coating deposition. At each heat treatment temperature the coated sample was suspended a minimum of three times, sampling the ringdown of each mechanical resonance multiple times at each stage. Fig 6.4 shows ϕ_{coating} for each heat treatment temperature.

Frequency (Hz)	$\frac{E_{\text{sub}}}{E_{\text{coat}}}$
2790.3	292.68
6355.4	293.38
10974.6	294.84
16577.2	296.94
16749.6	291.44
23094.9	299.52
24952.3	295.84
30459.1	294.71

Table 6.3: Measured resonant frequencies of the *Upper Stack* coated silica disk and the calculated ratio of energy stored in the coating and substrate from COMSOL.

After heat treatment at 300 °C, the coating loss for every shear dominated ($m \geq 2, n = 0$) mode was measured to have lower loss, with bulk dominated ($m \geq 1, n = 0$) modes showing a change of ∼4%. Again increasing the heat treatment temperature to 400 °C, ϕ_{coating} decreased across the measured frequency range. A linear dependence with frequency is still observed, with losses measured at 30.4 kHz showing the smallest change. At 500 °C the coating loss decreased further, where a linear relationship with frequency can still be assumed. The final heat treatment of this coating was carried out at 550 °C, below the expected crystallisation temperature for RLVIP Ta_2O_5 (see chapter 4). At this temperature modes, below 16.8 kHz the loss decreases by ∼5%, above 22 kHz it increases by ∼7%, and no longer follows a linear frequency dependence. A full summary of the average ϕ_{coating} is detailed in Tab. 6.4.

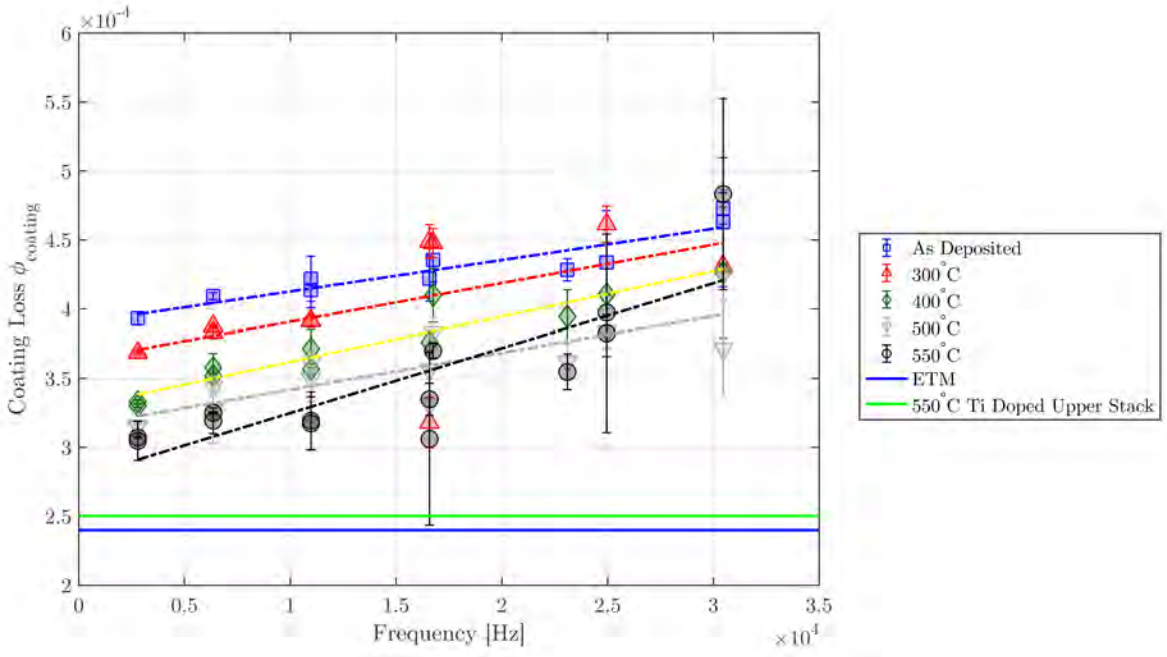

Figure 6.4: Calculated lowest coating loss values for each set of suspensions of the *Upper Stack* disk as a function of heat treatment temperature. Linear fits of ϕ_{coating} as a function of frequency are shown for each heat treatment. Error bars are denoted as the associated standard deviation in coating loss for that measured mode.

Heat Treatment (°C)	ϕ_0	m
As Deposited	3.90×10^{-4}	2.28×10^{-9}
300	3.63×10^{-4}	2.79×10^{-9}
400	3.29×10^{-4}	3.27×10^{-9}
500	3.15×10^{-4}	2.66×10^{-9}
550	2.78×10^{-4}	4.69×10^{-9}

Table 6.4: Coefficients ϕ_0 and m extracted from linear fits of $\phi_{\text{coating}}(f)$ for the *Upper Stack* coating for each heat treatment.

An increase in coating loss at higher frequencies when the samples were heat treated to higher temperatures could be a precursor to material crystallisation, however heat treatments above 550 °C were not tested. Measurements by A. Amato found that an IBS Ti:Ta_2O_5—SiO_2 aLIGO coating of broadly equivalent reflectivity to be 24.5% lower in loss [217]. It is expected that by doping the Ta_2O_5 layers with titania the coating loss of these layers could be reduced by 23%, as occurs with IBS

Ta_2O_5, allowing the RLVIP material to potentially have the same performance and the loss requirements of an ETM coating could be achieved.

These measurements of the coating loss can then be used to predict the magnitude of thermal noise it would produced in a gravitational wave detector. Thermal motion of a coating layer sensed by GW detector are dependent on the dominant mechanism of motion used by the material to dissipate vibrational energy [132]. The total coating loss ϕ_{total} measured can be decomposed into two components associated with bulk and shear energy, ϕ_{bulk} and ϕ_{shear} dissipated by the coating layer. As these values cannot be directly measured using the ringdown method, they must be inferred from a combination of total coating loss $\phi^{BS}_{\text{coating}}$ and the energy stored in bulk and shear motion. Standard minimisation techniques exist where a global minimisation of two unknown values be performed in order to calculate the 'best fit' values of ϕ_{bulk} and ϕ_{shear} to a model of the total coating loss:

$$\phi^{BS}_{\text{coating}} = \left(\frac{E_{\text{bulk}}}{E_{\text{total}}}\right)\phi_{\text{bulk}} + \left(\frac{E_{\text{shear}}}{E_{\text{total}}}\right)\phi_{\text{shear}}. \tag{6.4}$$

An example of the predictions for ϕ_{bulk} and ϕ_{shear} of the as-deposited *Upper Stack* coating is shown in Fig 6.5. Values marked with '△' are solutions of Eq. 6.4 which place wide bounds on possible values of ϕ_{bulk} and ϕ_{shear} between 1×10^{-9} and 0.1, which are further constrained by respective errors on each measured coating loss. While this minimisation produces a fit within the error bounds of the measured data, the wide parameter space (1×10^{-9} - 0.1) over which the minimisation is run does not allow the errors of the estimated values for ϕ_{bulk} and ϕ_{shear} to be extracted.

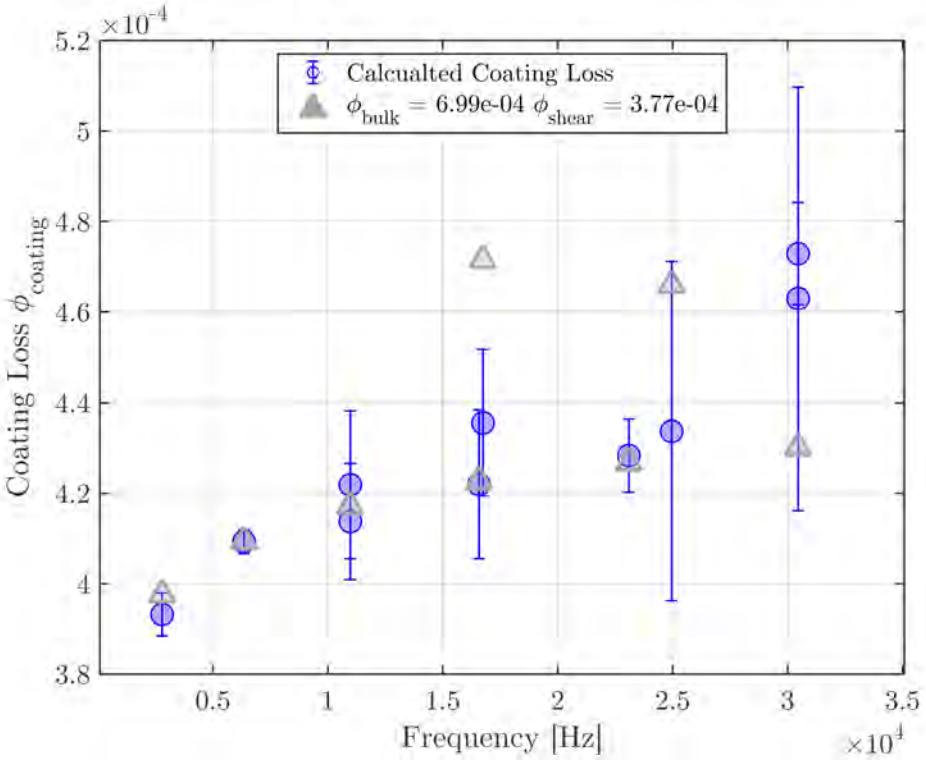

Figure 6.5: Coating loss of the *Upper Stack* coating in the as-deposited state compared to minimised solutions of ϕ_{bulk} and ϕ_{shear}.

In order to place confidence bounds on the generated values, the minimisation was run at the upper and lower error bounds of the measured loss values. This then allows the spread in predicted values of ϕ_{bulk} and ϕ_{shear} as a result of these errors to be estimated. An example of this method is shown in Fig 6.6 where the same measured data is compared to values of $\phi_{coating}$ calculated with the minimised values of ϕ_{bulk} and ϕ_{shear} in each case. A third comparison case is also produced by taking alternating sets of upper and lower bounds for each frequency.

Performing these minimisations produces four different estimated for ϕ_{bulk} and ϕ_{shear} for a given set of measured data allowing the mean and spread of these values to be calculated. It is interesting to note that the values of ϕ_{bulk} and ϕ_{shear} produced using the upper and lower bounds of the measured data do not immediately reflect those values, but instead follows a slightly different trend. In each case the values of $\phi_{coating}$ created using alternating error bounds lies between the values predicted using the upper and lower extremes.

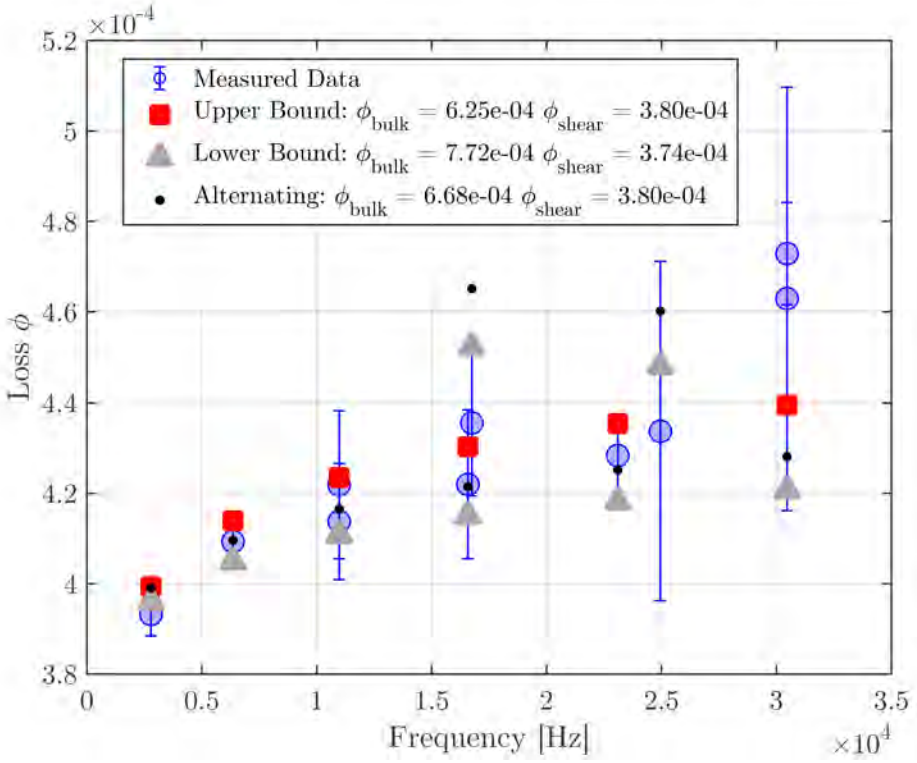

Figure 6.6: As deposited coating loss of *Upper Stack* coating in compared to minimised solutions of ϕ_{bulk} and ϕ_{shear} produced using upper, lower and alternating bounds of the measured ϕ_{coating}.

Parameter	Wide Bounded	Lower	Upper	Alternating
ϕ_{bulk}	7.0×10^{-4}	7.7×10^{-4}	6.3×10^{-4}	6.7×10^{-4}
ϕ_{shear}	3.8×10^{-4}	3.7×10^{-4}	3.8×10^{-4}	3.8×10^{-4}

Table 6.5: Comparison of ϕ_{bulk} and ϕ_{shear} values for the *Upper Stack* coating (as-deposited) calculated from each minimisation.

Following Eq. 6.4, ϕ_{coating} is a function of two loss components and the ratio of energies stored in bulk and shear motion. For cylindrical samples this ratio is weighted towards shear motion for all mode shapes (see section 5.1.1) and as such should have the largest influence on ϕ_{coating}.

Fig 6.7 details how each component of coating loss contributes to the total measured loss of the sample. This allows the thermal motion of the coating in a GW detector where the predicted mean values of ϕ_{bulk} and ϕ_{shear} have been scaled by the ratio of E_{bulk} and E_{shear} stored in the coating for each mode frequency, as

dictated by Eq. 5.11. Applying this correction to include the total energy for each mechanism, each loss component shows that the contributions of shear motion is higher than dissipation from bulk motion for each resonance of the coated sample. This contrary to what is expected at first glance using the values of ϕ_{shear} and ϕ_{bulk} shown in Tab. 5.4.

The dependence of ϕ_{shear} and ϕ_{bulk} on mechanical mode frequency follows the dependence of E_{bulk} and E_{shear} (see section 5.1).

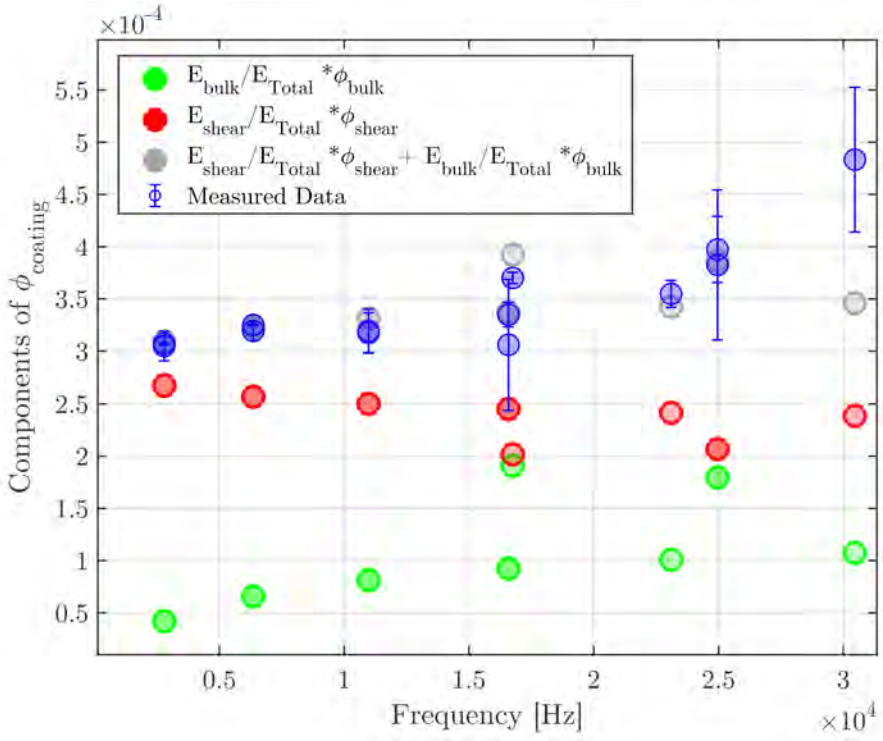

Figure 6.7: Values of ϕ_{bulk} and ϕ_{shear} scaled by the fraction of energy stored in each type of motion for the *Upper Stack* coating in the as-deposited state.

This analysis was repeated for all heat-treatment steps. The results are summarised in Fig 6.8, where the calculated ϕ_{bulk} and ϕ_{shear} values are compared to the trend in measured ϕ_{coating}. A decreasing trend in ϕ_{shear} is observed with heat treatment, moving from $(3.8 \pm 0.07) \times 10^{-4}$ as-deposited, to a minimum of $(2.9 \pm 0.02) \times 10^{-4}$ at 550 °C closely following the trend of ϕ_{coating}. ϕ_{bulk} increases after the first heat-treatment step, but also decreases with successive heat treatment steps, reaching a value of $(6.5 \pm 0.37) \times 10^{-4}$ between 500 °C and 550 °C. The apparent increase in bulk loss at 300 °C could possibly be explained by the lower

number of resonant frequencies which were measured on the sample in this state, worsening the final estimated values.

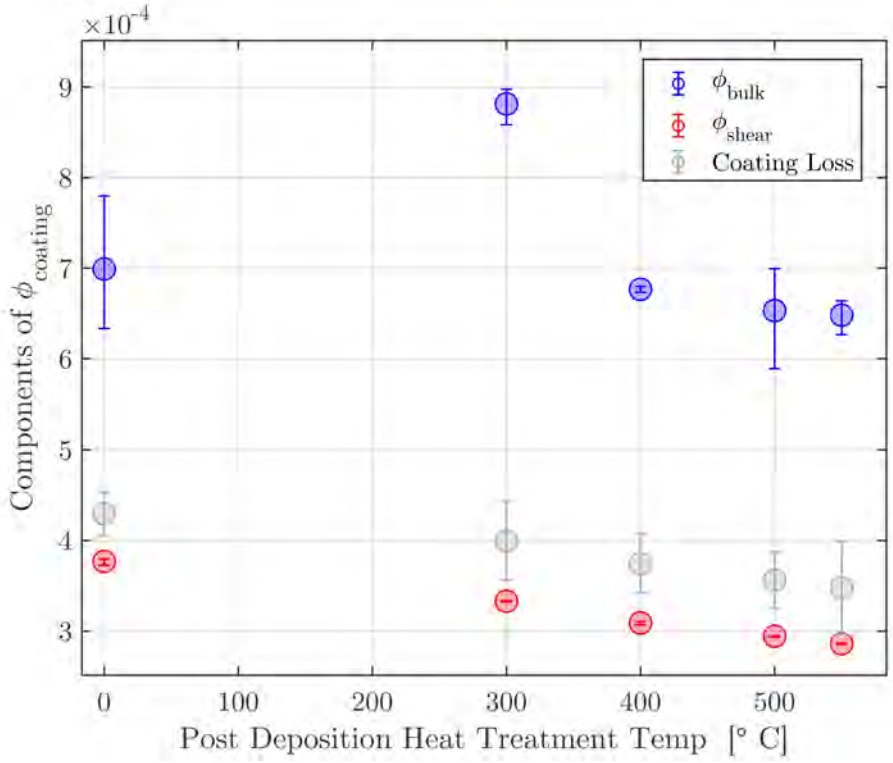

Figure 6.8: Calculated values of ϕ_{bulk} and ϕ_{shear} for different heat treated states of the *Upper Stack* coating. Error bars denote the standard deviation of values produced by different minimisations.

By applying the same scaling of ϕ_{bulk} and ϕ_{shear} to the strain energy stored in the coating the contribution ϕ_{bulk} and ϕ_{shear} to the coating loss as a function of heat treatment can be shown. In Fig 6.9 the scaled components of ϕ_{bulk} and ϕ_{shear} are shown as a function of heat treatment temperature where the respective errors are calculated from the standard deviation of each scaled loss component for all measured frequencies. All values for ϕ_{bulk} and ϕ_{shear} are detailed in Tab. 6.6 In this regime, as in Fig 6.7, the loss of the *Upper Stack* coating is dominated by ϕ_{shear} for each heat treatment.

Heat Treatment [°C]	Parameter	Wide Bounded	Lower	Upper	Alternating
As Deposited	ϕ_{bulk}	7.0×10^{-4}	7.7×10^{-4}	6.3×10^{-4}	6.7×10^{-4}
As Deposited	ϕ_{shear}	3.8×10^{-4}	3.7×10^{-4}	3.8×10^{-4}	3.8×10^{-4}
300	ϕ_{bulk}	8.8×10^{-4}	9.0×10^{-4}	8.6×10^{-4}	9.0×10^{-4}
300	ϕ_{shear}	3.3×10^{-4}	3.3×10^{-4}	3.3×10^{-4}	3.3×10^{-4}
400	ϕ_{bulk}	6.8×10^{-4}	6.7×10^{-4}	6.8×10^{-4}	6.7×10^{-4}
400	ϕ_{shear}	3.1×10^{-4}	3.1×10^{-4}	3.1×10^{-4}	3.1×10^{-4}
500	ϕ_{bulk}	6.5×10^{-4}	7.1×10^{-4}	6.0×10^{-4}	6.9×10^{-4}
500	ϕ_{shear}	2.9×10^{-4}	2.9×10^{-4}	2.9×10^{-4}	2.9×10^{-4}
550	ϕ_{bulk}	6.5×10^{-4}	6.7×10^{-4}	6.3×10^{-4}	6.6×10^{-4}
550	ϕ_{shear}	2.9×10^{-4}	2.9×10^{-4}	2.9×10^{-4}	2.9×10^{-4}

Table 6.6: Comparison of ϕ_{bulk} and ϕ_{shear} values for the *Upper Stack* coating for each heat treatment, calculated from each minimisation.

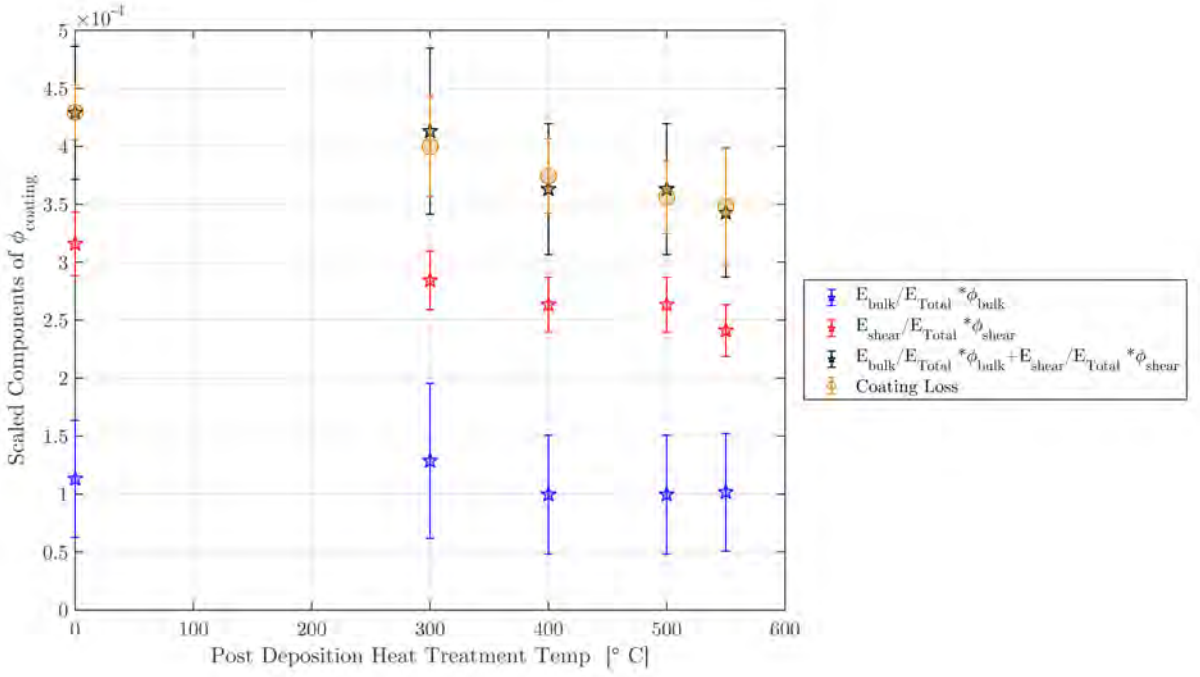

Figure 6.9: Average values ϕ_{bulk} and ϕ_{shear} scaled by the fraction of energy stored in each type of motion for the *Upper Stack* coating for each heat treatment temperature. Error bars are the standard deviation of the scaled contributions for each resonant frequency.

While the contributions of shear loss decreases as a function of heat treatment temperature, the level of bulk loss remains constant within its respective errors. By taking the product of the bulk and shear loss components the decreasing trend of calculated coating loss is produced. This result shows that minimisation method used to calculate ϕ_{bulk} and ϕ_{shear} is accurate, and suggests that heat treatment of the *Upper Stack* primarily reduces shear mechanical dissipation.

6.0.3 Lower Stack Coating - RLVIP aSi-SiO_2

As the loss of single layers of RLVIP SiO_2 (measured by P. G. Murray on cSi cantilevers) and aSi have a similar magnitude of mechanical loss [188] any analysis of changes in loss to the *Lower Stack* coating with heat treatment must consider changes in both materials. Fig 6.11 shows the calculated level of ϕ_{coating} for the *Lower Stack* as a function frequency and heat treatment temperature.

Frequency (Hz)	$\frac{E_{\text{sub}}}{E_{\text{coat}}}$
2803.3	306.41
6384.1	307.23
11022.9	308.86
16648.5	311.14
16813.2	305.44
23188.5	313.92
25044.7	310.16
30580.1	308.99

Table 6.7: Measured resonant frequencies of the *Lower Stack* coated silica disk and the calculated ratio of energy stored in the coating and substrate from COMSOL.

After heat treatment at 300 °C a $\approx 4\%$ reduction in coating loss was observed, with all modes <20 kHz following this trend. Modes at higher frequencies showed a slight increase in ϕ_{coating}. This increase in loss is a consequence of the non-linear minimisation of beating decays, where the second loss component of the 25.05 kHz and 30.5 kHz mode pairs show much higher levels of coupling between modes compared to lower frequency modes. Heat treating the sample at a 400 °C shows a

larger improvement in coating loss, with approximately 12% reduction for all modes, save the calculated coating loss at 25.05 kHz showing a 3.44% reduction. A 500 °C heat treatment produced a decrease of 22.5% in ϕ_{coating} across the measured frequency range. This heat-treatment had a larger effect on the loss of the higher frequency modes (>20 kHz), previously the effect of heat treatment on these modes has less of an effect.

The mechanical loss of the SiO_2 substrate can be presumed largely unaffected by heating as these samples were heated to 950 °C for four hours before the uncoated loss measurements. Comparing this trend with single layers of aSi measured on silica cantilevers by M. Fletcher [203], the same frequency dependence can be observed. This observation potentially hints at an underlying loss mechanism inside the layers of aSi which are largely unaffected by heat treatments below 500 °C.

Comparing the total reduction in loss between the sample in the as-deposited state to the sample after heat treatment at 500 °C the average mechanical loss has been lowered by 35%, which is not the same level of reduction shown from measurements of single-layer aSi [188], where more than a 70% reduction in loss was observed. Further heat tenements of the sample at 550 °C, higher than the optimum temperature for aSi, caused $\phi(f)$ to increase by 36%. A full summary of the average coating loss and frequency dependence for the *Lower Stack* coating is detailed in Tab. 6.8.

Cryogenic mechanical loss measurements of 1 µm thick RLVIP aSi carried out by R.Robie [225], show an increase in loss at 300 K, with similar measurements carried out on SiO_2 showing little contributions to ϕ_{coating} at these temperatures. This suggests that with heat treatment the aSi loss peak is reduced at temperatures below 550 °C. Studies have been carried out by J. Steinlechner *et al.* [188] which suggest that the optical absorption of single layer aSi coatings when heat treated between 350 °C - 600 °C can be explained by the change in the materials amorphous structure. Heat treatments at these temperatures can cause the formation of micro-crystals in the coating layer. It could then be inferred that if micro-crystallite structures were produced in the coating layers, this could produce a higher measured ϕ_{coating}.

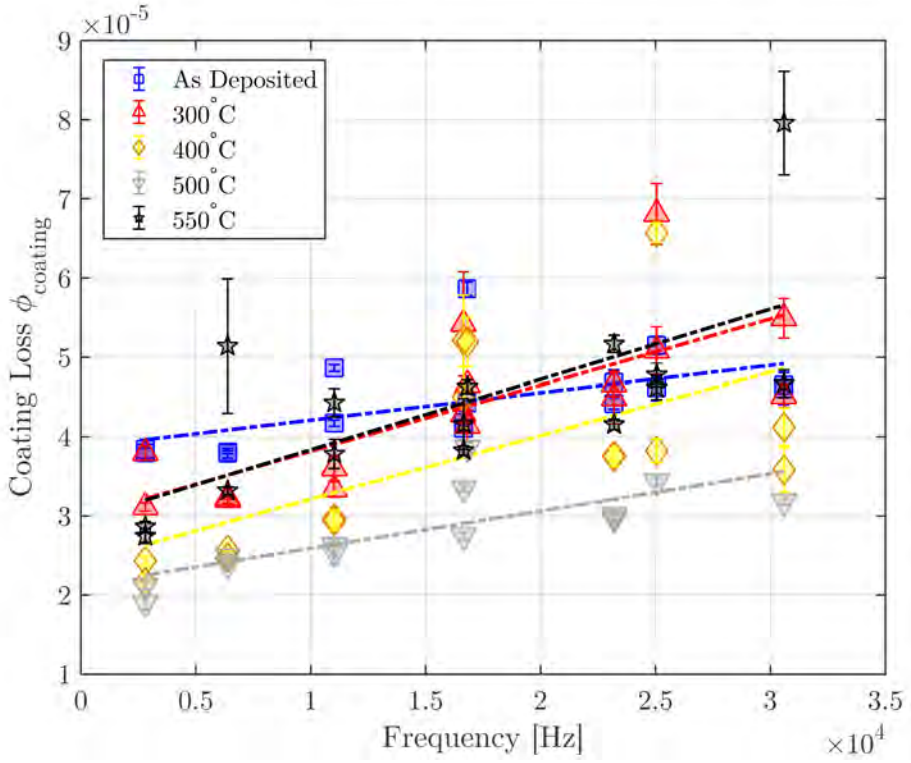

Figure 6.10: Calculated lowest coating loss values for each set of suspensions of the *Lower Stack* disk as a function of heat treatment temperature. Linear fits of ϕ_{coating} as a function of frequency are shown for each heat treatment. Error bars are denoted as the associated standard deviation in coating loss for that measured mode.

Referring to measurements of optical absorption single layer RLVIP aSi taken by J. Steinlechner [188], it was found that heat treatment above 500 °C caused the optical absorption of the coating layers to increase by ≥20% but did not report any changes in Raman spectra that would indicate that the coating had fully crystallised. Given that changes from an amorphous coating structure to a partially crystallised structure has been linked with an increase in mechanical loss [226], it would be expected that the presence of micro-crystallites in the coating layers would be noted at temperatures above 400 °C if the mechanical loss of the aSi layers was not affected by the surrounding SiO_2 layers. However, in another paper [149] where the single layer aSi and an SiO_2—aSi bi-layers were deposited by IBS, both coating materials show a similar increase in absorption at 1550 nm at temperatures above 400 °C. If we compare the changes in optical absorption from these measurements detailed in Chapter 4 to the changes in ϕ_{coating} measured on the *Lower Stack* coating, there is a suggested correlation that the formation of nano-crystallises in

the material could explain the increase in ϕ_{coating} measured at 550 °C.

Heat Treatment (°C)	ϕ_0	m
As Deposited	3.87×10^{-5}	3.4×10^{-10}
300	2.98×10^{-5}	8.3×10^{-10}
400	2.41×10^{-5}	8.0×10^{-10}
500	2.12×10^{-5}	4.7×10^{-10}
550	2.95×10^{-5}	8.9×10^{-10}

Table 6.8: Coefficients ϕ_0 and m extracted from linear fits of $\phi_{\text{coating}}(f)$ for the *Lower Stack* coating for each heat treatment.

Any changes in ϕ_{bulk} or ϕ_{shear} as a function of heat treatment may then also reveal possible changes which corroborate this hypothesis. The method used in the previous section (see Fig 6.11) is again employed, comparing ϕ_{coating} of the *Lower Stack* after 400 °C heat treatment with values predicted by minimising Eq. 6.4 to the spread in the measured data.

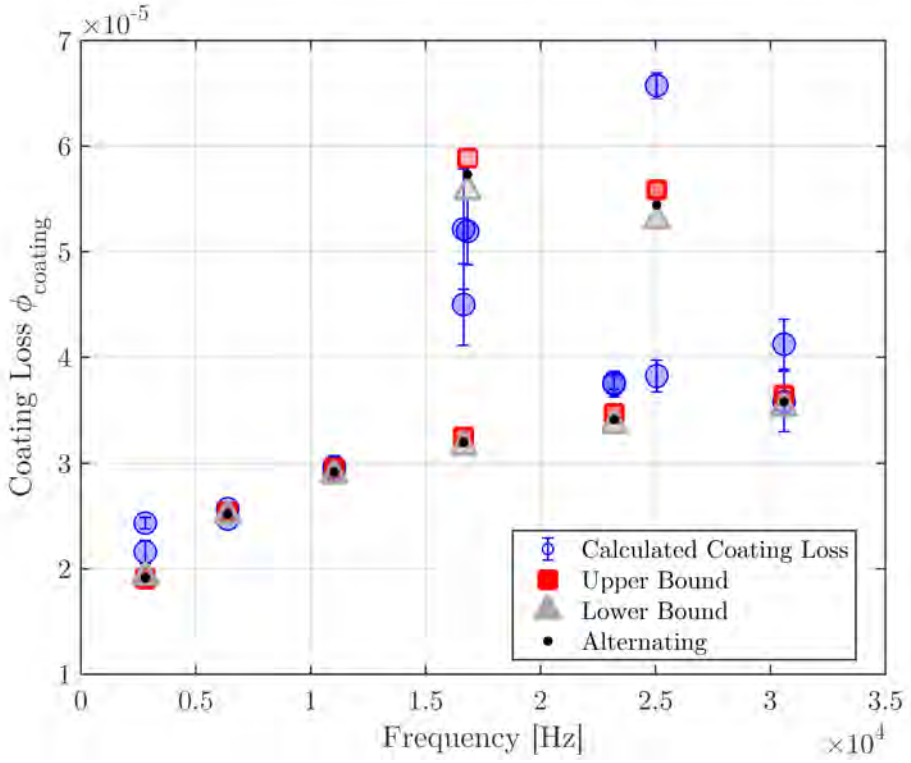

Figure 6.11: Coating loss of the *Lower Stack* coating heat treated to 400 °C, compared to minimised solutions of ϕ_{bulk} and ϕ_{shear} produced using upper; lower and alternating bounds of the measured ϕ_{coating}.

The observations of $\phi(f)$ of the *Lower Stack* follows a similar increasing trend as a function of frequency, with modes at 16.6 kHz, 16.8 kHz and 25.1 kHz showing higher levels coupling with their respective mode pairs. It is thought that this then distorts predictions at higher frequencies, over predicting losses at 25 kHz and an under prediction of those at 30 kHz. Fig 6.12 then compares the average values of ϕ_{bulk} and ϕ_{shear} for each heat treatment state of the *Lower Stack* sample with the average calculated coating loss. For transparency the level of ϕ_{bulk} and ϕ_{shear} predicted by each minimisation is also detailed in Tab 6.9.

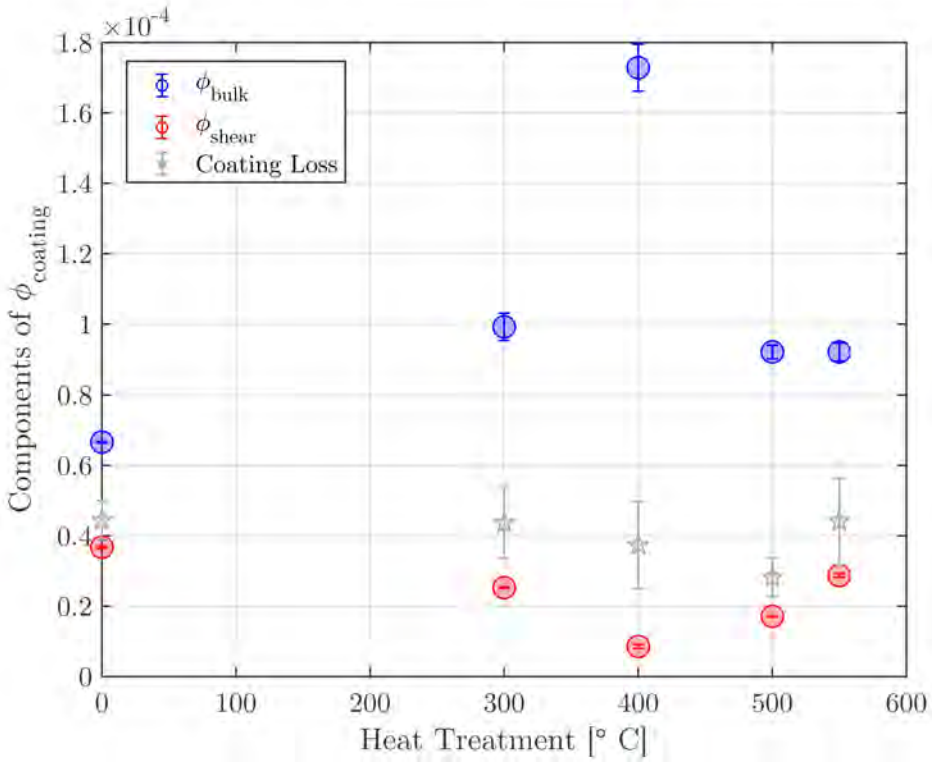

Figure 6.12: Calculated values of ϕ_{bulk} and ϕ_{shear} for different heat treated states of the *Lower Stack* coating. Error bars denote the standard deviation of values produced by different minimisations.

Heat Treatment [°C]	Parameter	Lower	Upper	Alternating	Wide Bounded	STD
As Deposited	ϕ_{bulk}	6.7×10^{-5}	6.6×10^{-5}	6.7×10^{-5}	6.7×10^{-5}	1.9×10^{-7}
As Deposited	ϕ_{shear}	3.7×10^{-5}	3.6×10^{-5}	3.7×10^{-5}	3.7×10^{-5}	3.5×10^{-7}
300	ϕ_{bulk}	1.0×10^{-4}	9.5×10^{-5}	1.0×10^{-4}	9.9×10^{-5}	3.9×10^{-6}
300	ϕ_{shear}	2.5×10^{-5}	2.5×10^{-5}	2.5×10^{-5}	2.5×10^{-5}	1.6×10^{-7}
400	ϕ_{bulk}	1.8×10^{-4}	1.7×10^{-4}	1.7×10^{-4}	1.7×10^{-4}	6.7×10^{-6}
400	ϕ_{shear}	7.9×10^{-6}	9.1×10^{-6}	8.5×10^{-6}	8.5×10^{-6}	5.7×10^{-7}
500	ϕ_{bulk}	9.4×10^{-5}	9.0×10^{-5}	9.2×10^{-5}	9.2×10^{-5}	1.9×10^{-6}
500	ϕ_{shear}	1.7×10^{-5}	1.7×10^{-5}	1.7×10^{-5}	1.7×10^{-5}	6.6×10^{-8}
550	ϕ_{bulk}	1.1×10^{-4}	1.1×10^{-4}	1.1×10^{-4}	9.2×10^{-5}	2.6×10^{-6}
550	ϕ_{shear}	2.6×10^{-5}	2.5×10^{-5}	2.6×10^{-5}	2.9×10^{-5}	5.3×10^{-7}

Table 6.9: Values of ϕ_{bulk} and ϕ_{shear} calculated for each heat treatment of the *Lower Stack* sample using different minimisation bounds.

A decreasing trend in ϕ_{shear} is calculated for the *Lower Stack* coating with heat treatment, moving from $\phi_{\text{shear}} = 3.7\times10^{-5} \pm 6.8\times10^{-7}$ as-deposited to a minimum of $\phi_{\text{shear}} = 8.5\times10^{-6} \pm 1.1\times10^{-6}$ at 400 °C. In direct contrast with this trend, the predicted values of ϕ_{bulk} which shows a large increase over the same heat treatment range. To probe the outlying values bulk and shear loss produced at 400 °C the minimisation was run again, while excluding modes which are known for MMIs, 16.6 kHz, 16.8 kHz and 25.1 kHz. In doing so it was found that this did not change the predictions of ϕ_{bulk} and ϕ_{shear} by more than 1%.

The 100 °C difference between the measured minimum in ϕ_{shear} (400 °C) and ϕ_{coating} (500 °C) is then apparent, as this may suggest a shift in the underlying coating structure which does not emerge in the measured ϕ_{coating} data.

Increasing the heat treatment temperature to 500 °C and 550 °C the measured and predicted values of ϕ_{shear} and ϕ_{coating} increase in kind, while the predicted values of ϕ_{bulk} also decrease, but thereafter are unchanging to within more than 1%. These predictions then suggest that this method of probing ϕ_{bulk} and ϕ_{shear} is more sensitive to changes in ϕ_{bulk} than the latter.

Scaling the predicted values of ϕ_{bulk} and ϕ_{shear} shown in Fig 6.12 by their respective strain energy ratios, the contribution of each loss mechanism to the measured coating loss is produced. The results of these calculations are shown in Fig 6.13. While the scaled contributions of ϕ_{shear} and ϕ_{bulk} are able to reproduce the measured coating loss within experimental error, the predicted variations in each mechanism are much larger than was *Upper Stack* coating (see Fig 6.9). As the values of ϕ_{shear} and ϕ_{bulk} from 6.12 are used, the scaled value of ϕ_{bulk} at 400 °C is again much higher than at other heat treatments.

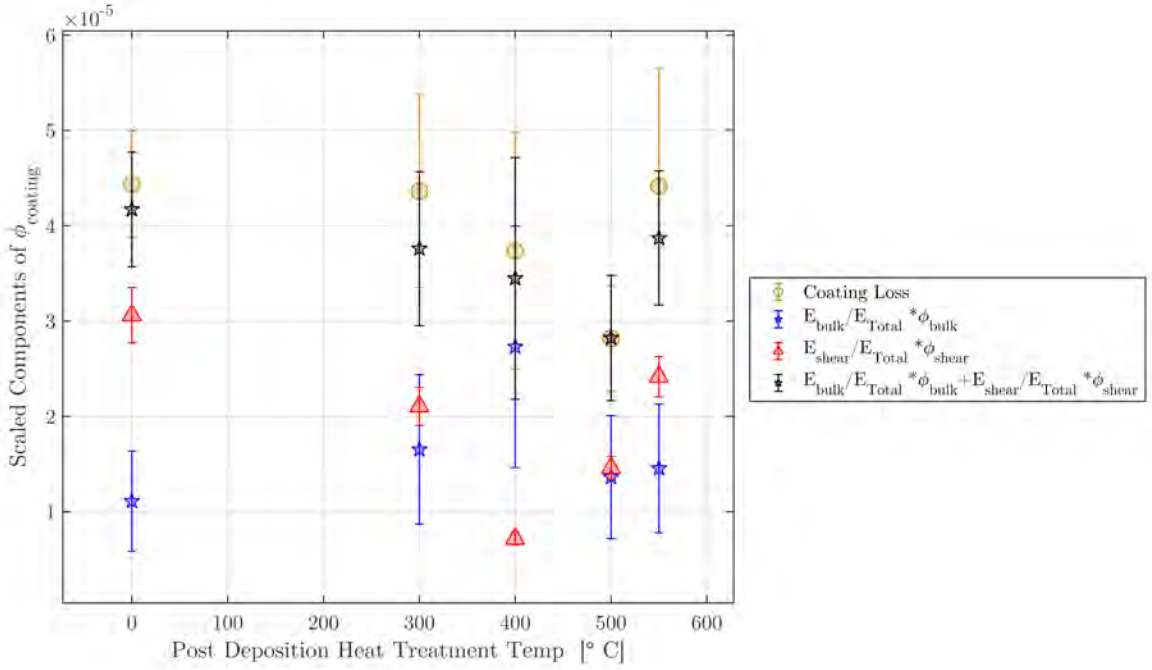

Figure 6.13: Average values ϕ_{bulk} and ϕ_{shear} scaled by the fraction of energy stored in each type of motion for the *Lower Stack* coating for each heat treatment temperature. Error bars are the standard deviation of the scaled contributions for each resonant frequency.

If the material began to crystallise after a heat treatment at 400 °C, the differences between ϕ_{shear} and ϕ_{coating} could corroborate this hypothesis if the resulting crystal structure favoured mechanical dissipation through ϕ_{shear}. However, as the values of both ϕ_{shear} and ϕ_{bulk} at after being heat treated at 400 °C are significantly different from the values calculated after heat treating the sample at different temperatures, it seems unlikely that such a large change in the samples structure could be produced by heat treatment at this temperature. If these spurious points are discounted from the final comparison, there is a clear decrease in the scaled ϕ_{shear} contributions to coating loss as a function of heat treatment temperature, which reaches a minimum between 400 °C and 500 °C, coinciding with the lowest measured coating loss. At 500 °C, the contributions of ϕ_{shear} and ϕ_{bulk} overlap, with the mean value of ϕ_{shear} calculated to be slightly higher than ϕ_{bulk}. While contributions of ϕ_{bulk} appear nominally unchanging as a function heat treatment temperature, the large error bars in its predictions could suggest that the GeNS suspension method is much more sensitive to changes in ϕ_{shear} than in ϕ_{bulk}.

The difference between ϕ_{shear} and ϕ_{bulk} for the low loss coating also suggests that

the mechanical loss of both aSi and SiO_2 are dominated by ϕ_{shear}. This is consistent with work carried out by M. Fletcher whose work suggested that RLVIP aSi was dominated by ϕ_{shear} [203]. However, as the mechanical losses of RLVIP aSi and SiO_2 have similar magnitudes when deposited and heat treated to 500 °C, it is difficult to disentangle the contributions of each material to the measured coating loss.

6.0.4 Full Stack Coating - RLVIP Ta_2O_5-SiO_2/SiO_2-aSi

A blank disk (Disk 3) coated with the *Full Stack* coating was also studied in the same manner as the *Upper* and *Lower* stack coatings. From the measurements of $\phi(f)$ of the *Upper* and *Lower* stack coatings, with progressive heat treatment, predictions of the *Full Stack* coating could be made. As in each of these cases, the higher refractive index material had a larger component of $\phi(f)$ compared to the lower index SiO_2, the effects of each material would dictate the changes of the *Full Stack* coating. It would be expected that its mechanical loss would show a similar frequency dependence to that observed on the *Upper Stack*, its magnitude dictated by the higher loss Ta_2O_5 layers.

Frequency (Hz)	$\frac{E_{\text{sub}}}{E_{\text{coat}}}$
2938	130
6712.9	130
11630	131
17642	132
17680	129
24697	134
26402	132
29595	131

Table 6.10: Measured resonant frequencies of the *Full Stack* coated silica disk and the calculated ratio of energy stored in the coating and substrate from COMSOL.

Fig 6.14 shows the calculated coating loss for the *Full Stack* coating as a function of frequency and heat treatment temperature using the values detailed in Tab. 6.10. These values are compared with the measured loss of an aLIGO ETM coating [217].

In the as-deposited state, a similar frequency dependence in ϕ_{coating} is observed to that of the *Upper Stack* with a frequency dependence coefficient, m of 1.6×10^{-5} (see Tab. 6.11) compared to the 3.3×10^{-6} and 2.2×10^{-5} observed on the *Lower Stack* and *Upper Stack* coatings respectively. It is noted that the frequency dependence of the *Full Stack* coating follows the same trend as the *Upper Stack*, as this dictates that the dominant source of loss in the coating stack is also responsible for its frequency dependence.

After heat treatment at 300 °C, the coating loss for all measured frequencies decreases, with modes susceptible to higher order coupling showing a larger spread in coating loss. The frequency dependence of ϕ_{coating} appears to be dominated by that of the *Upper Stack*, with a m of 2.5×10^{-5} and 1.9×10^{-5} on the *Upper* and *Full Stack* samples respectively.

Heat treatment of the coating stack at 400 °C produced the largest relative change in ϕ_{coating} from any other previous state with an $\sim$ 27% reduction. In this state, there is also a noticeable change in the measured frequency dependence of the ϕ_{coating}. Measurements of the *Full Stack* coating producing an $m = 9.8\times10^{-6}$ values are more comparable with the lower values of $\phi_{\text{coating}}(f)$ measured on the *Lower Stack* (4.2×10^{-6}) and *Upper Stack* (2.1×10^{-5}). This decrease in measured frequency dependence is in large due to the change in loss of modes measured below 16.6 kHz.

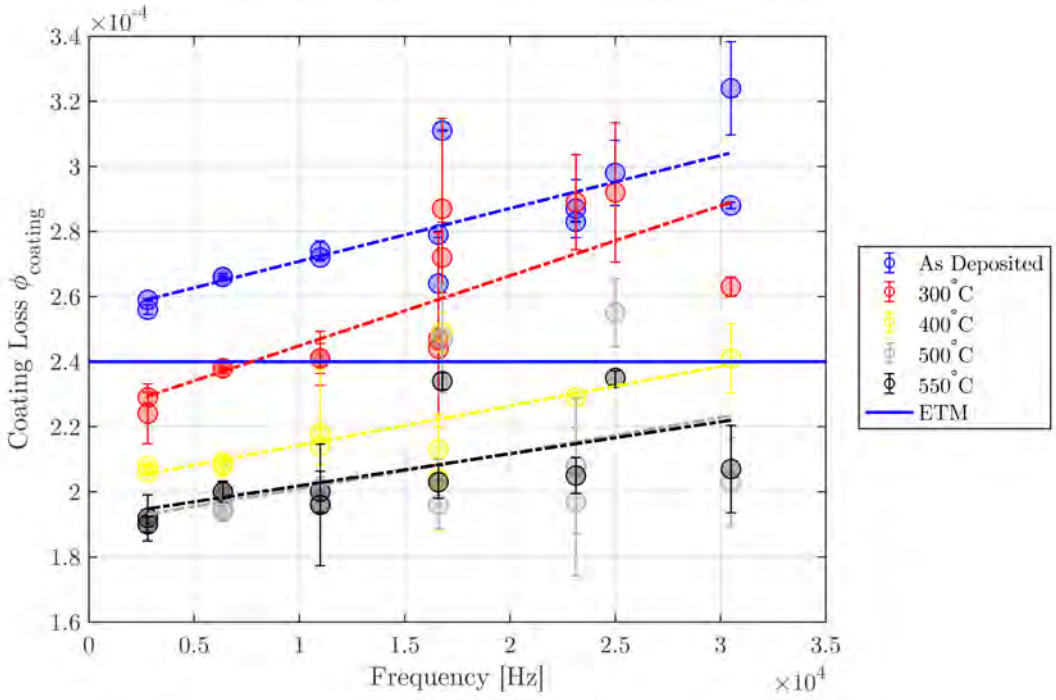

Figure 6.14: Calculated lowest coating loss values for each set of suspensions of the *Full Stack* disk as a function of heat treatment temperature. Error bars are denoted as the associated standard deviation in coating loss for that measured mode.

Heat treating the *Full Stack* coating again to 500 °C further reduced the loss and its dependence on frequency. As at lower heat-treatment temperatures, the lowest loss was measured at 2.7 kHz with this mode giving a loss of 2.6×10^{-4}. Mechanical modes at 16.8 kHz and 23.1 kHz again exhibit larger values of ϕ_{coating} consistent with the expected degenerate coupling behaviour measured on other samples at all heat treatment values. These measured frequencies yield spurious values of ϕ_{coating}, comparable to those from the *Full Stack* coating in the as-deposited state. After heat treating the same at this temperature, all measured values of ϕ_{coating} fall below the measured coating loss values of current ETM coatings [142].

Heat Treatment (°C)	ϕ_0	m
As Deposited	2.55×10^{-4}	1.6×10^{-9}
300	2.23×10^{-4}	2.2×10^{-9}
400	2.02×10^{-4}	1.2×10^{-9}
500	1.90×10^{-4}	1.1×10^{-9}
550	1.92×10^{-4}	9.8×10^{-10}

Table 6.11: Coefficients ϕ_0 and m extracted from linear fits of $\phi_{\text{coating}}(f)$ for the *Full Stack* coating for each heat treatment

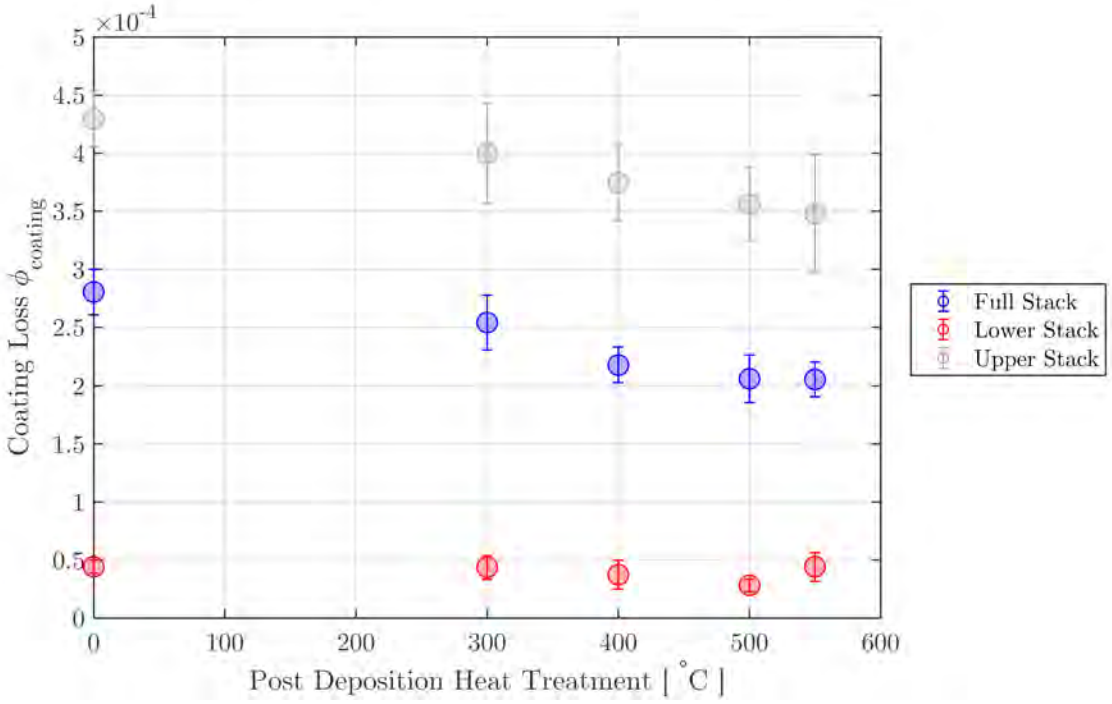

Figure 6.15: Comparison of average coating losses measured on the *Lower Stack*, *Upper Stack* and *Full Stack* coatings as a function of post deposition heat treatment temperature.

Fig 6.15 shows a comparison of the coating loss from the *Full*, *Upper* and *Lower Stack* coatings as a function of heat treatment temperature. This result proves the concept that in the implementation aSi layers in place of Ta_2O_5 layers in the coating stack can reduce the levels of ϕ_{coating} below the level of loss from coatings currently used for aLIGO ETMs. Before heat treating the *Full Stack* to higher temperatures,

the observed trends in the *Upper* and *Lower* stacks were considered. Upon heat treatment the *Upper Stack* coating saw the largest reduction in ϕ_{coating} at 550 °C. However, when the *Lower Stack* coating was subjected to the same temperatures, there was a notable 36% increase in ϕ_{coating}. Given that the layers of aSi in the *Full Stack* have been responsible for the reduction in ϕ_{coating} up to this point, it would therefore be expected that a similar increase would be observed in the loss of the *Full Stack* coating.

Under the assumption that the *Lower Stack* and *Upper Stack* coatings will behave in exactly the same manner when deposited atop one and other, they can be used to calculate the coating loss of the *Full Stack* sample. This can be achieved by calculating a thickness weighted average of the measured loss values at each frequency on the two component stacks

$$\phi_{\text{coating}} = \frac{Y_1 t_1 \phi_1 + Y_2 t_2 \phi_2}{Y_{\text{coating}} t_{\text{coating}}}. \tag{6.5}$$

This equation takes into account the composite Young's moduli', thicknesses and mechanical losses of the two component coatings, where Y_1, Y_2 and Y_{coating} were calculated from the single layer properties of each coating using Eq. 6.1, and t_1, t_2 and t_{coating} are the thicknesses of the *Upper*, *Lower* and *Full Stack* coatings respectively. In Fig 6.16 the predicted coating loss of the *Full Stack* coating is calculated from the *Upper*, *Lower Stack* coating losses at different heat treatment states. All material properties used are detailed in Tab 6.1. These predictions agree with the mean value of the measured coating loss of the Full Stack to within 2%. Mechanical modes which exhibit degenerate coupling at 16.8 kHz and 25 kHz show a much higher loss than the single layer predictions due to this effect. If the difference between degenerate and non-degenerate modes is subtracted from these values, this would increase the accuracy of these predictions. The coating loss measured at 2.8 kHz is also measured at a statistically significant value lower than these calculations predict, but also provides a reasonable match if the upper error bounds of the measured coating loss values are considered.

If we now compare the calculated values of coating loss from the component *Upper Stack* and *Lower Stack* coatings there are noticeable differences between the measured and calculated values. The most striking difference is the change in the predicted frequency dependence of the *Full Stack*, from the two components, with

the coating loss at 30 kHz 31% higher than was measured. It is also interesting to note that at 550 °C measurements of the *Lower Stack* coating showed an increase measured coating loss, indicating that if loss arising from micro-crystallites in the aSi layers is a true artifact in the *Full Stack* coating, its contribution is far smaller than contributions from layers of Ta_2O_5. As was previously shown for the *Upper Stack* coating in Fig 6.4, the frequency dependent coating loss for Ta_2O_5 was much greater than for either of the other two materials and as such provides a greater contribution to these calculated values.

If instead of comparing coating loss values of the *Upper Stack*, *Lower Stack* when both heat treated to 550 °C we compare the 4 combinations of 500 °C and 550 °C measurements to the measured *Full Stack* coating the contributions of each material can be understood. At 500 °C and 550 °C the frequency dependence of the *Upper Stack* coating over predicts the coating loss at frequencies above 20 kHz with the *Lower Stack* coating reducing these predictions. It can be shown that the sum of square error (SSE) between the measured *Full Stack*, a quantitative distance between each prediction and measurement, that closest prediction is produced when each stack is heat treated between 500 °C and 550 °C. At these temperature the loss of the *Lower Stack* was 36% lower than the equivalent loss at 550 °C, with the magnitude of loss on the *Upper Stack* differing on average by ∼2%. At 500 °C the frequency dependence of the *Upper Stack* coating was also reduced by 43%.

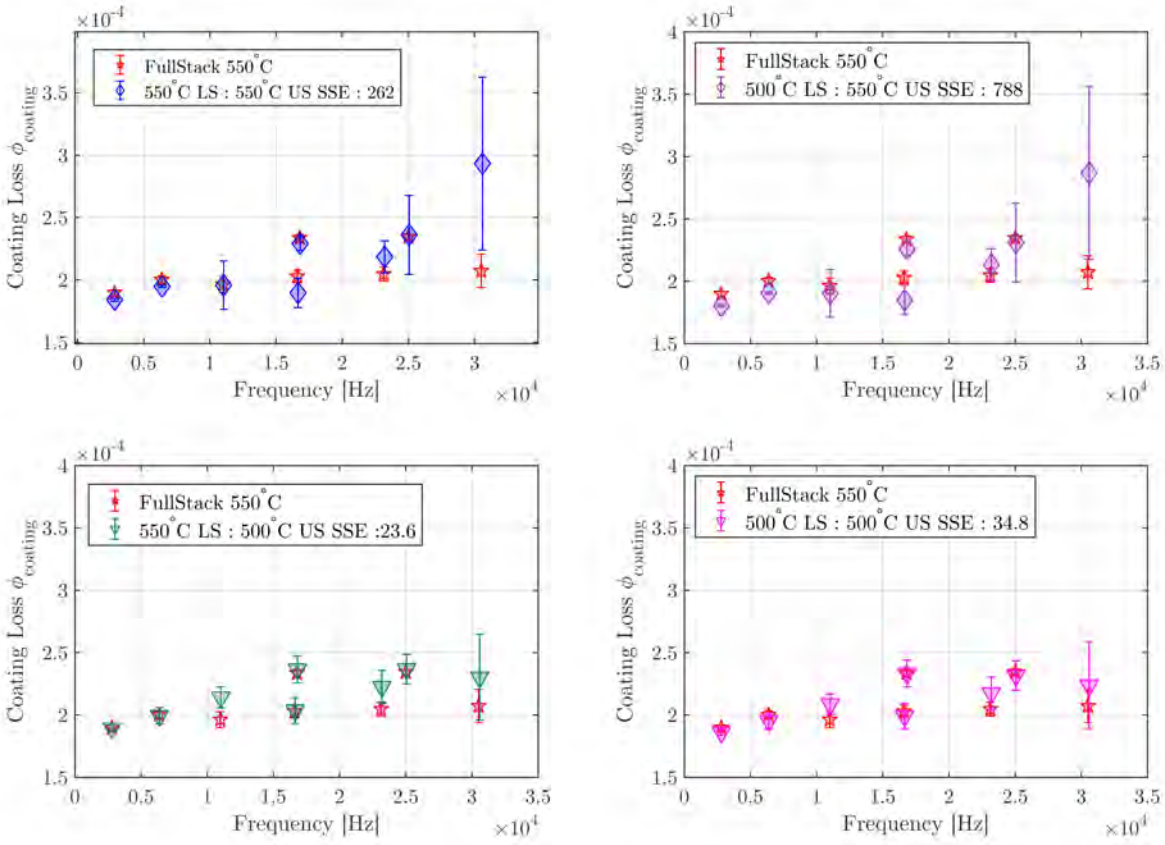

Figure 6.16: Comparison of *Full Stack* coating losses to calculated predictions using lowest single layer losses and summation of *Upper* and *Lower Stack* losses at 500 °C and 550 °C.

As comparing the different heat treatment states of each component stack does not fully reproduce the measured *Full Stack* coating loss values, these results are open to interpretation. However, the differences between measured coating loss on each component stack and the full coating stack suggests that there are structural differences between each combination of materials. Calculating the loss of the *Full Stack* coating heat treated to 550 °C with losses from lower heat treatment temperatures on the two component stacks showed that values which were less frequency dependent, with a lower magnitude of coating loss provide a better depiction of the *Full Stack* coating. Both of these characteristics which increased on the *Upper* and *Lower Stacks* when heat treated to 550 °C suggest that the changes in structure in the aSi and Ta_2O_5 layers responsible for these properties has not occurred in the same manner in the *Full Stack* coating.

This leads to the hypothesis that the deposition of multiple coating layers with different thermal expansion properties imparts a stress on the total coating layer. The imparted thermal stress and the presence of the SiO_2 'boundary' layers increases the energy required for high index layers in the HR stack such as aSi and Ta_2O_5 to change to a lower potential, lower mechanical loss state. As each coating stack is then heat treated, the coating layers in the *Full Stack* coating are more constrained than layers in the *Upper* or *Lower* stacks, inhibiting the formation of micro-crystallites which would otherwise increase the loss of the coating. This hypothesis is consistent with similar investigations into crystallisation suppression of SiO_2—TiO_2 nanolayers, ≈120 nm, thick carried out by H.Pan *et al.* [227]. In this study it was found that by increasing the number of coating layers in a nanolayer stack could suppress the crystallisation of TiO_2 layers by more than 100 °C.

Currently, due to experimental restrictions, experimental evidence to support this hypothesis is not available. It would be expected that measurements of these samples using a non-invasive technique such as Raman spectroscopy or a crystalline structure analysis through x-ray diffraction would show a clear difference in the material structure between the *Upper* and *Lower* stacks when compared to the *Full Stack* coating. Given that cantilever samples have been produced for these coatings (See Chapter 8) these can also be used to determine the coating stress as a function of heat treatment through deflection measurements. These measurements would provide enough information to qualitatively determine if the miss-match in thermal expansion coefficients between coating layers can affect the coating loss of a multi-material sample.

To provide further insight into the loss components of the *Full Stack*, the coating loss was decomposed into its components of ϕ_{bulk} and ϕ_{shear} following the method used in sections 6.0.2 and 6.0.3. Fig 6.17 shows an example for the *Full Stack* when heat treated to 550 °C.

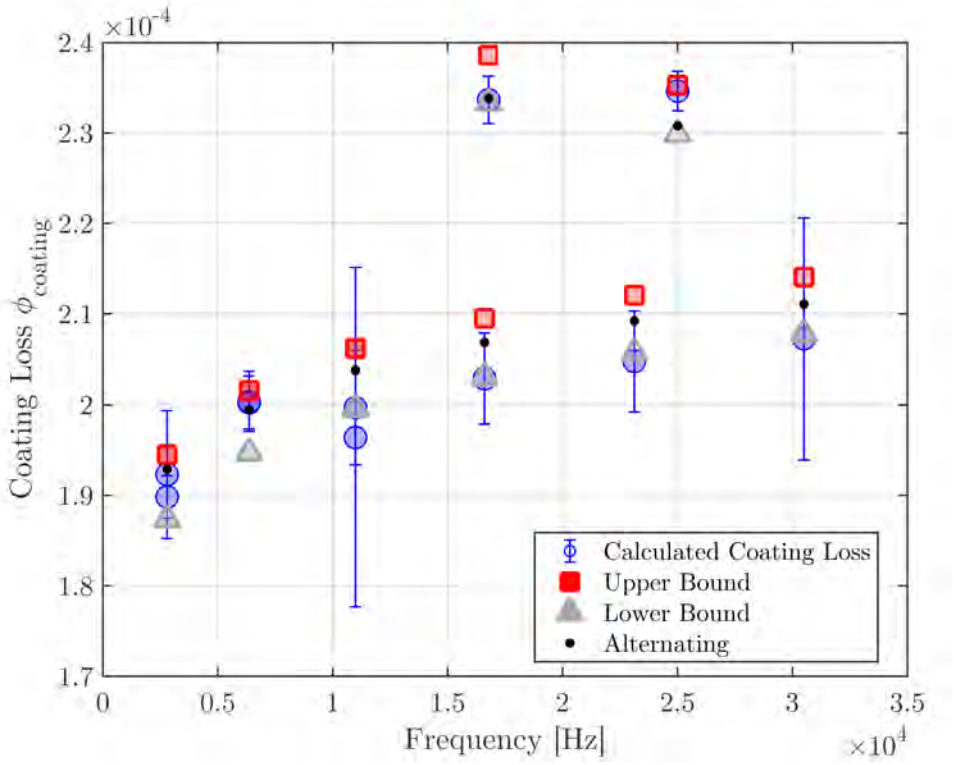

Figure 6.17: Coating loss of the *Full Stack* coating heat treated to 550 °C, compared to minimised solutions of ϕ_{bulk} and ϕ_{shear} produced using upper; lower and alternating bounds of the measured ϕ_{coating}.

Heat Treatment [°C]	Parameter	Lower	Upper	Alternating	Wide Bounded	STD
As Deposited	ϕ_{bulk}	4.7×10^{-4}	4.7×10^{-4}	4.7×10^{-4}	4.7×10^{-4}	8.3×10^{-7}
As Deposited	ϕ_{shear}	2.5×10^{-4}	2.4×10^{-4}	2.5×10^{-4}	2.4×10^{-4}	5.5×10^{-7}
300	ϕ_{bulk}	5.1×10^{-4}	4.5×10^{-4}	4.9×10^{-4}	4.8×10^{-4}	3.1×10^{-5}
300	ϕ_{shear}	2.1×10^{-4}	2.1×10^{-4}	2.1×10^{-4}	2.1×10^{-4}	3.0×10^{-6}
400	ϕ_{bulk}	4.4×10^{-4}	4.4×10^{-4}	4.4×10^{-4}	4.4×10^{-4}	2.1×10^{-6}
400	ϕ_{shear}	1.9×10^{-4}	1.9×10^{-4}	1.9×10^{-4}	1.9×10^{-4}	5.3×10^{-7}
500	ϕ_{bulk}	3.4×10^{-4}	2.9×10^{-4}	3.2×10^{-4}	3.2×10^{-4}	2.6×10^{-5}
500	ϕ_{shear}	1.8×10^{-4}	1.8×10^{-4}	1.8×10^{-4}	1.8×10^{-4}	1.4×10^{-6}
550	ϕ_{bulk}	3.7×10^{-4}	3.7×10^{-4}	3.6×10^{-4}	3.7×10^{-4}	8.4×10^{-6}
550	ϕ_{shear}	1.8×10^{-4}	1.7×10^{-4}	1.8×10^{-4}	1.8×10^{-4}	4.2×10^{-6}

Table 6.12: Values of ϕ_{bulk} and ϕ_{shear} calculated for each heat treatment of the *Full Stack* sample using different minimisation bounds.

Through each heat treatment ϕ_{coating} of the *Full Stack* follows closely with the predicted values of ϕ_{shear} decreasing from $\phi_{\text{shear}} = (2.4 \pm 0.0055) \times 10^{-4}$ as-deposited to a minimum of $\phi_{\text{shear}} = (1.8 \pm 0.0042) \times 10^{-4}$ at 550 °C. The magnitude of ϕ_{bulk} also decreases with each heat treatment with the largest change predicted when the sample was heat treated at 500 °C. Upon closer inspection it was found that by excluding points at 16.6 kHz, 16.8 kHz and 25.1 kHz the predictions of ϕ_{bulk} and ϕ_{shear} did not change by more than 1% and thus cannot fully explain the observed changes. However, such a decrease in ϕ_{bulk} does not produce a coinciding change in ϕ_{shear} in this state, making the result improbable. This result also suggests that like the *Lower* and *Upper* stack coatings, the *Full Stack* material is also dominated by ϕ_{bulk} and could therefore have implications on its level of coating thermal noise in a gravitational wave detector [132].

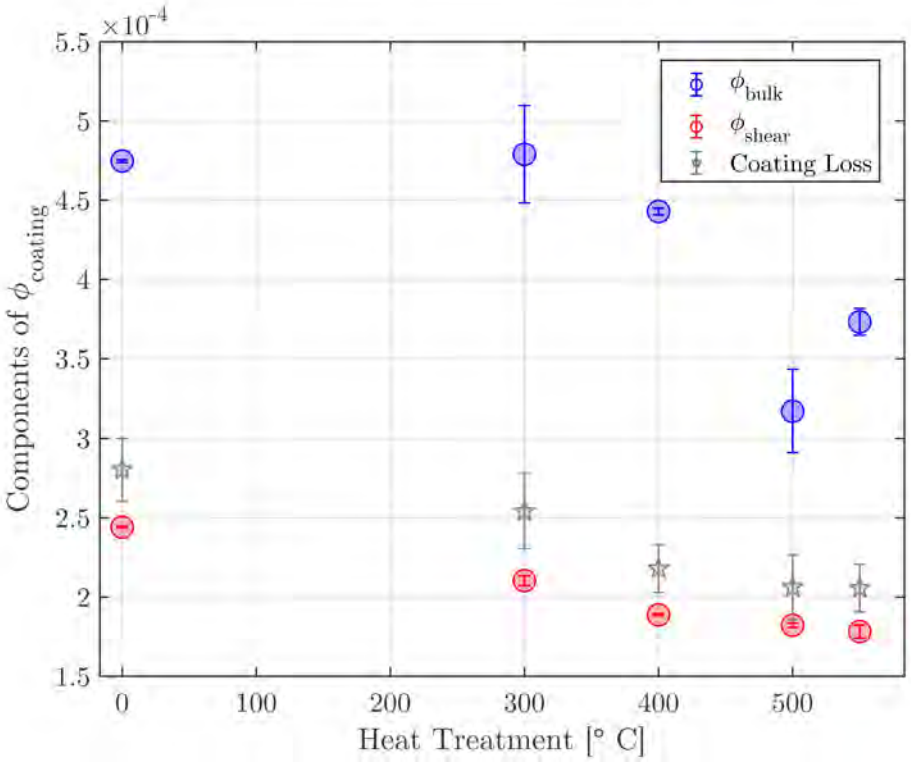

Figure 6.18: Calculated values of ϕ_{bulk} and ϕ_{shear} for different heat treated states of the *Full Stack* coating. Error bars denote the standard deviation of values produced by different minimisations.

This analysis shows a clear decline in ϕ_{shear} and ϕ_{bulk} for each heat treated state of the *Full Stack*, in corroboration with the changes observed in each component stack in sections 6.0.3 and 6.0.2. These results are collated in Fig 6.19.

Indeed the behaviour of ϕ_{shear} and ϕ_{bulk} for each coating stack follows a decreasing trend reaching a minimum at 550 °C, save the *Lower stack* which shows a minimum at 400 °C. The level of ϕ_{bulk} predicted from the *Upper* and *Full* stack show some parallels, with their respective magnitudes differing between 32 % - 51% with the largest difference occurring after heat treatment at 500°C. The *Upper* and *Full* stacks levels of ϕ_{shear} both show a decreasing trend with heat treatment where values differ by 35 % - 37%.

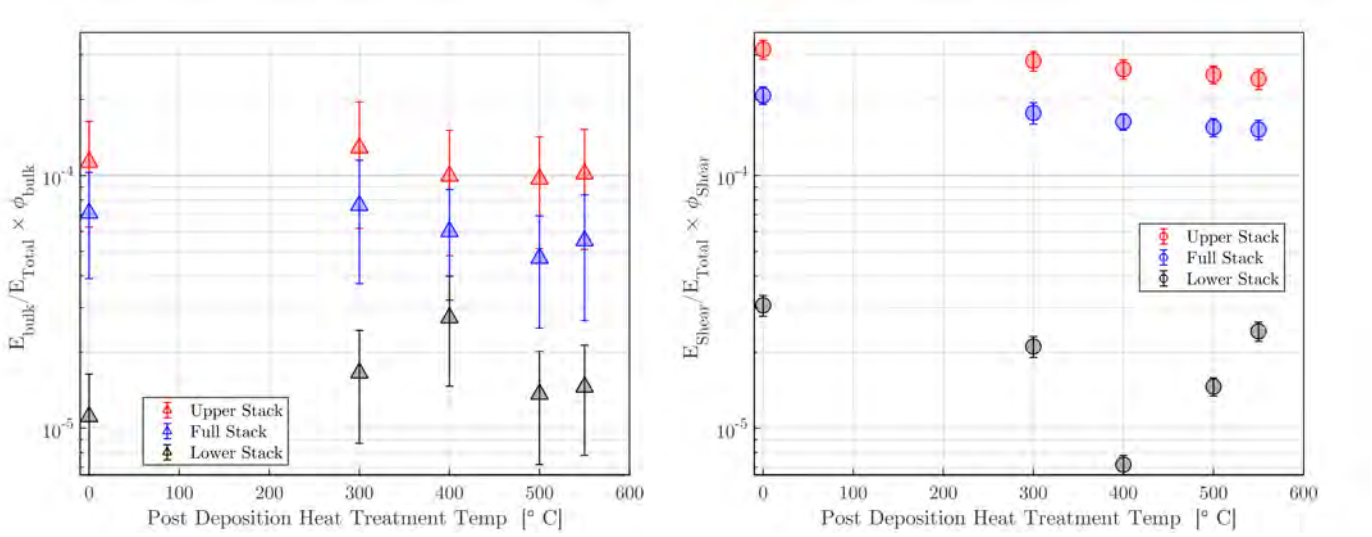

Figure 6.19: Scaled values of ϕ_{bulk} and ϕ_{shear} for different heat treated states of the *Full Stack*, *Lower Stack* and *Upper Stack* coatings. Error bars are the standard deviation of the scaled contributions for each resonant frequency.

The scaled values of ϕ_{bulk} for each coating stack shown in Fig 6.19 has a large spread, producing no clear trend with increasing heat treatment temperature. While each scaled value of ϕ_{shear} for the *Full Stack* and its component show clear changes with each heat treatment, the contributions of each decrease with increasing temperature. The *Lower Stack* material, with the lowest coating loss, shows the largest decrease in ϕ_{shear} between 400 °C and 500 °C. These results show that in each coating stack losses associated with ϕ_{shear} dominate over those produced by ϕ_{bulk} and highlight the sensitivity of the GeNS technique to measure changes in both loss mechanisms. As the dominating source of loss in the *Full Stack* is associated with ϕ_{shear}, its interactions with a sensing laser in a GW detector will be reduced compared to a coating dominated by ϕ_{bulk}, and will therefor create less thermal

noise.

In the next chapter, measurements of the aforementioned cantilever samples will be detailed as a function of environmental temperature (4 K- 293 K) and heat treatment temperature. Work carried out by I. Martin [204] and others [93, 205] have shown that measuring the mechanical loss of a coating layer at cryogenic temperatures can provide structural information on the coating material and can be used as an indication of crystallisation.

6.0.5 Thermal Noise of Multi-Material Coatings

The development of a multi-material coating such as the *Full Stack* has shown that the use of layers of aSi close to the substrate in place of Ta_2O_5 can reduce the level of mechanical loss produced by the coating stack, without significantly increasing the absorption of the coating (see Chapter 4). As described in Chapter 2 the level of thermal noise from such a multi-material coating stack could also be reduced due to these layers. In this section the thermal noise of the *Full Stack* coating material will be discussed in relation to applications in room temperature (aLIGO) [157] gravitational wave detectors.

In order to calculate the total Brownian thermal noise contributions of a multi-material coating stack, the material properties of the substrate and coating and properties of the main interferometer laser beam are required. As was discussed in more detail in Chapter 2 the Brownian noise for a multi-material coating can be calculated using Eq. 2.35 and Eq. 2.36 where the light penetration into the coating layers is dictated by the thickness and refractive index of each layer is considered as constants.

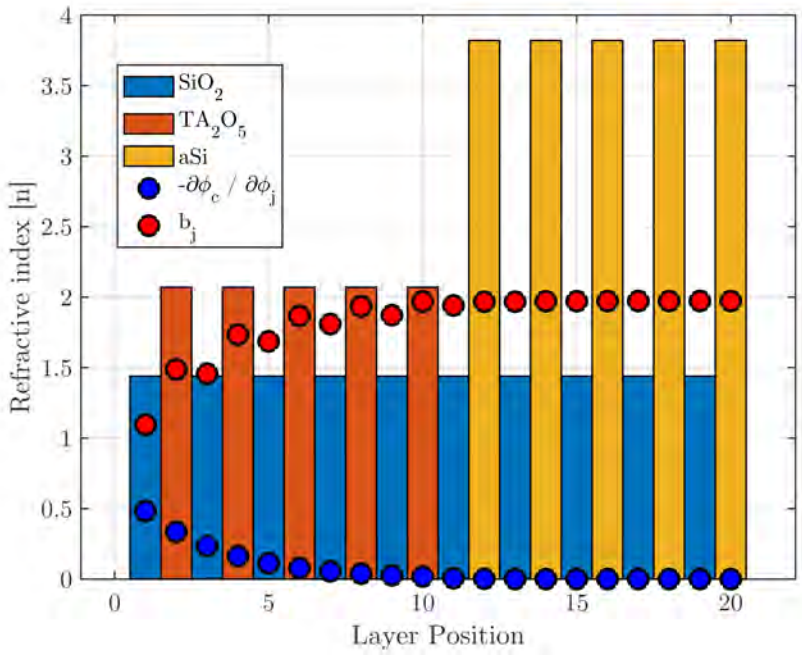

Figure 6.20: Illustration of the refractive index of each coating layer in the *Full Stack* coating. Calculated values of $\frac{\partial\theta_c}{\partial\theta_j}$ and b_j are shown for each coating layer.

By recursively calculating the round trip phase noise $\frac{\partial\theta_c}{\partial\theta_j}$ and b_j terms for the *Full Stack* coating shows one of the main benefits of this multimaterial design, where coating stack layers closer to the substrate contribute more to the coating thermal noise. This effect is clearly shown by viewing b_j as a function of layer position in Fig 6.20.

The calculated coating thermal noise produced by the *Full Stack* coating for a single ETM at the coating design wavelength of $\lambda_{\mathrm{IFO}} = 2000\,\mathrm{nm}$ at a temperature of 290 K is shown in Fig 6.21. As the level of coating loss in the *Upper Stack* and *Lower Stack* coatings reach a collective minimum at 500 °C the lowest level of coating Brownian thermal noise occurs at this heat treatment temperature. Comparing CTN level at 100 Hz, between the optimum *Full Stack* and current aLIGO coating shows a 49.5% lower CTN level. After heat treating to 550 °C the predicted Brownian thermal noise increases due to expected crystallisation of the *Lower Stack* coatings in this state.

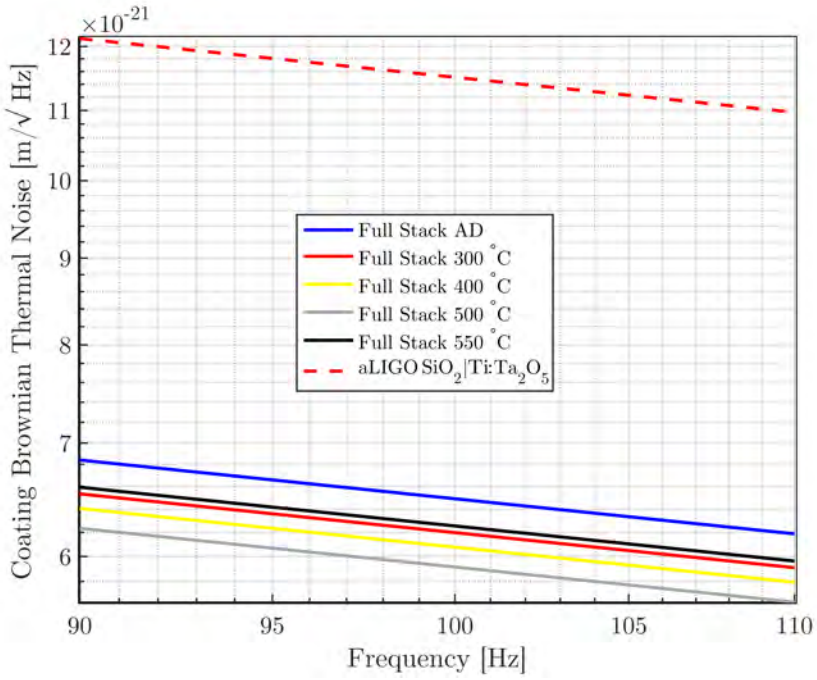

Figure 6.21: Calculated thermal noise produced by the *Full Stack* coating material using measured *Upper Stack* and *Lower Stack* losses.

These measurements of a multi-material coating containing five bi-layers of aSi:SiO_2 are the first experimental verification that such a configuration can reduce the mechanical dissipation of a HR coating stack. Through multiple heat treatments it was found that after the 500 °C heat treatment, the loss of the *Full Stack* reduced by 26.6% to a value of 1.9×10^{-4}. The coating thickness of the *Full Stack* is much thinner compared to an aLIGO HR stack of equivalent reflectivity. The reduction in mechanical loss and coating thickness gives a 20% reduction in coating Brownian thermal noise compared to current aLIGO HR coatings. Using the method derived by Yam [134] and Hong [132] the total coating Brownian noise was produced for each heat treated state of the *Full Stack* coating. In each case the lower mechanical loss of the aSi layers showed higher contribution the CTN with a predicted total reduction of 49.5% at 100 Hz.

6.1 Conclusions

Thermal noise created by HR mirror coatings currently limits gravitational wave detector sensitivity between $\sim$50 - 150 Hz. Measurements of three silica disks,

coated with the *Full Stack*, *Lower Stack* and *Upper Stack* were carried out using the GeNS apparatus where the mechanical loss of each coating was characterised as a function of heat treatment temperature. It was found that the mechanical loss of each coating stack decreased as a function of heat treatment temperature, reaching a minimum after heat treatment at 500°C. By evaluating the combined loss of the *Upper Stack* and *Lower Stack* coatings, remarkable agreement with the measured loss of the *Full Stack* is found at 500°C. Repeating this process for different heat treated states of each coating suggests changes observed in the *Upper* and *Lower Stack* coatings do not occur to the same extent when configured in the *Full Stack*.

The measured mechanical loss of each coating stack was separated into its bulk and shear loss components. Results from these calculations suggest that ϕ_{shear} decreases as a function of heat treatment temperature with values of ϕ_{bulk} limited by the sensitivity of the measurement technique. Differences in the trends of ϕ_{bulk}, ϕ_{bulk} and ϕ_{coating} could also suggest that mechanical loss of coating layers in the *Full Stack* are affected differently by heat treatment compared to the other, thinner coating materials.

The Brownian thermal noise of the *Full Stack* used in an aLIGO-like detector was calculated as a function of heat treatment temperature, showing a 47.6% reduction in Brownian thermal noise compared to an RLVIP coating stack comprised of SiO_2-Ta_2O_5. This result proves the working principal of using layers of low mechanical loss situated close to the substrate can be employed to produce thinner, lower thermal noise coating stack with 99.9999% reflectivity.

Chapter 7

Cryogenic Mechanical Loss

Future third-generation gravitational wave detectors such as the Einstein Telescope: Low-Frequency detector [78] (ET-LF) and LIGO Voyager [104] are planned to operate at cryogenic temperatures (10 K and $\sim$123 K), reducing the total thermal energy available to produce thermal noise (see Chapter 2).

Coating loss can be strongly temperature dependent, with evidence for mechanical loss peaks at low temperature (Ta_2O_5 at 40 K [204] and SiO_2 at 30 K [93, 204, 228]) in the coatings currently used in aLIGO and AdV [208] and new coatings are required to allow cryogenic detectors to meet their sensitivity goals [104, 201].

This chapter extends the room temperature study detailed in Chapter 5, adapting the automated GeNS support to measure the mechanical loss of cSi substrates between 80 K and 300 K.

7.1 Silicon at Low Temperatures

The Young's modulus, Y, of crystalline silicon (cSi) is temperature dependent [96, 229, 230]. This, along with any changes in the Poisson ratio, will result in the vibrational mode frequencies of a silicon sample changing with temperature. In an anisotropic material such as cSi, the strain (ε) and stress (σ) are defined by matrices where each matrix element is expressed in terms of stiffness C or compliance S.

$$\varepsilon = S\sigma, \quad \sigma = C\varepsilon, \tag{7.1}$$

Hooke's law then takes the form of a 3rd rank tensor which describes the stress and strain constants for an anisotropic material where each value of C and S are defined by their orientation to the Miller indices of the cubic crystal structure [230].

$$\sigma_{ij} = Y_{ijk}\varepsilon_{kl}, \quad \varepsilon_{ij} = Y_{ijk}\sigma_{kl}. \tag{7.2}$$

Each subscript *i, j, k* in Eq. 7.2 relates to the value of C or S in each crystal plane of the same subscript. Using condensed notation defined by [230] the $3^3 = 27$ terms needed to describe the anisotropic properties can be expressed in terms of C and S for each axis:

$$\begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ \sigma_5 \\ \sigma_6 \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{44} \end{bmatrix} \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \\ \varepsilon_5 \\ \varepsilon_6 \end{bmatrix}. \tag{7.3}$$

All cSi samples described in this chapter are produced from a larger wafer, such that their larger faces are aligned with the <110> orientation. In this crystal axis the Youngs moduli, Poisson ratios and bulk moduli are defined in that plane, i.e. Y_x, $Y_y = 169\,\text{GPa}$, but $Y_z = 130\,\text{GPa}$. Similarly for Poisson ratio $\nu_{yz} = 0.36$, $\nu_{zx} = 0.28$, $\nu_{xy} = 0.064$ and the Shear modulus $G_{yz} = G_{zx} = 9.6\,\text{GPa}$, $G_{xy} = 50.9\,\text{GPa}$ [230]. It has been shown by [230] that the following strain matrix is produced:[1]

$$\begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ \sigma_5 \\ \sigma_6 \end{bmatrix} = \begin{bmatrix} 165.7 & 63.9 & 63.9 & 0 & 0 & 0 \\ 63.9 & 165.7 & 63.9 & 0 & 0 & 0 \\ 63.9 & 63.9 & 165.7 & 0 & 0 & 0 \\ 0 & 0 & 0 & 79.6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 79.6 & 0 \\ 0 & 0 & 0 & 0 & 0 & 79.6 \end{bmatrix} \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \\ \varepsilon_5 \\ \varepsilon_6 \end{bmatrix}. \tag{7.4}$$

Finite element analysis (Ansys / COMSOL) was used to allow rapid, practical mode-

[1] Each value in units of GPa

frequency calculations.THe Young's Modulus Y [96, 229, 230], thermal conductivity κ, emissivity α [231] and heat capacity at constant pressure C_{p} [232] and thermal expansion [232] of silicon, all exhibit temperature dependencies. Changes in these material properties can manifest themselves as changes in the stress and curvature of the sample [203]. For uncoated substrates, this effect can be considered negligible as the material has a singular thermal expansion coefficient. For coated substrates, where the coating properties are also temperature dependent, temperature dependent stress and curvature can occur [233, 234].

7.2 Cryogenic Gentle Nodal Support

The auto-loading gentle nodal support (GeNS), which was shown to improve the repeatability of room-temperature loss measurements in section 5.1.3 was adapted for use in a cryogenic environment.

7.2.1 System Design

Due to the confined space inside the available cryostat[2], it was not possible to install cryogenic motors which would directly mimic the room-temperature setup. Instead a smaller voice coil actuator (VCA) - a permanent magnet encasing a smaller electromagnetic coil, was used to drive the translation stage. Mounting the coil in a fixed position below the translation stage allowed the stage to be raised and lowered by controlling the current flowing through the coil. It was found that the VCA required to be driven with a current of $< 0.01\,\mathrm{A}$ precision to produce smooth translation and ensure accurate final placement. To fit the VCA inside the nodal support the height of the 'table' (see Fig 7.1) and length of the PTFE translation rods were increased, lengthening the overall range of the translation stage.

[2]The cryostat system was a custom system (NDL-HDL-12) produced by IR Labs [235]

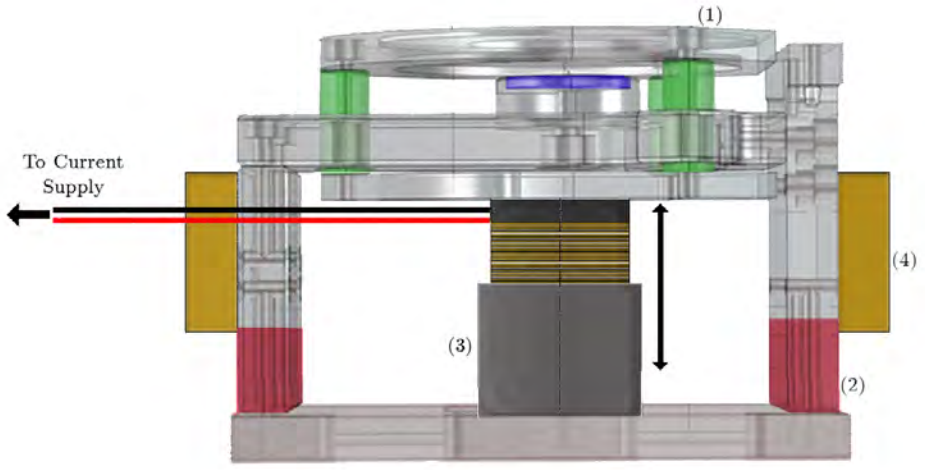

Figure 7.1: Schematic diagram (side view) showing the voice coil electromagnet assembly used for actuation of the CryoGeNS translating stage. (1) Changes to the upper locking plate were made to accommodate different disk dimensions. (2) Changes to the system height and shorter translation length of the stage rods are used to accommodate the VCA. (3) VCA (4) 2× 25Ω Resistive Heaters.

The cryostat was originally designed to be inverted after loading so that access to the cryogenic filling ports can be gained from above (see Fig 7.2). This is not compatible with the nodal support design which relies on balancing a disk on the retention plate so the cryostat was adapted to allow it to be loaded from below, with no need to invert it after sample installation. A thin aluminium plate previously used for conductive cooling/radiation shielding at the bottom of the cryostat was replaced by a thicker plate which could be used as a breadboard for the nodal support apparatus.

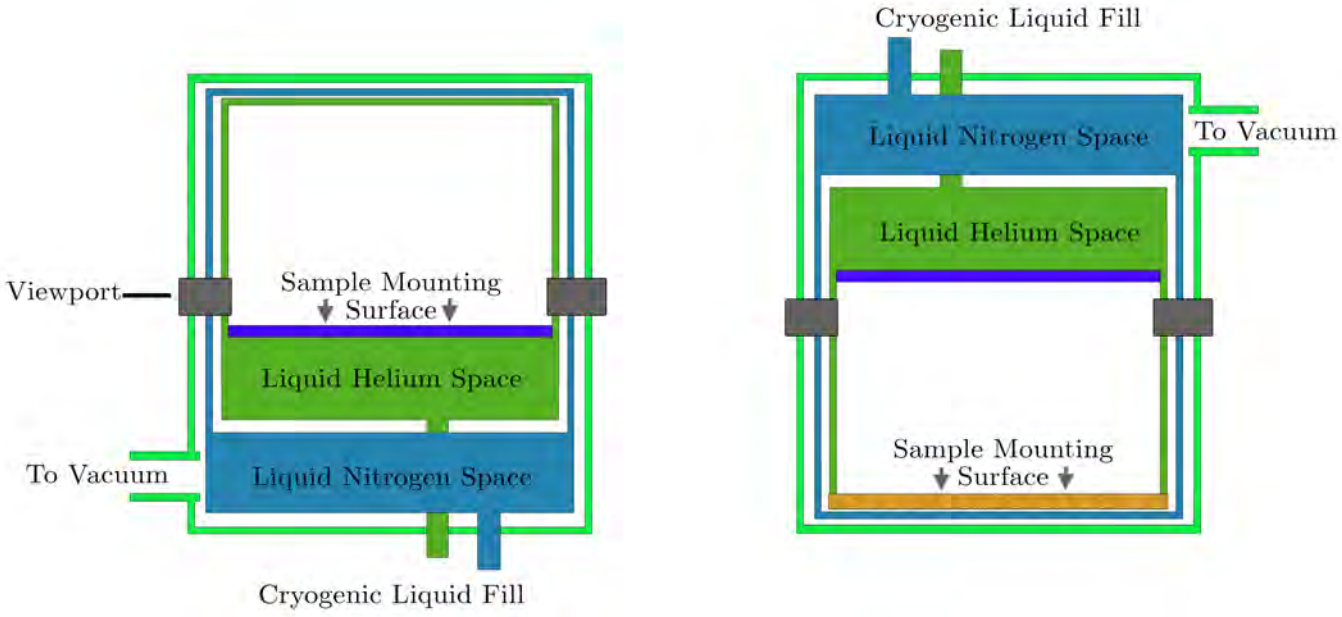

Figure 7.2: Rendering of IR labs cryostat (LHS) shows by design 'inverted' loading, (RHS) shows orientation required to load the nodal support.

FEA was used to determine a suitable thickness of the plate, to ensure that it would not deform excessively when mass was added and to ensure no resonances at problematic frequencies were introduced. Fig 7.3 shows the fundamental bending mode of the optimised plate with the nodal support attached. A plate thickness of 5 mm was found to produce a fundamental plate resonance at 521 Hz: far from the fundamental mode frequency of the samples to be measured, which occur at $f \approx 1$ kHz.

Figure 7.4 shows the Cryogenic Gentle Nodal Support (CryoGeNS) experiment rotated 90° anticlockwise with respect to Fig 7.1. The cryostat contains two nested chambers for cryogenic liquids - one cooling the entire inner experimental space (T$\approx$4 K), the other cooling a radiation shield surrounding this space(T$\approx$80 K). Once the nodal support is fixed to the 5 mm thick aluminium plate, the plate and system are raised into position and fixed to the outer walls of the inner experimental chamber (see Fig 7.4). The cryostat is mounted vertically on a supporting frame for ease of access to the system.

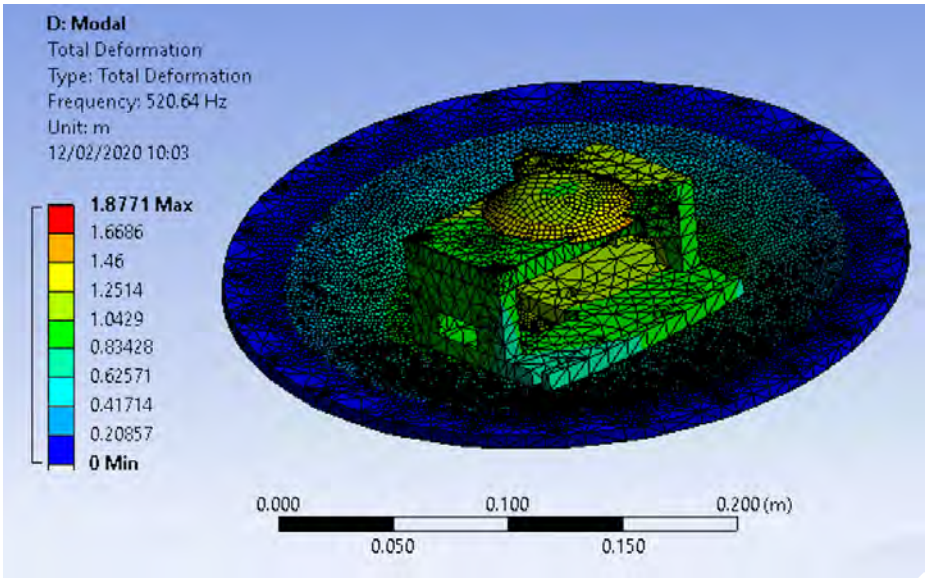

Figure 7.3: Ansys model showing the fundamental bending mode (521 Hz) of the optimised baseplate and mounted nodal support. Colour scale depicts relative deformation, with red indicating a large relative deformation and blue showing little/no deformation.

To load a sample into the system, the chamber is opened and disk sample is placed into the nodal support with the translation stage raised. A $\sim 5\,\mathrm{mW}$, HeNe (633 nm) sensing laser is then aligned to the edge of the disk such that the motion of the reflected spot can be measured with an external split photo-diode. Once the laser is aligned, the apparatus is raised into the cryostat and fixed into the system. The pressure inside the system is reduced using a turbo-molecular pump to $<1\times10^{-5}\,\mathrm{mbar}$.

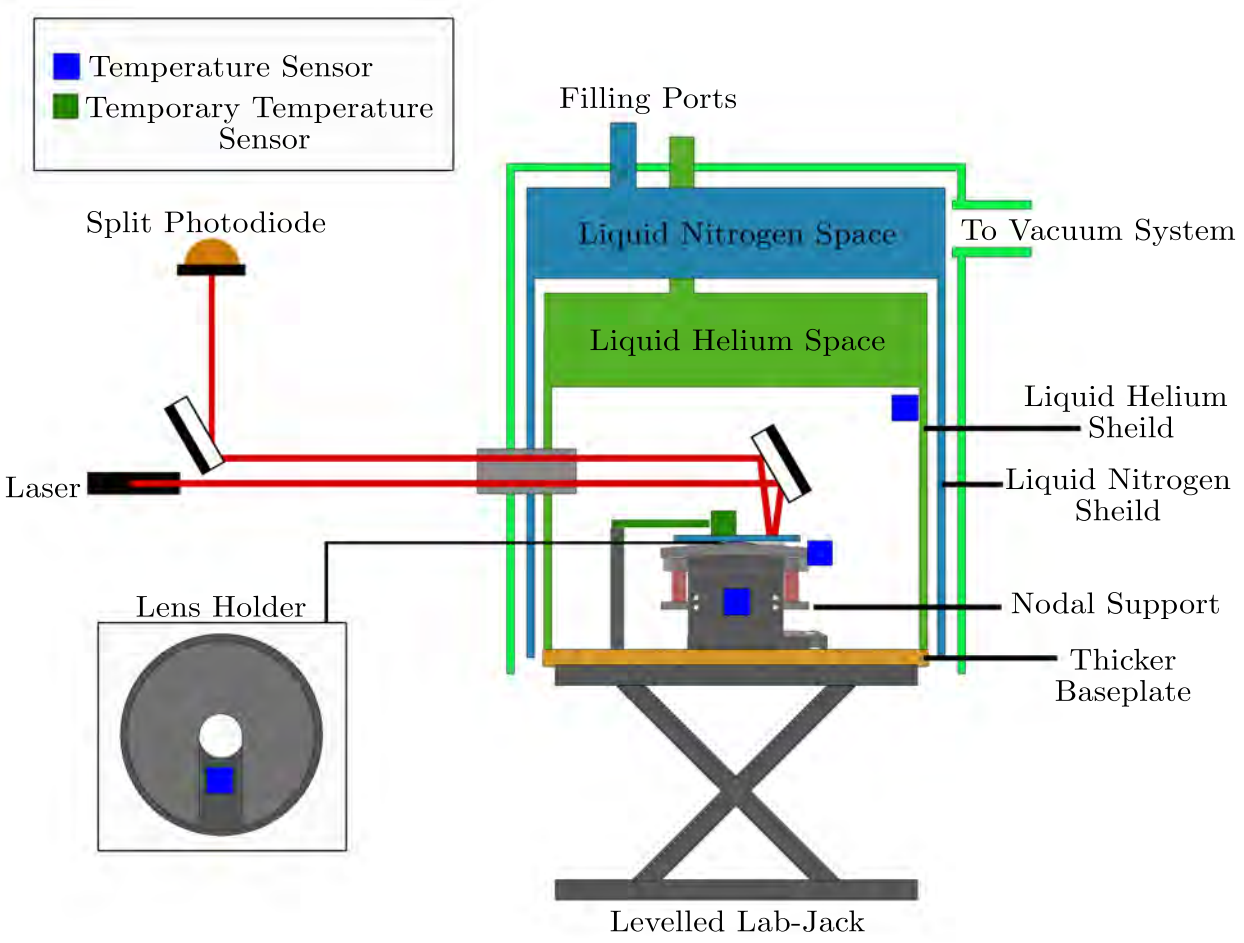

Figure 7.4: Representation of experimental CryoGeNS set-up (not to scale). Placement of temperature sensors are shown. A Top-down view of the lens holder is also shown for temperature sensor placement. Cryostat supporting apparatus and optics bench are not shown.

During the cool-down process, the pressure further reduces to $<1\times10^{-6}$ mbar. Four Lakeshore DT-670 silicon diode sensors [236], shown in Fig 7.4, monitor the temperature of the cryostat and nodal support. One sensor is mounted in direct contact with the LHe shield, one on the front of the support, one on the side of the Disk Retention plate and one directly below the balancing lens. The Lens sensor is the closest sensor to the suspended sample.

Fig 7.5 shows the recorded temperatures as the system cooled over 24 hours with the VCA deactivated and the stage in the lowered position. The LHe shield sensor reduces from 300 K to 81.8 K after 23 hours. The ‘GeNS Body’ and ‘Lens’ sensors also show a similar decrease in temperature over the measured time, stabilising after

10 hours. However, the ‘ Disk Retention’ plate sensor cooled at a much slower rate over the run and did not reach equilibrium.

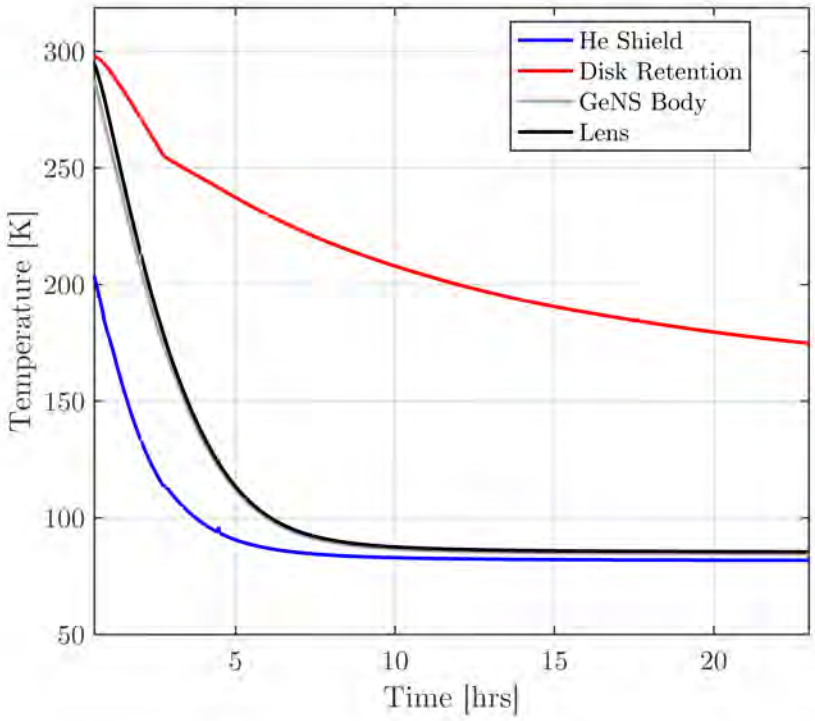

Figure 7.5: Temperature recorded from sensors mounted on the helium shield (blue), Disk Retention plate (red), GeNS Body (grey) and Lens (black) as the system was cooled with liquid nitrogen. The positions of each sensor is shown in Fig 7.4[3].

As the ‘Disk Retention’ plate was to be used as a means to estimate the temperature of a mechanical loss sample during the cool-down process, this trend could remove this avenue of sample temperature estimation. In the next section a study involving FEA modelling will be used to try to discern the cause for the observed difference in temperature for the ‘Disk Retention’ plate and other sensors in the apparatus.

[3]N.B Temperature recording started after the cryostat had been filled, therefore all temperature sensors would have started at 300 K.

7.2.2 Heat Flow Modelling

To understand the significantly different cooling rate of the 'Disk Retention' plate, the upper 'table' of the nodal support was modelled using a Heat Transfer analysis COMSOL® using the geometry of the nodal support system drawn in SolidWorks.

As shown in the previous section, the 'Disk Retention' plate cooled much more slowly than other parts of the CryoGeNS system. To better understand this disparity, the nodal support's response to heating was measured experimentally, allowing the rate of heating for different parts of the system to be calculated. This was verified using FEA modelling.

The nodal support has two primary sources of heating: the two 25 Ω resistive heaters on either side of the system, and heat arising from the VCA (see Fig 7.1). When the stage is fully raised (consistent with its initial position during system cooling), the 'table' and bottom of the translation stage should be in physical, and therefore thermal contact. Under these conditions heat applied from either the VCA or resistive heaters should allow the temperature of the 'Disk Retention Plate' to be increased.

In order to find which source of heating has a dominant effect on the temperature of the 'Disk Retention plate' the temperature change of the stage was measured in response to the heaters which were powered on full for a period of 149 s. These results are shown in Fig 7.6. With these measurements, the mass, thermal conductivity and dimensions of the system, a total equivalent heating power of 17.4 W ($0.02\,\mathrm{K\,s^{-1}}$) was calculated with both heaters fully active. The nodal support's expected response to heating with the VCA was calculated from the number of turns in the coil and the 0.4 A required to raise the stage ($4.8\times10^{-3}\,\mathrm{K\,s^{-1}}$). From this calculation a heating rate is estimated.

With two different heating rates calculated from different sources in the stage, these values could then be used as inputs for FEA modelling in COMSOL. Two final models of different heating regimes were made and compared to the data recorded as the system was heated. In each simulation (i.e. model 1 - application of resister heating and model 2 - heating by VCA) all bodies are considered in perfect thermal

and physical contact. This simplifies each simulation and reduces computation time.

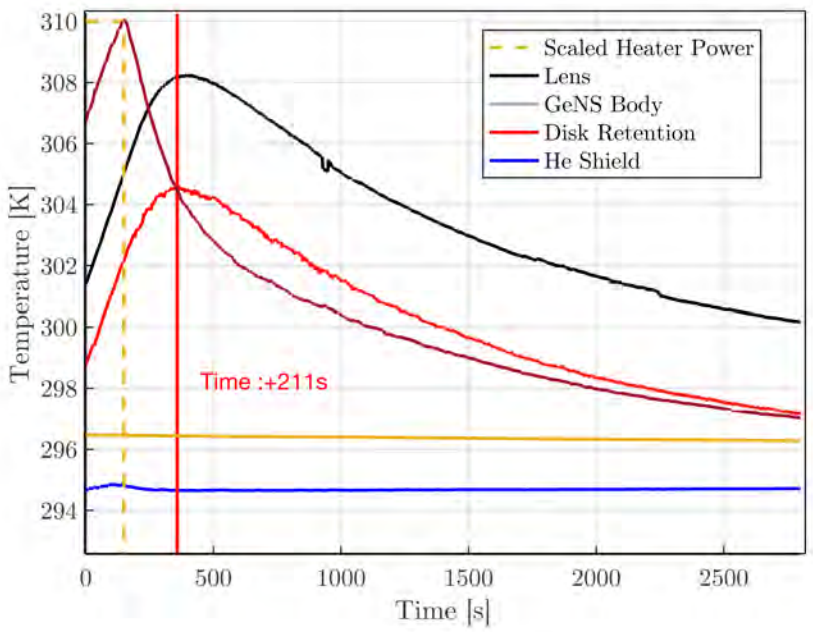

Figure 7.6: The Measured heating response of the CryoGeNS system with full heating power, starting from T = 300 K. The total power flowing through each heater is scaled to show their state at each time, their response can be considered binary.

Only the upper locking plate, central table plate, and PTFE roads were simulated as these parts are expected to be the most susceptible to the applied heating in each model. The placement of 25 Ω resistive heaters (see Fig 7.1), was replicated by adding heat sources at the front and rear of the central table plate. Each component in the model had a surface emissivity ϵ defined and was allowed to radiate thermal energy to the surroundings, which were specified to be at 80 K.

Figure 7.7 shows the simulated system temperature in model 1, with heat flow simulated from the resistive heaters on either side of the system. Heat was applied for 7000 s and the temperatures at the locations of the temperature sensors were extracted from the simulation. These results are shown below in Figure 7.8.

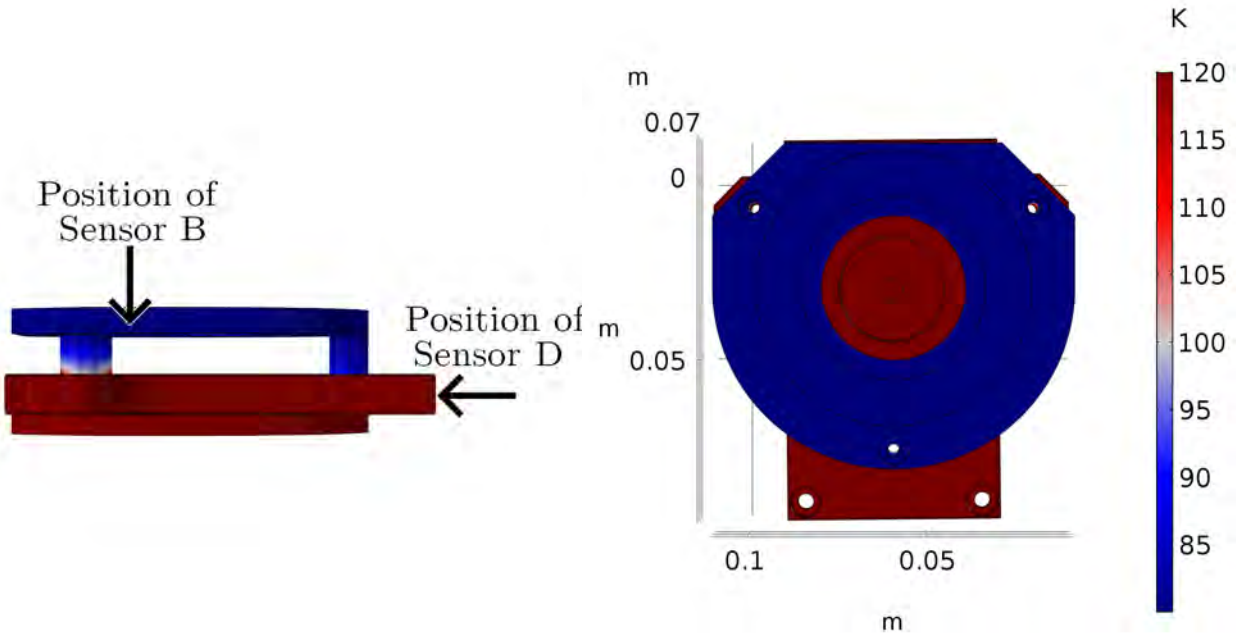

Figure 7.7: Finite element model of heating applied to the table plate of the CryoGeNS after 1000 s. This was to replicate the heat flow through the apparatus when 2×25 W resistive heaters were active.

In model 1 it can be seen that the temperature of the lower portion of the stage, central table plate and regions of PTFE in contact with this plate have increased in temperature reaching temperatures between 115 K- 120 K. In contrast the upper locking plate remains below 85 K, with the heat flow through the system to the Disk Retention plate inhibited mainly by the PTFE rods in the translation stage. The simulation was run for a total of 7000 s where the final temperature at B reached 109.2 K, 158 K lower than Sensor D.

In model 2, the total heating rate of $4.8\times10^{-3}\,\mathrm{K\,s^{-1}}$ was applied directly to the centre of the lower locking plate emulating the position of the VCA, calculating the temperature response of the system again for 7000 s. This simulation yields a similar trend with each analysed component below the PTFE rods reacting to the increase in temperature, with the Disk Retention plate remaining mostly unchanged over the simulation time, increasing by only 15 K.

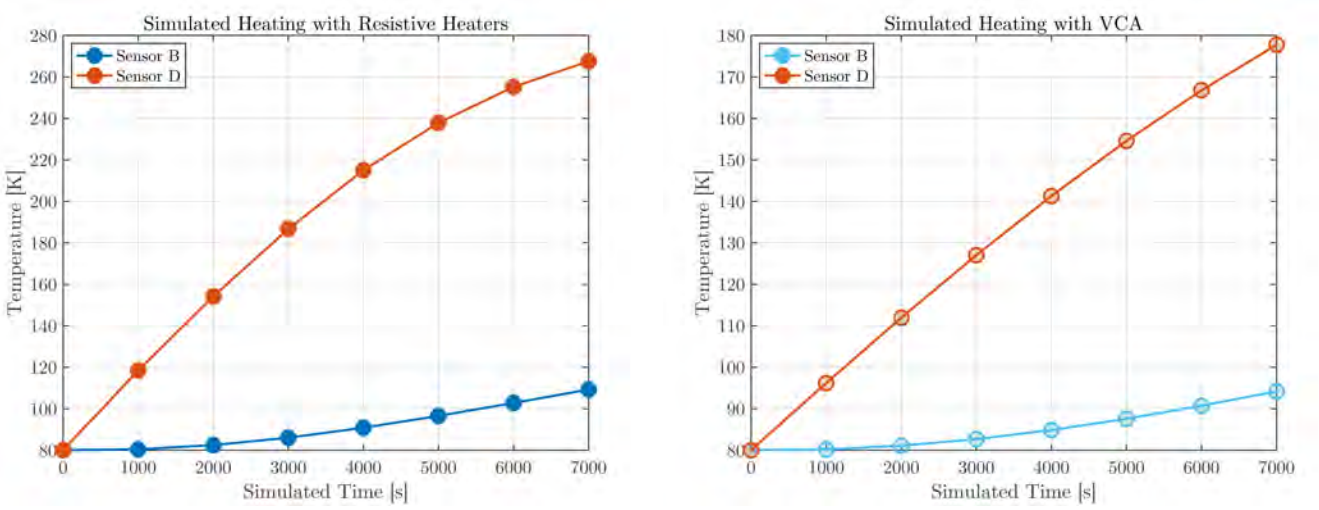

Figure 7.8: The calculated temperature response from the equivalent sensor positions of Sensor B (Disk Retention Plate) and Sensor D (CryoGeNS Body): (LHS) Heat flow replicated through stabilisation heaters (RHS) Heat Flow replicated from VCA heating.

The change in temperature with time from models 1 and 2 can be compared to the measured cooling rate of the Disk Retention sensor shown in Fig 7.5. The simulated results shows that in both cases, the temperature of the upper locking plate changes at a rate much slower than the rest of the simulated geometry. This analysis also shows that the VCA can heat the 'table' plate of the CryoGeNS by direct conduction, but has minimal impact on the temperature of the sample when the stage is raised. These results approximate the rate of $\frac{dT}{dt}$ measured in the nodal support system (see Figure 7.8) where a prediction of $+5.8\,\mathrm{K}$ was simulated over 150 s, close to the $+3.4\,\mathrm{K}$ which was measured. It should be noted that as the experimental results were produced at 300 K, the expected increase with heater power will be smaller than if the experiment started at 80 K.

Both of these simulations suggest that the experimental discrepancy in Disk Retention plate temperature is due to low thermal conduction of the PTFE rods, which limit the heat flow during cooling. When the translation stage is raised, the plate and sensor are no longer in good thermal contact with the rest of the system (due to the PTFE rods). The limited heat flow from the Disk Retention plate will limit not only the temperature of the Disk Retention plate, but the temperature of any samples which it is in physical contact with. In subsequent iterations of the

nodal support, it may be advantageous to replace these rods with a higher conducting material, with similar bearing properties.

Subsequently, due to the levels of heat produced from the VCA during its operation, once the cryostat cooling process has started, it was powered off - lowering the stage and placing the sample in direct thermal contact with the balancing lens, decreasing any potential limitation.

Once the Lens and Body temperature sensors have reached thermal equilibrium, the mechanical loss $\phi(f,\mathrm{T})$ of a balanced sample can be measured using the ringdown method detailed in section 5.1.3. The temperature of the nodal support is then actively increased using the two 25 Ω resistive heaters placed on the front and rear of the support, their output controlled by a proportional–integral–derivative (PID) controller[4]. Having a temperature sensor on the sample during measurements would introduce excess damping, dominating over the loss of the sample itself. For this reason, the surrounding sensors were used to estimate of the temperature of the sample.

7.3 Sample Temperature Calibration

A temperature sensor was affixed to a $\varnothing = 2$" diameter, 360 µm thick silicon disk, as shown in Fig 7.9, to provide a calibration between the disk temperature and the temperature of the various parts of the nodal support. This calibration can then be used to calculate the sample temperature when no measurement sensor is in direct contact with the sample during loss measurements. To counteract the mass of the affixed sensor and keep the disk horizontal to the balancing lens (as would be required during later loss measurements) a small amount of cryogenic compatible varnish (ICEles03) was used [238] at the interface between the curved lens and the disk sample. The small volume was used to mimic the original contact area between the horizontal sample and curved lens.

The use of the varnish would have an effect on the thermal conduction between the monitored sample and the rest of the apparatus. An understanding of the effect that

[4]PID control was carried out using a Lake Shore Cryotronics model 336 temperature controller [237].

the varnish layer has on thermal conduction is required. Assuming no deformation in their contact, r_c, the contact radius, is given by Hertz theory of contacting bodies [239]:

$$r_c = \left[\frac{3}{4} M g R_{\text{oc}} \left(\frac{1-\sigma_{\text{disk}}^2}{E_{\text{disk}}} + \frac{1-\sigma_{\text{lens}}^2}{E_{\text{lens}}}\right)\right]^{1/3}, \tag{7.5}$$

where M is the mass of the disk, g is the acceleration due to gravity, R_{oc} is the radius of curvature of the balancing lens E_{disk}, E_{lens}, σ_{disk}, σ_{lens} are the Youngs Moduli and Poisson ratios of the disk and lens respectively. If r_c is known, the total heat flow between the disk and lens can be estimated. As the amount of varnish used attempts to mimic the area as calculated with Eq. 7.5 it was assumed that the delay time between the disk and lens sensors would be largely unaffected. As the thermal conductivity of the varnish is less than that of the silicon/silicon interface ($k_{\text{varnish}} = 0.44\,\text{W/mK}$ at $300\,\text{K}$) [238], then the time taken to reach equilibrium will be underestimated by this procedure.

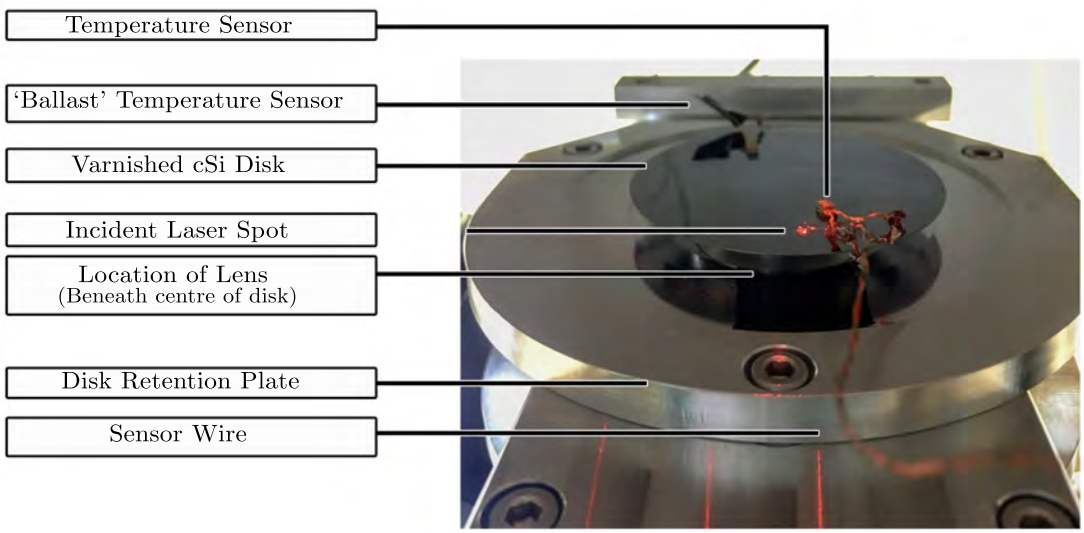

Figure 7.9: Image of temperature sensor position relative to the incoming laser spot on a $\varnothing = 2$" silicon disk during the varnished disk calibration run.

Figure 7.10 shows the recorded values from each temperature sensor during a heating test. Prior to the first heating run the system was cooled overnight, with the helium baseplate reaching an ultimate temperature of 79 K, 1 K higher than the

desired temperature 78 K. The Lens, Disk and GeNS Body sensors had reached thermal equilibrium. The heat was applied to the body of the nodal support with the resistive heaters, increasing the temperature of the system in 10 K and then 20 K intervals.

It should be noted that during this measurement the Lens, Disk and GeNS Body sensors reached thermal equilibrium faster than the Disk Retention sensor. This sensor did not reach a stable temperature due to poor thermal conduction of the PTFE rods (see section 7.2.1).

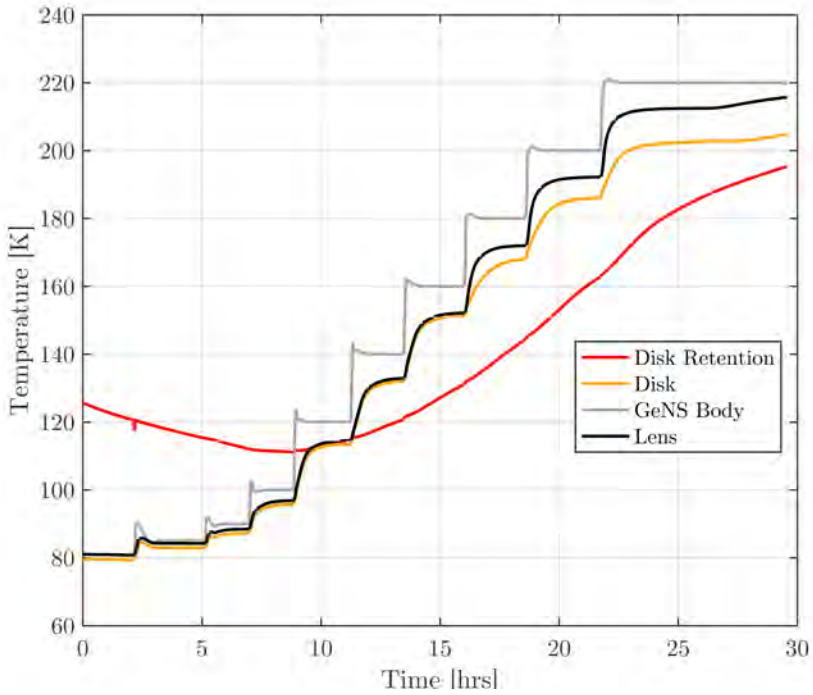

Figure 7.10: Recorded temperature of all tank sensors during the 'Varnished Disk' calibration run.

The results show that the GeNS Body sensor, which is closest to a heating element, has much faster response to heating than other parts of the system. The response of the Lens sensor and Disk sensor were found to be similar, due to these sensors being located close to each other.

Fig 7.11 shows the rate of change of temperature $\frac{dT}{dt}$ for the Lens and Disk sensors. During times where the system is stable, $\frac{dT}{dt}$ oscillates around 0. Each

heating step appears as a large peak in the data. At temperatures below 150 K, the Disk and Lens sensors heat and stabilise at the same rate. In this range, the temperature offset between both sensors is $\Delta T = 1.4\,K$ at 80 K to $\Delta T = 0.9\,K$ at 160 K. Above this temperature, the level of liquid nitrogen inside the cryostat began to decrease, causing the sensor relationship to deviate.

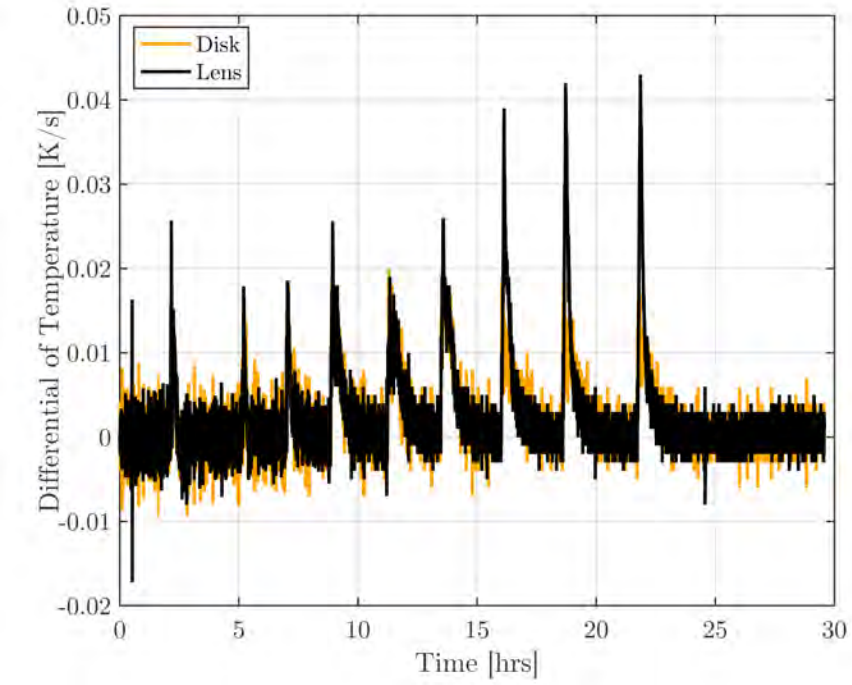

Figure 7.11: Differential of recorded temperature on Lens and Disk Sensors for 'Varnished Disk' calibration run over a 30 hrs period.

The data shown in Fig 7.10 was used to relate the measured disk temperature to the temperatures measured at different points of the nodal support. By producing a relationship which would allow each sensor to be expressed as functions of the other, the disk temperature could be estimated during loss measurements, when it was not possible to have a sensor on the disk.

The LabView program used to measure and record ringdown data uses the temperature of the GeNS Body sensor for PID control, and all ringdown measurements were recorded as a function of this temperature. In order to convert these values of temperature from the GeNS Body sensor to Disk temperature, a

transfer function from GeNS Body and Lens sensors was produced using datapoints where the cryostat temperature was deemed stable ($\frac{dT}{dt}$ $<0.01\,\mathrm{K\,hr^{-1}}$). The relationship between GeNS Body and Lens was then approximated using a 6^{th} order polynomial fit, shown in Fig 7.12 allowing temperature points to be interpolated.

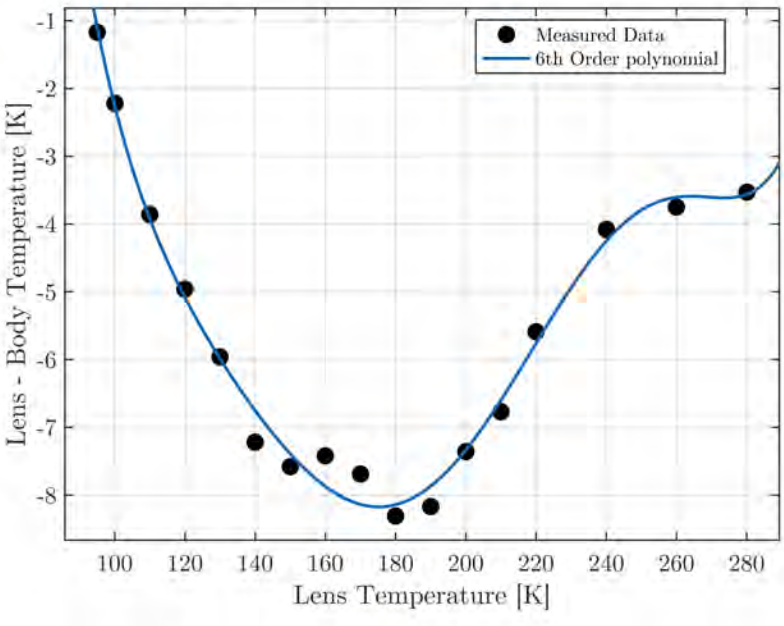

Figure 7.12: Measured difference between GeNS Body Sensor and Lens sensor as a function of Lens sensor temperature.

Once this relationship had been established, a similar process was used to produce a conversion from Lens sensor to the measured temperature of the disk.

This method of translating the PID controlled Body temperature to the Disk temperature was subsequently tested, using measurements from the 'Varnished Disk' calibration run at times when the system was thermally stable. Fig 7.13 compares the directly measured disk temperature to the disk temperature calculated using the relationship derived between the Lens, Body and Disk sensors.

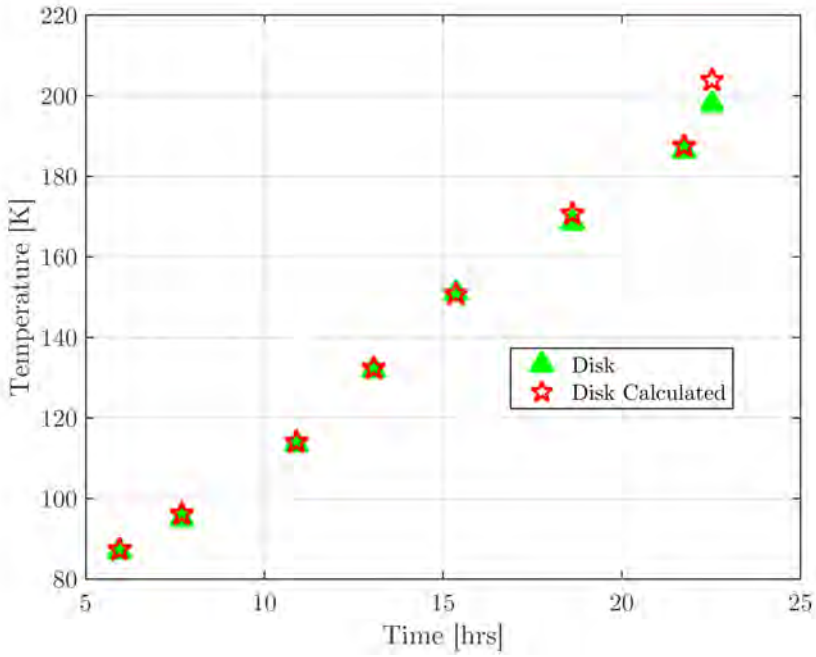

Figure 7.13: Comparison of measured Disk sensor temperatures and those calculated from the equivalent Body temperature.

Throughout the measured temperature range (80 K and 220 K), this method of conversion is able to reproduce the temperature of the disk sensor from other points in the system. It, therefore, can be used to infer the temperature of the disk where no sensor is present within this temperature range. However, as this calibration run was only carried out at the lower temperature range at which this system is to operate, a second method is required to calculate the temperature of the disk at $T > 220\,K$.

This method uses the temperature dependent properties of the sample such as $Y(T)$ and change in its resonant frequencies with temperature ($\frac{\Delta f}{dT}$) to infer the temperature of the sample (see section 7.1). During mechanical loss measurements, the temperature of the disk sample is then estimated using a combination of the low temperature calibration and the measured frequency shifts of the modes where the disk is allowed to oscillate freely.

Fig 7.14 shows the measured frequency shift of a 1251 Hz First family (2,0) mode between 80 K and 300 K. Measuring the variation in the frequency between 80 K and

300 K shows an inverse relationship between the measured frequency and the temperature.

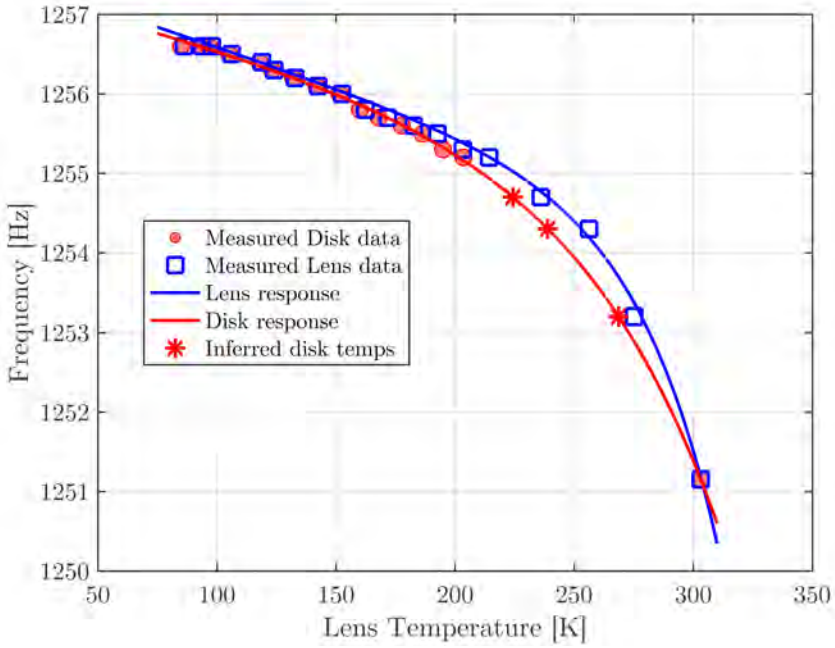

Figure 7.14: Measured frequency shift of a 1251 Hz First family (2,0) mode between 80 K and 300 K, assumed with a 2 term exponential fit.

As the method shown in Fig 7.13 allows for conversion between Lens and Disk temperatures (T<220 K) this can then be the basis to ascertain the disk temperature outwith this range. At 300 K the cryostat is in thermal equilibrium and all sensors are in agreement, and therefore the measured temperature is unchanging in the Lens or Disk reference frame. In Fig 7.14 the resonant frequency of a 1.2 kHz mode is plotted as a function of Lens temperature (shown in blue), points at indefinite temperatures between 220 K and 280 K are highlighted. It was found that $f(\mathrm{T_{lens}})$ data can be assumed by a two-term exponential function (of the form $Ae^{aT} + Be^{bT}$). Shown in red are the same values of frequency, translated to disk temperature. These values are approximated by a second function of the same form. By evaluating this second function at the measured frequencies between 220 K and 280 K, the equivalent disk temperature is calculated.

7.4 Characterisation of uncoated cSi Disks

To obtain ϕ_{coat}, the loss of the uncoated sample must first be measured in the same manner as is detailed in section 5.1.1. Initial measurements inside the system were used to test the repeatability of ϕ_{mech} and temperature stability of the apparatus. Fig 7.15 details the dimensions of the silicon disk substrates used for cryogenic loss measurements in this section.

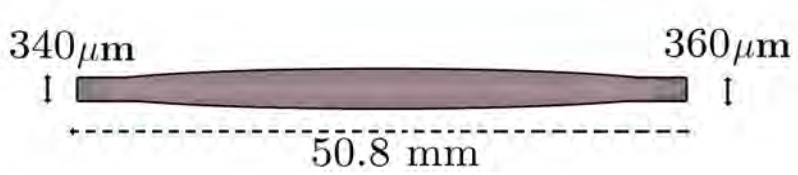

Figure 7.15: Illustration of $\varnothing = 2$" (50.8 mm), $t \approx 360$ µm silicon disk used for cryogenic mechanical loss measurements.

The disks were inspection polished on both flat surfaces, with the <100> crystal plane oriented with the flat surfaces of the sample. The surface flatness of the samples was tested using a Zygo GPI XP interferometer [240] resulting in the profiles shown in Fig 7.16.

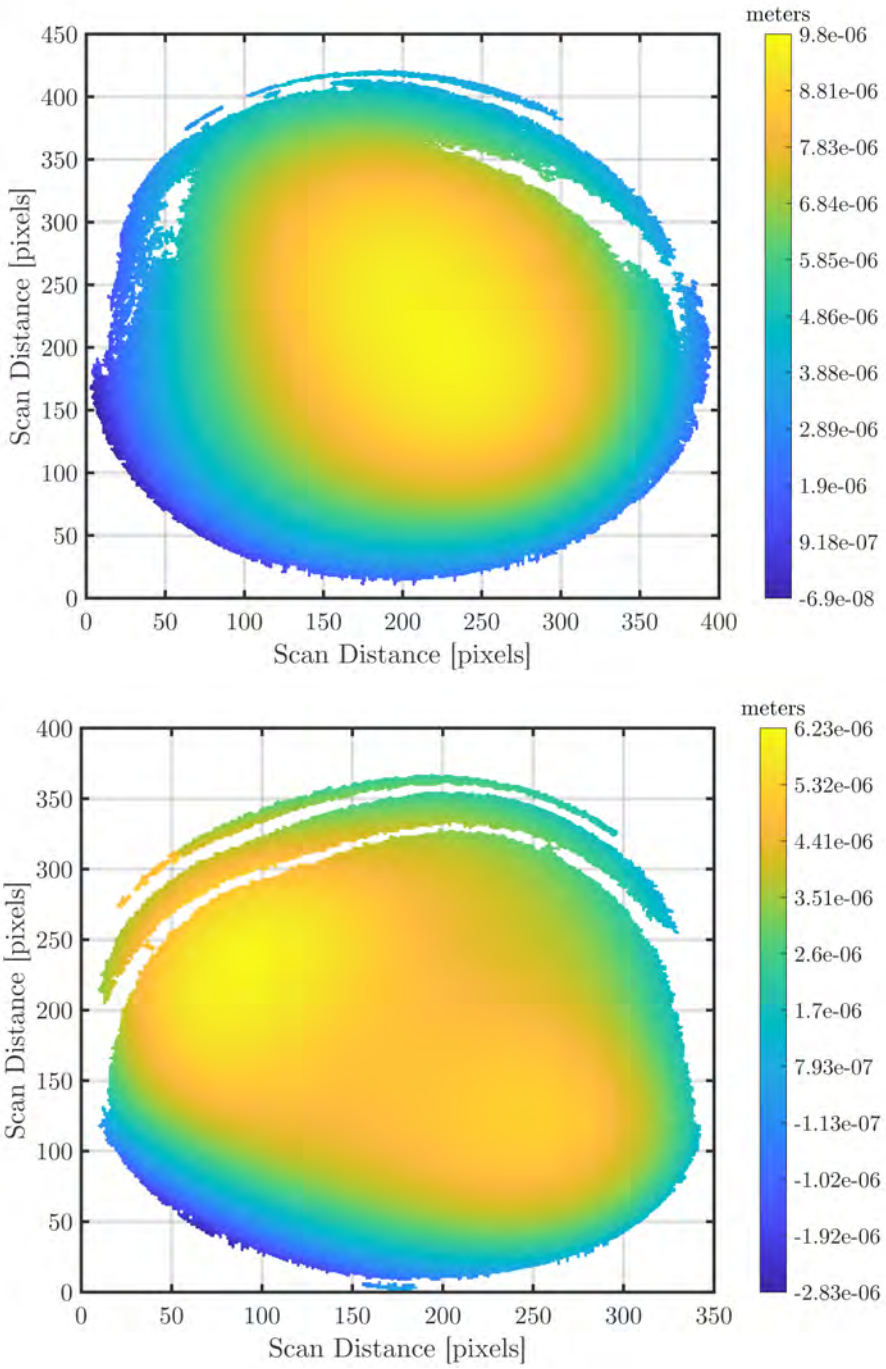

Figure 7.16: Surface map of the front(a) and rear(b) faces of a cSi disk ⌀ =2" (50.8 mm) measured using a ZYGO GPI-XP interferometer [240].

The surface form factor shows the 'front' face has an average peak to valley height of ± 100 nm with an almost Gaussian distribution of thickness across its radius. The 'rear' face presents a much larger average spread of ± 200 nm. The scan

reveals a 'bi-nodal', variation in the sample thickness, with two regions of increased thickness across the diameter of the sample. The total variation in the surface also shows that regions towards the edge of the sample were measured at a more considerable distance from the ZYGO compared to its centre, suggesting the regions are thinner. It can be speculated that such surface inhomogeneities could contribute higher surface losses to modes which have radial deformation anti-nodes such as Second Family 'Drum modes' ($m \geq 0, n \geq 0$) or Third family modes ($m \geq 1, n \geq 0$).

During cryogenic measurements of these samples, it was imperative to ensure that the measured values of $\phi_{\mathrm{mech}}(f, \mathrm{T})$ are unaffected by temperature gradients throughout the sample. Loss measurements were started after the temperature of the GeNS Body sensor had stabilised ($\frac{d\mathrm{T}}{dt} < 0.1\,\mathrm{K\,hr^{-1}}$). From the initial temperature calibration run (see section 7.3) it was known that when this sensor reached stability, the disk temperature was also stable to a similar degree.

7.4.1 Thermoelastic loss of cSi Disks

As a sample deforms during vibration the coupling of the stress and temperature fields inside the system allow for energy dissipation. Conversion of mechanical energy into an irreversible heat flow is a source of dissipation. This effect - thermoelastic loss - is particularly prevalent in close-packed metallic materials such as silicon which is of great interest for third-generation cryogenic gravitational wave detectors [78, 104].

Fig. 7.17 shows an example of heat flow across a body under deformation during oscillation at resonance. The frequency dependence of this loss is related to the characteristic time taken for heat to flow though the thickness of the body τ.

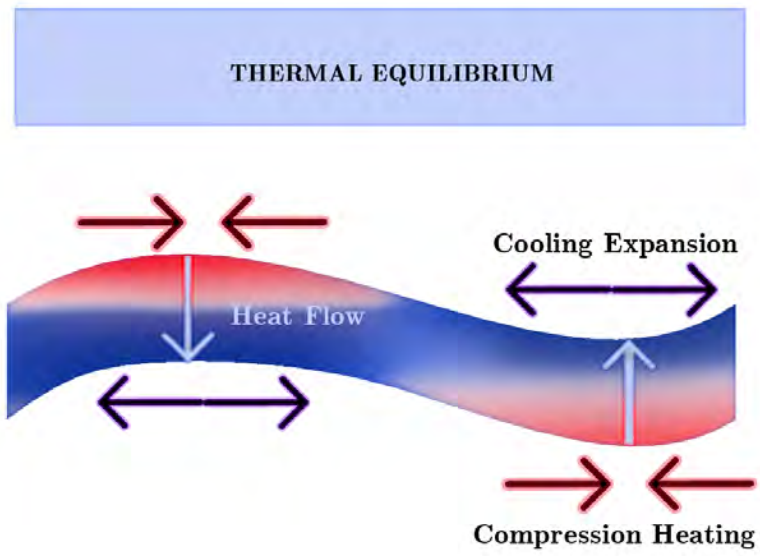

Figure 7.17: Illustration of thermoelastic heat flow across an oscillating body. Highlighting the temperature gradients created by thermal expansion and contraction compared to a body in thermal equilibrium.

As work in later chapters of this thesis rely on the thermoelastic contributions of cylinder of thickness t, the formulaton of ϕ_{therm} by G. Cagnoli *et al.* [135] can be used:

$$\phi_{D,\text{therm}} = \mathbf{D}_{\text{dil}} \cdot \phi_{\text{therm}} = \mathbf{D}_{\text{dil}} \cdot \frac{(3\alpha)^2 K T_0}{\mathrm{C}_P \nu} \frac{\omega \omega_{\text{peak}}}{\omega^2 + \omega_{\text{peak}}^2}. \tag{7.6}$$

Here $\mathbf{D}_{\text{dil}}$, the 'dilution factor' is the ratio of the elastic strain energy stored in bulk motion to that stored in shear motion, $\frac{E_{\text{bulk}}}{E_{\text{Total}}}$. By including the ratio of $\frac{E_{\text{bulk}}}{E_{\text{Total}}}$, the total change in volume for each mode shape is taken into consideration. 3α is also employed as the volumetric thermal expansion coefficient of the material. K is the isotropic bulk modulus of the material, T_0 is the environmental temperature of the cylinder, C_{P} is the heat capacity of the material at constant pressure multiplied by ρ the density.

As was shown by Zener and others [241–246] the magnitude of thermoelastic damping exhibited in an oscillating body is defined by the frequency of the oscillating body compared to the dissipation rate τ. When the frequency of the oscillating strain field ω and τ are equivalent, the heat flow and body deformation are out of phase

producing the maximum level of internal friction. This phase relationship between the heat flow and deformation decreases exponentially as ω moves further from τ. The frequency at which the maximum phase difference between these two quantities occurs is

$$\omega_{\text{peak}} = \frac{\kappa}{\mathrm{C}_P \rho}\left[\left(\frac{k\pi}{t}\right)^2 + \left(\frac{\mu_m^{(n)}}{r_0}\right)^2\right] \tag{7.7}$$

at a frequency denoted by ω_{peak}. In this formulation $\mu_m^{(n)}$ is the derivative of a Bessel function J_n of order n where m denotes the order of zero from the function [135] and κ is the thermal conductivity of the material. In the case of a thin cylinder, such that $t << r_0$ the calculation of ω_{peak} can be simplified such that $J_n'\left(\mu_m^{(n)}\right) = 0$ and $k = 1$:

$$\omega_{\text{peak}} = \frac{\kappa}{\mathrm{C}_P \rho}\left(\frac{\pi}{t}\right)^2. \tag{7.8}$$

Given that the material properties of silicon vary with temperature [230, 247], these must also be taken into account. The thermal conductivity, κ is mostly dependant on the purity of silicon [247, 248], the level of doping, and can vary by up to 70% at low temperatures. Fig 7.18 shows the predicted $\phi(f,T)_{\text{thermoelastic}}$ for an uncoated cSi $\varnothing = 2$" (50.8 mm), $t = 360\,\mu$m disk at 1000 Hz. This calculation takes into account the temperature dependence of C_P, ν, Young's Modulus, α and κ and its variance on doping.

For this experiment, which studied mechanical loss between 80 K- 300 K, the uncertainty due to material doping will be small, with a slightly larger effect at the lower end of the temperature range. However, as multiple mode frequencies were measured in this temperature range, variations in both frequency and $\mathrm{E}_{\text{bulk}}/\mathrm{E}_{\text{shear}}$ for different modes are of consequence. Figure 7.19 shows $\mathrm{E}_{\text{bulk}}/\mathrm{E}_{\text{shear}}$ for a $\varnothing$=50.8 mm $\times$ $t = 360\,\mu$m <100> Si disk calculated using COMSOL. Each set of mode shapes or 'families' are shown to follow distinct trend with their value of $\mathrm{E}_{\text{bulk}}/\mathrm{E}_{\text{shear}}$.

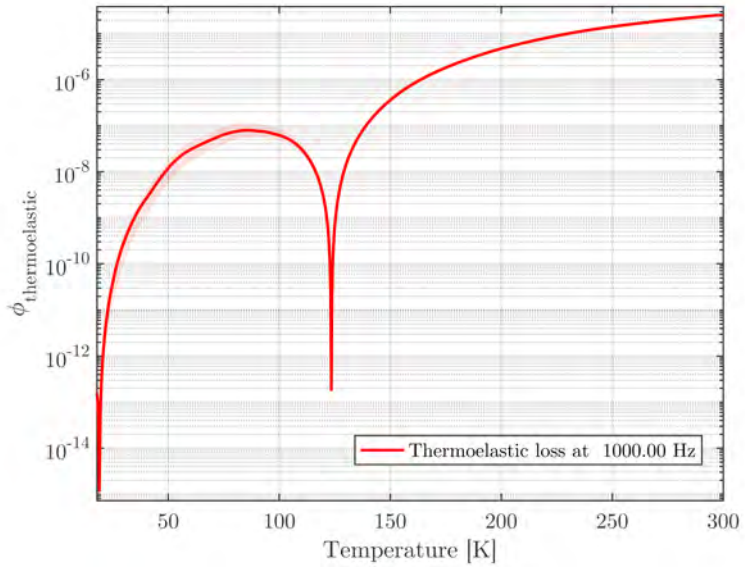

Figure 7.18: Upper and lower bounds for thermoelastic loss of a $\varnothing = 2$" (50.8 mm), $t \approx 360\,\mu$m calculated at 1000 Hz at temperatures between 20 K and 300 K. Dilution factor $\mathbf{D}_{\text{dil}}$ =0.047 is used, typical for a typical (2,0) mode at 1 kHz.

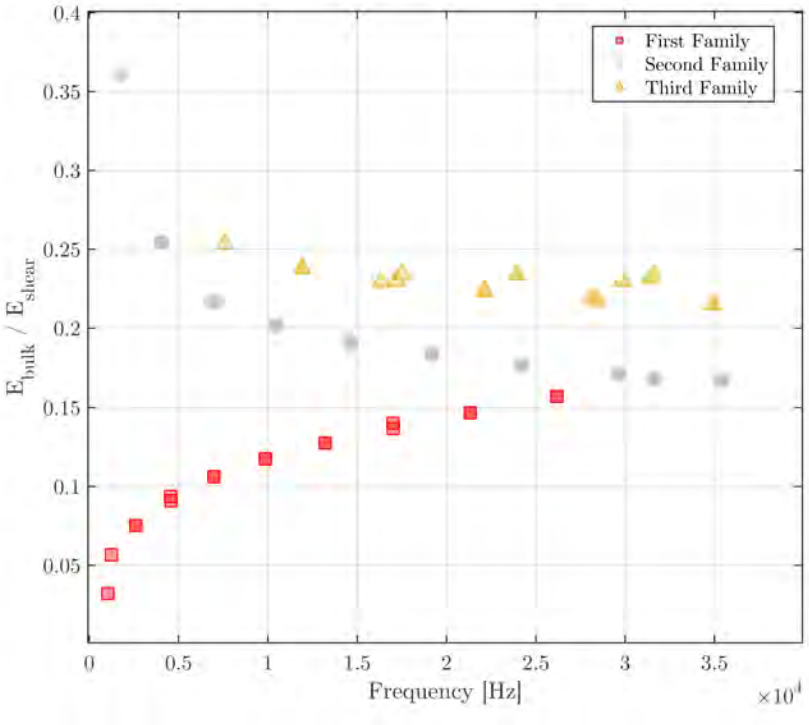

Figure 7.19: Calculated values of E_{bulk}/E_{shear} for a $\varnothing = 2$" cSi disk as a function of frequency, mode families are differentiated by marker.

Using these values in the calculation of thermoelastic loss shows that as the frequency of modes increase, the thermoelastic loss of that mode will be scaled differently depending on its mode family (see Chapter 5 for more information on disk mode families). As a result, the thermoelastic loss of each mode family will exhibit different frequency-dependent behaviour (see Eq. 7.6). This is illustrated in Figures 7.20, 7.21 and 7.22, showing the upper and lower bounds of $\phi_{\text{thermoelastic}}$ for each mode family.

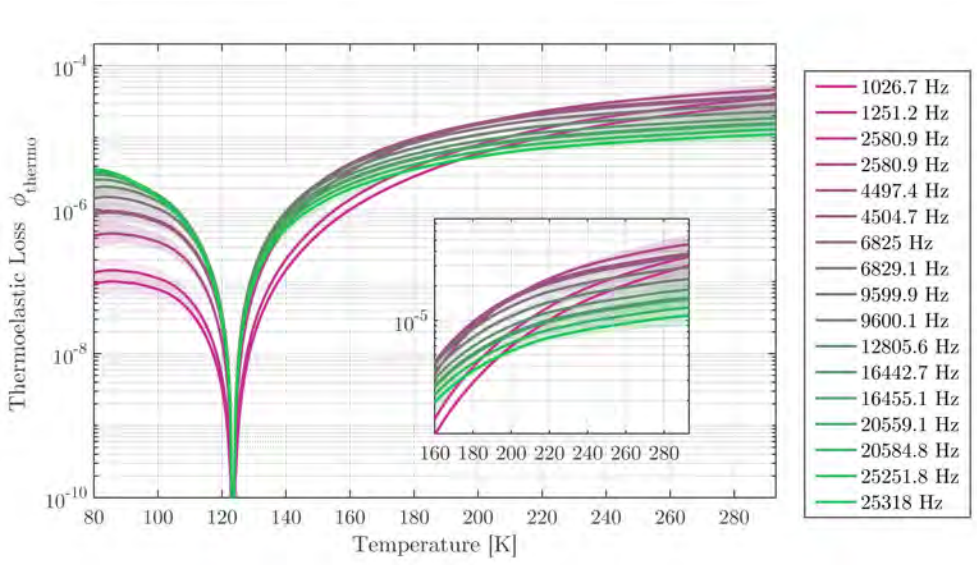

Figure 7.20: Thermoelastic loss of (n,0) First Family modes for a silicon disk $\varnothing = 2$" (50.8 mm), t ≈360 µm calculated as a function of frequency at temperatures between 77 K and 300 K. The inset figure shows regions where curves intersect.

It can be seen from these calculations that the differences in shear and bulk energies by mode family and frequency lead to different magnitudes of thermoelastic loss in each case. At temperatures below 120 K, Third Family modes have the largest thermoelastic loss, but the smallest variation with frequency due to the largely consistent distribution of $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$. At temperatures above 120 K, the thermoelastic loss of Second Family mode frequencies varies over an order of magnitude (9×10^{-5} to 9×10^{-6}), with the range of First family showing less variation, and third family modes less still. These thermoelastic effects are likely to lead to the loss of each mode family being dependent on frequency, $\frac{E_{\text{bulk}}}{E_{\text{shear}}}$ and temperature.

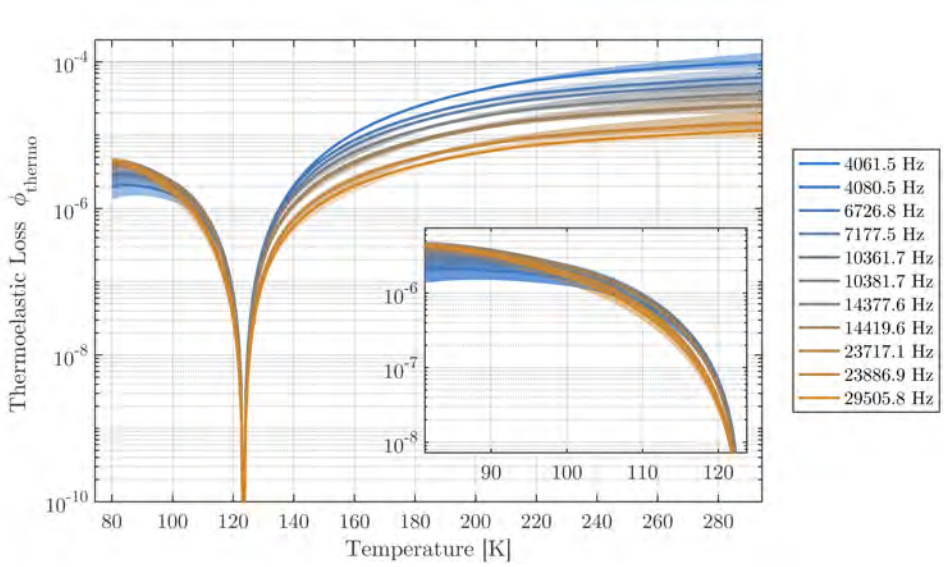

Figure 7.21: Thermoelastic loss of (n,1) Second Family modes for a silicon disk $\varnothing = 2$" (50.8 mm), $t \approx 360$ µm calculated as a function of frequency at temperatures between 77 K and 300 K. The inset figure shows regions where curves intersect.

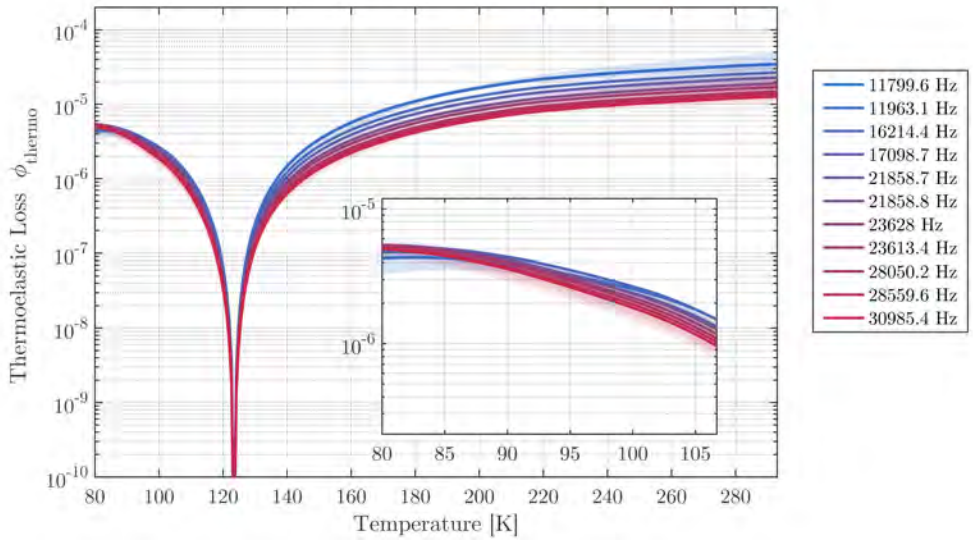

Figure 7.22: Thermoelastic loss of (n ,0) First Family modes for a silicon disk $\varnothing = 2$" (50.8 mm), $t \approx 360$ µm calculated as a function of frequency at temperatures between 77 K and 300 K. The inset figure shows regions where curves intersect.

7.4.2 Cryogenic Measurements of an uncoated cSi Disk

The first measurements of $\phi(f, \mathrm{T})_{\mathrm{mech}}$ using the CryoGeNS were carried out using an uncoated cSi disk, measured between 80 K and 300 K. The system was allowed to cool to a stable baseplate temperature of 82.5 K and the system incrementally heated to a new set-point temperature and left to stabilise for more than 4 hrs, before the loss was measured. $\phi(\mathrm{T})$ was measured for mechanical resonances between 1 kHz and 32 kHz to allow for any dependencies between $\phi(\mathrm{T})$, mode shape and frequency to be investigated.

Fig 7.23 shows the average loss of several modes measured over four ringdowns for each mode. The error-bars representing the max-min deviation. Mechanical modes which exhibit 'coupling' (i.e. the interaction of ringdowns between one or more mode frequencies) have been removed for clarity. It would be expected that the measured mechanical loss values for each mode will follow $\phi_{\mathrm{thermoelastic}}$ of the silicon substrate, reaching a minimum at $\mathrm{T} \approx 120$ K.

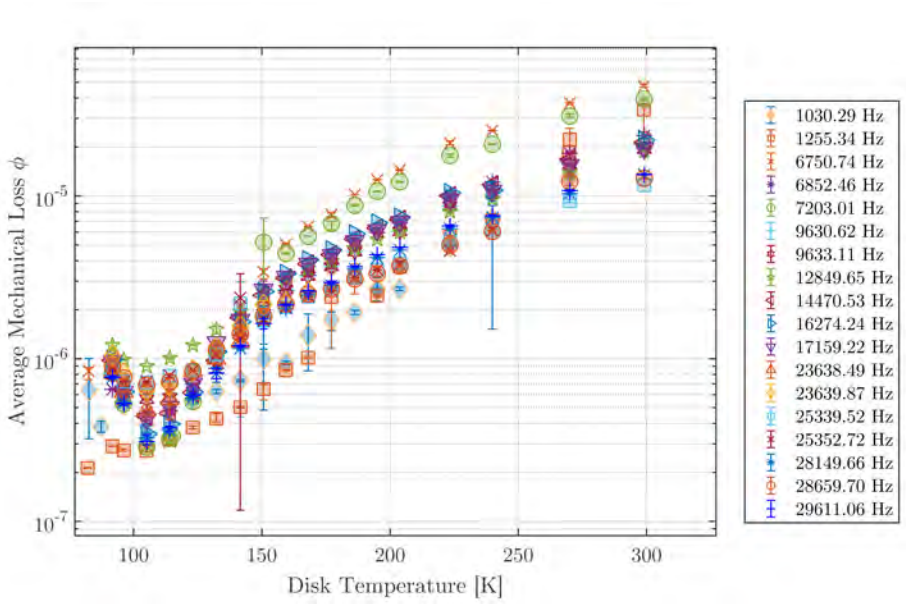

Figure 7.23: Mechanical loss of an $\varnothing = 2$" uncoated cSi disk between 80 K and 300 K measured using the CryoGeNS system.

For all measured modes the mechanical loss decreases with temperature, in most cases limited by the level of $\phi_{\mathrm{thermoelastic}}$ at that frequency, showing a minimum loss at an average disk temperature of 106 K - somewhat lower than the expected 123 K

(determined by the zero in thermal expansion at this temperature). For each mode family the same effect is observed suggesting either one of two possible explanations:

- there is a difference in the surface/bulk loss of the sample, which is altering the functional form of the measured loss, or
- the sample temperature is higher than inferred via the temperature calibration process (see section 7.3).

7.4.2.1 Surface Losses

In this section possible contributions to the functional form of $\phi_{\text{mech}}(T)$ will be considered, focusing on surface losses from either circular face of the sample and ϕ_{barrel}, a source of loss from the samples circular edge [249].

From the work of Granata *et al.* [142] on $\varnothing = 3$" silica disks, it can be hypothesised that the contribution of ϕ_{barrel} is the ratio of the surface area of the sample and the surface area of the barrel scaled for each mode shape. For ϕ_{barrel} to dominate, if the ratio of the surface area of the barrel A_{barrel} to the circular faces A_{circ} is more than 1/56, the polishing quality of its surface significantly impacts the loss of the sample. In Granata's method, if the loss of First Family modes at a fixed temperature can be assumed by the function $af^b + \epsilon d\phi_{\text{barrel}}$ the difference between the loss of these modes and those outwith the trend are contributions of ϕ_{barrel}. Inspection of $\phi_{\text{mech}}(f)$ for each temperature step show no such relationship. This suggests that the sample does not conform to the same frequency dependence, indicating that ϕ_{barrel} does not impact the functional form of $\phi(f, \text{T})_{\text{mech}}$ measured on this substrate.

This then leaves contributions from the circular faces of the disk as another influence which could effect the functional form of $\phi(f, \text{T})_{\text{mech}}$. To investigate this premise, the loss of the uncoated substrate must be analysed at temperatures where $\phi_{\text{thermoelastic}}$ has no significant contribution to the measured losses.

In each of the following plots, $\phi(f, \text{T})$ has been presented for each mode family. As the surface deformation of each family generates nodes and anti-nodes at different points on the circular surface, if a clear difference in loss between each

family is found, this could indicate localised surface loss contributions.

First-Family modes between 1 kHz and 10 kHz are shown in Fig 7.24, where frequencies exhibiting multi-modal interactions (MMI) have been removed. This mode family has vibrational anti-nodes close to the edge of the disk, with minimal deformations close to the disks centre. The mechanical loss of frequency shows a steady decrease in loss with temperature between 150 K and 300 K. Between 100 K and 115 K the measured loss inflects reaching a minimum of $\phi = 2.7\times10^{-7}$ at 106 K. This change in loss is far less pronounced compared to mode shapes which have less surface deformation at the circumference of the sample.

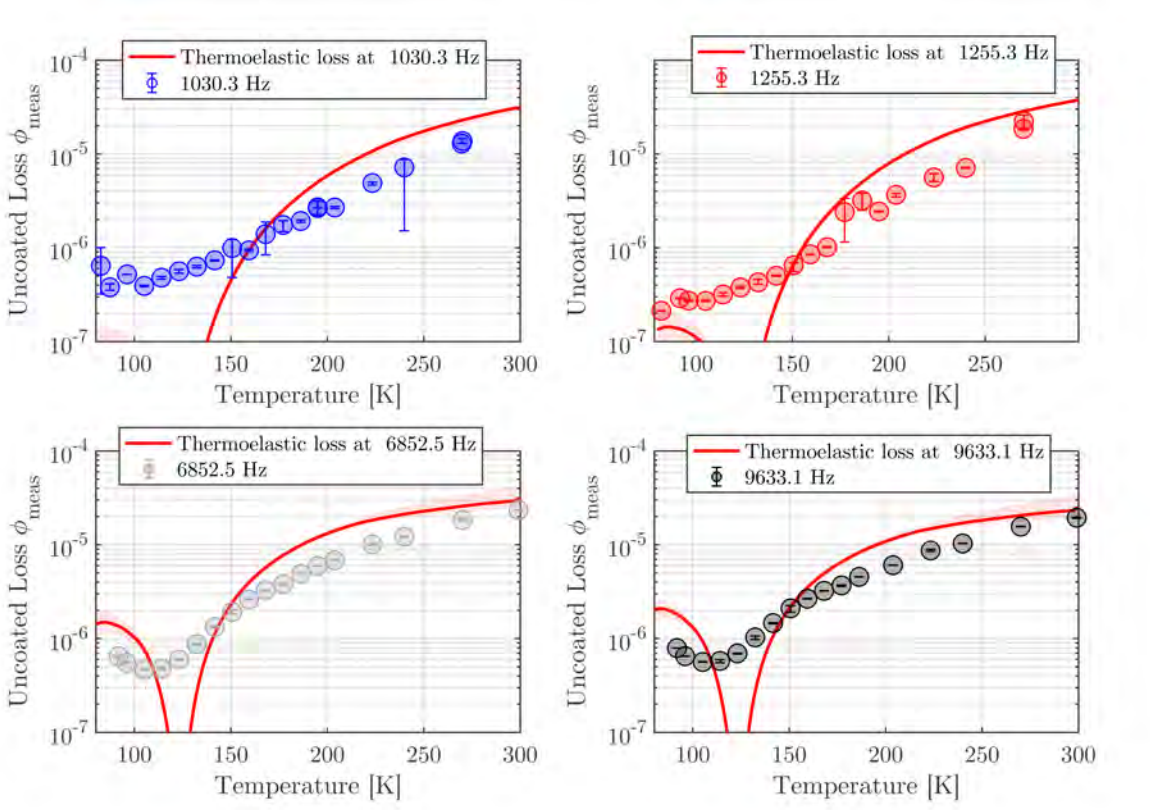

Figure 7.24: Mechanical loss of First Family modes on an $\varnothing = 2$" uncoated cSi disk between 80 K and 300 K. Each set of measured data is compared to predictions for thermoelastic loss.

Figure 7.25 shows values of $\phi(f, T)$ measured from Second Family modes which have a larger contribution of volumetric deformation. Due to the nature of the measurement technique, (0,1) and (0,2) modes could not be measured as these are damped by the balancing lens. As was shown in Fig 7.21 the predicted

$\phi(f,T)_{\text{thermoelastic}}$ for Second Family modes is much higher at temperatures 120 K compared to First Family modes.

At temperatures below 130 K the average loss of all non-coupling modes decreases with a much sharper gradient compared to the level of loss measured on First Family modes in this temperature range, with both families reaching approximately the same minimum loss ($\phi = 2.8\times10^{-7}$ and $\phi = 2.6\times10^{-7}$) at 106 K.

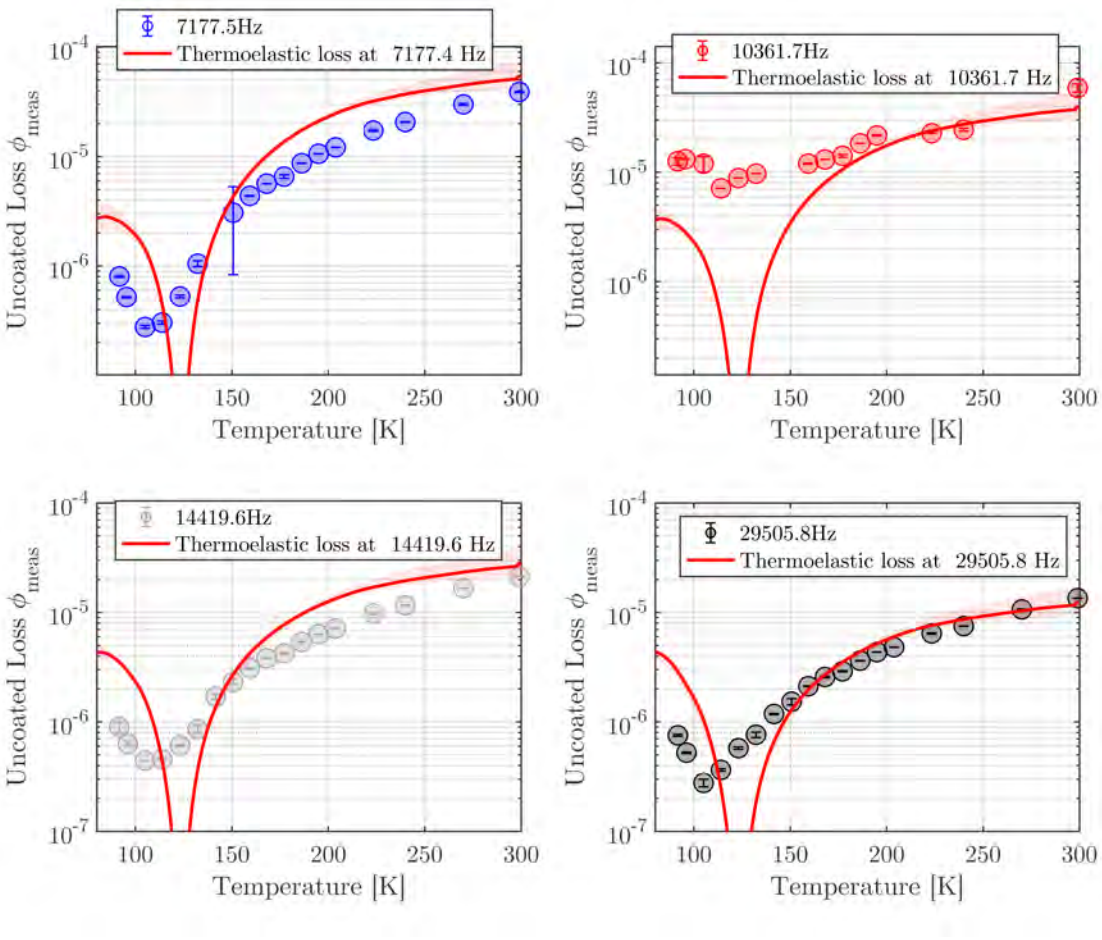

Figure 7.25: Mechanical loss of Second Family modes on an $\varnothing = 2$" uncoated cSi disk between 80 K and 300 K.

Viewing Third Family modes in the same manner (see Fig 7.26), a very similar trend in loss to modes from the Second Family is found. As these mode shapes are not produced on this geometry at f below 12 kHz there are fewer measurements of such modes. Above 120 K, the measured loss follows the expected functional form

of $\phi_{\text{thermoelastic}}$, but is lower than predictions in most cases. At temperatures below 120 K a trend in loss akin to Second Family modes is found, with the lowest losses produced at 6.7 kHz of $\phi = 3.3\times10^{-7}$ at 106 K.

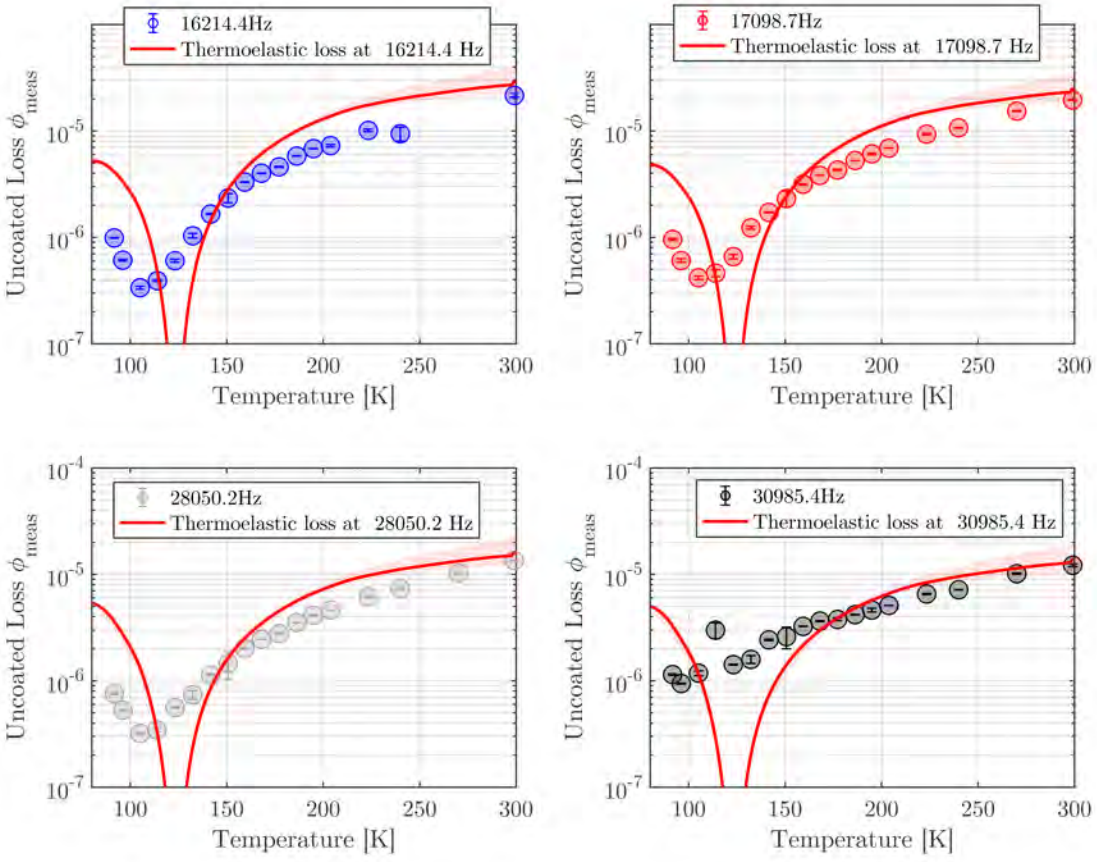

Figure 7.26: Mechanical loss of Third Family modes on a $\varnothing = 2$" uncoated cSi disk between 80 K and 300 K.

Differences in the surface quality of each sample (shown in Fig 7.16) shows large variations in thickness across the surface of each uncoated disk. By analysing the lowest measured losses and their respective temperatures from each mode family a compelling difference in loss as a function of mode family is suggested. The lowest losses from First and Second mode families are both measured at 106 K and agree within 4%. While the lowest loss from Third Family modes also occurred at the same temperature, the loss was measured to be >20% higher. This could suggest that this mode shape is subject to additional damping which does not effect other modes. This result suggests that the lowest losses measured on this disk could be

limited by its surface polishing quality. However, as the lowest loss for each mode shape are consistently measured at 106 K it seems unlikely that this affects the functional form of $\phi(f, \mathrm{T})$.

In previous sections, modes which showed no clear signs of interactions with other frequencies were described. This produces an acceptable description of the behaviour of $\phi(f, \mathrm{T})_{\mathrm{mech}}$, but does not entirely describe the sample. With a view to producing such a description, all modes must be analysed and any uncovered discussed.

Section 7.4.2.1 revealed that it was convenient to view $\phi(f, \mathrm{T})_{\mathrm{mech}}$ in terms of mode shape. Each mode which shows significant signs of MMIs is itemised in Tab. 7.1. These modes can add effects such as beating in the ringdown amplitude or artificially higher measured loss for interacting modes. Throughout the measured temperature range, the material properties of the substrate change, altering the resonant frequency of each mode. This change can cause modes to interact at certain temperatures, and not at others. Fig 7.25 and 7.26 details measurements of $\phi(f, \mathrm{T})_{\mathrm{mech}}$ from Second Family and Third Family modes, compared to the predicted $\phi_{\mathrm{thermoelastic}}$.

The temperature dependence of the average mechanical loss for each mode family shows that not all mode shapes respond in the same manner as the system changes temperature. Comparing measurements from different mode families, there is a discrepancy between the ϕ_{meas} and thermoelastic predictions throughout the measured temperature range. However, it is interesting to note that despite a similar number of degenerate mode pairs being measured from each mode family, the average loss of the Second Family modes are still consistently higher than the other two families.

As the sample temperature approaches the dip in $\phi_{\mathrm{thermoelastic}}$ at $\sim$120 K the average loss of Third Family modes also decreases below the average of First Family modes, with the most considerable difference in average loss of 53% at 106 K. Given that the number of coupling modes within the measured frequency range are similar for each mode family, but the total spread in measured losses differs, the cause of the measured spread still remains. By comparing modes from each family which do

not exhibit MMIs, the total variation in loss with frequency can be compared directly. The average loss and spread of 11 modes were measured from both the Second and Third mode family and 13 First Family modes are shown in Figure 7.27.

Measured	$\frac{E_{\text{bulk}}}{E_{\text{shear}}}$	(m,n)	Mode Family
1030.9	0.047	(2,0)	First Family
1251.3	0.056	(2,0)	First Family
2580.9	0.089	(3,0)	First Family
4061.5	0.277	(1,1)	Second Family
4497.4	0.109	(4,0)	First Family
4504.7	0.102	(4,0)	First Family
10361.7	0.225	(3,1)	Second Family
10381.7	0.225	(3,1)	Second Family
11935.7	0.259	(1,2)	Third Family
11975.4	0.259	(1,2)	Third Family
14377.6	0.210	(5,1)	Second Family
16442.7	0.153	(8,0)	First Family
16455.1	0.149	(8,0)	First Family
20559.1	0.156	(9,0)	First Family
20584.8	0.156	(9,0)	First Family
21937.2	0.244	(3,2)	Third Family
21961.3	0.244	(3,2)	Third Family
30985.4	0.253	(2,3)	Third Family

Table 7.1: Summary of mode frequencies measured on a $\varnothing = 2$" uncoated cSi disk which exhibit multi-mode interactions.

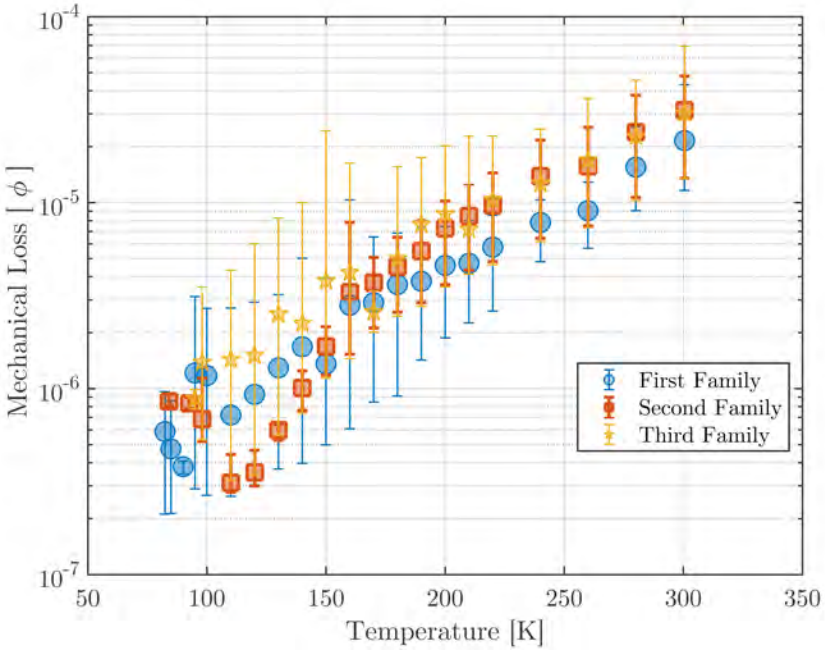

Figure 7.27: Average mechanical loss of a $\varnothing = 2$" uncoated cSi disk at each temperature step eliminating coupling modes. Error bars denote the max-min variation of ϕ_{mech} for each mode family.

This relatively large spread in mechanical dissipation for First and Third family modes suggests that there is an introduction of excess damping which affects only these mode families. A comparison of the surface deformation of mode-shapes from each family shows that regions of deformation during vibration for Second family modes manifest as radial anti-nodes, with smaller interactions on the circumference of the disk. First Family modes have large nodes around the circumference of the disk, with Third Family modes exhibiting both circumferential and radial nodes, which are much smaller in comparison to those found on other families.

7.4.2.2 Temperature Stability and Thermoelastic Predictions

To assess if inaccurate estimations of disk temperature could distort the expected functional form of $\phi(f, \mathrm{T})_{\text{mech}}$ compared to the expected $\phi(\mathrm{T})_{\text{thermoelastic}}$, the measured trend of Δf can be used. The temperature dependence of each mode frequency is related to both the material properties of the sample, the type of mode shape and its orientation relative to the cSi crystal structure. As a result the frequency spacing

between previously separated mode pairs can change with temperature, and in some cases introduce coupling and create temperature dependent degenerate mode pairs.

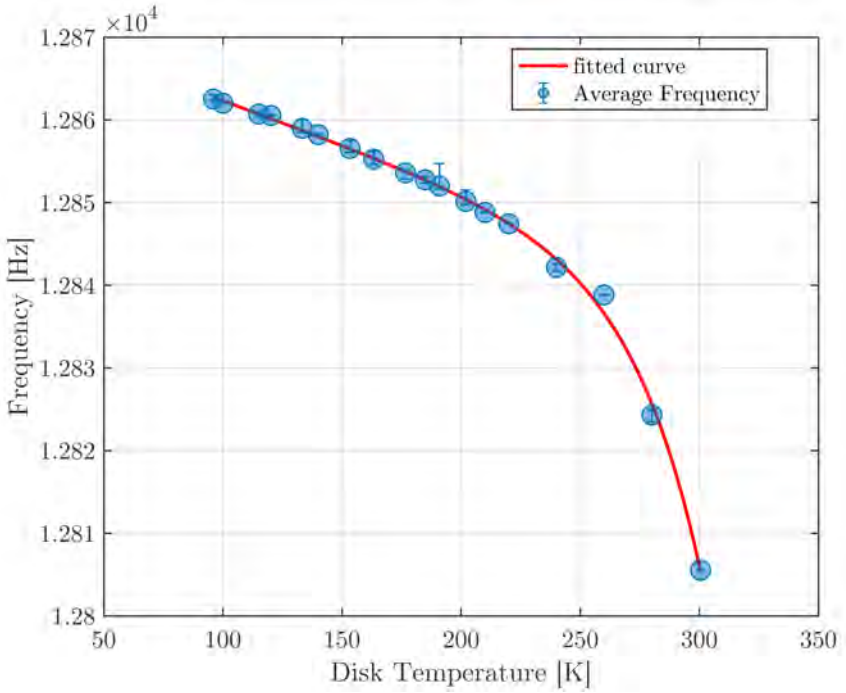

Figure 7.28: Average measured frequency of a 12.8 kHz mode from uncoated 2" cSi disk at stabilised temperature steps between 80 K and 300 K.

Fig 7.28 shows average measured frequency of a 12.8 kHz resonant mode as a function of temperature. This data is then fit with a two term exponential function dictated by the temperature dependence of silicons Young's Modulus. In each temperature step the measured frequency of the 12.8 kHz mode changes, on average by <0.5 Hz, with a maximum change of 3.1 Hz between 170 K and 180 K. Agreement with the exponential function validates that the change in Young's modulus is responsible for the observed change in frequency. The maximum measured deviation is coincident with a point in the measurement run where the cryostat's liquid nitrogen reservoir was running low, allowing the experimental chamber to warm more than desired. The liquid nitrogen was topped up, and the next measurements at 200 K shows the second-largest change in frequency of 0.99 Hz, indicating that the system did not reach a state of thermal equilibrium before measurements were resumed. Analysis of the frequency shift of each mode around this period shows that the temperature stability of the disk is correlated with the level of liquid nitrogen in the cryostat. This result then suggests that the stability of the system

at T>200 K could then affect the expected functional form of $\phi(f, \mathrm{T})$, this is discussed in more detail in section 8.3.

The sample's behaviour at lower temperatures must be explained. The zero thermal expansion of cSi at $\sim$123 K, leading to minimum in thermoelastic loss, is well known. The minimum loss being measured at lower temperatures most likely suggests that the inferred disk temperature is incorrect. This possibility was investigated by comparing theoretical $\phi_{\mathrm{thermoelastic}}$ predictions to measurements from an independent system.

Fig 7.29 shows the comparisons of the theoretical calculation of $\phi_{\mathrm{thermoelastic}}$ carried out by the author and measurements of $\phi(f, \mathrm{T})_{\mathrm{mech}}$ carried out on similar substrates ($\varnothing$ =3" (76.2 mm) $\times$ 459 µm) by M. Granata et al. [250]. Four modes were measured as a function of temperature, as the temperature approaches the thermoelastic zero, there is a reduction in mechanical loss. However, unlike in the data measured by the author, the mechanical loss measured here does not reduce below thermoelastic predictions calculated using Eq. 7.6.

As the exact composition of the sample is not known, these deviations could be produced by variations in material properties which are not taken into account and could affect the magnitude of the predicted thermoelastic at T<120 K. As no equivalent thermoelastic calculations were produced by [250] no comparison between theoretical predictions can be produced.

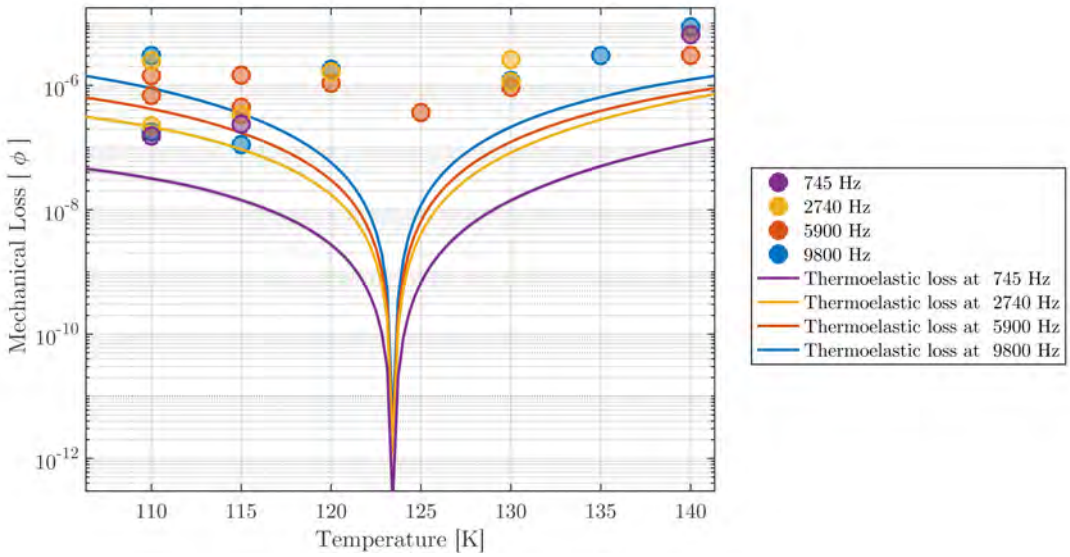

Figure 7.29: Mechanical loss carried out by M. Granata *et al.* of an $\varnothing = 2" \times 459\,\mu$m thick uncoated cSi disk between 110 K and 140 K each with thermoelastic predictions.

The fact that all measurements of all mode families show this discrepancy, and that the measured losses have a similar functional form to the predicted thermoelastic loss, suggests that the sample temperature is actually higher than is being estimated by the temperature calibration - perhaps due to additional conductive cooling via the applied varnish during the calibration run.

In the next chapter, this discrepancy is investigated further using a GeNS system installed in a Pulse Tube cooled cryostat. This apparatus does not require cryogenic liquid cooling and thus has greater control over temperature stability allowing the discrepancy between $\phi_{\text{thermoelastic}}$ and measured ϕ_{mech} in this chapter to be scrutinised.

7.5 Conclusions

The GeNS apparatus developed for room-temperature loss measurements has been successfully developed to operate at cryogenic temperatures. Measurements of uncoated cSi disks were carried out between 80 K and 300 K, to test the apparatus and provide reference samples for future coating measurements. The measured loss of the samples followed the predicted thermoelastic trend at temperatures above 140 K. However at temperatures below 140 K there is disagreement between the

measured $\phi(f, \mathrm{T})$ and predicted thermoelastic losses which is consistent for each mode family and could suggest temperature instabilities in the system.

Comparison to loss measurements produced on other systems shows that the magnitude of $\phi(f, \mathrm{T})$ of equivalent samples is analogous and validates the experimental methods provided by the author. However, the fact that measurements of $\phi(f, \mathrm{T})$ from all mode families show this discrepancy in temperature, and that the measured losses have a similar functional form to the predicted thermoelastic loss, suggesting that the sample temperature is actually higher than is being estimated by the temperature calibration procedure. Due to experimental restrictions after these measurements were completed, the temperature calibration could not be repeated.

Chapter 8

Cryogenic Mechanical Loss of Multi-Material Coatings

As discussed in Chapter 2, Brownian thermal noise can be reduced by lowering the mirror temperature as done in KAGRA [73] and is planned for several future detectors [108, 198, 201]. However improved coatings are required for these detectors to meet their sensitivity goals. In this section cryogenic mechanical loss measurements of the prototype mulit-material coating already discussed in Chapters 4 and Chapter 5. Three coating materials known as the *Full Stack* (Ta_2O_5-SiO_2/SiO_2-aSi), *Lower Stack* (SiO_2-aSi) and *Upper Stack* (SiO_2-Ta_2O_5) deposited on cSi cantilevers and $\varnothing$=2" cSi disks (see section 7.2.1) were studied.

8.0.1 Single Layers on Cantilever Substrates

As measurements of coating loss are traditionally carried out using coated silicon cantilevers, cantilever samples coated in single layers of RLVIP deposited SiO_2, aSi, Ta_2O_5 and samples coated in the *Upper*, *Lower* and *Full* Stack coatings were studied in parallel with *Upper*, *Lower* and *Full* Stack coatings deposited on disk substrates. Comparing these measurements for the different substrates at similar frequencies allows any differences in coating loss due to stress and substrate curvature to be gauged. This effect has been shown by M. Fletcher [203] to alter the level of loss when calculating ϕ_{coating} when the deformation is not accounted for.

Fig 8.1 shows the measured coating loss from three cantilever samples coated in single layers of RLVIP SiO_2, aSi, Ta_2O_5 between $\approx 4\,K$ and $293\,K$ by Peter Murray (using the measurement technique described in [93]). The coating loss of the Ta_2O_5 layer has the largest coating loss of all materials in the measured temperature range. In this data, a peak is observed at $\sim 60\,K$, behaviour indicative of the cryogenic loss peak for Ta_2O_5 [204]. The loss of SiO_2 shows a peak at $\sim 20\,K$ consistent with other studies of silica carried out in [105, 150]. The aSi layer shows consistently low loss, which decreases at lower temperatures. Also shown in Fig 8.1 is the coating loss measured from the *Full Stack* material, and a comparison to the expected loss of this coating, calculated using Eq. 6.5 based on the measured losses of its constituent materials.

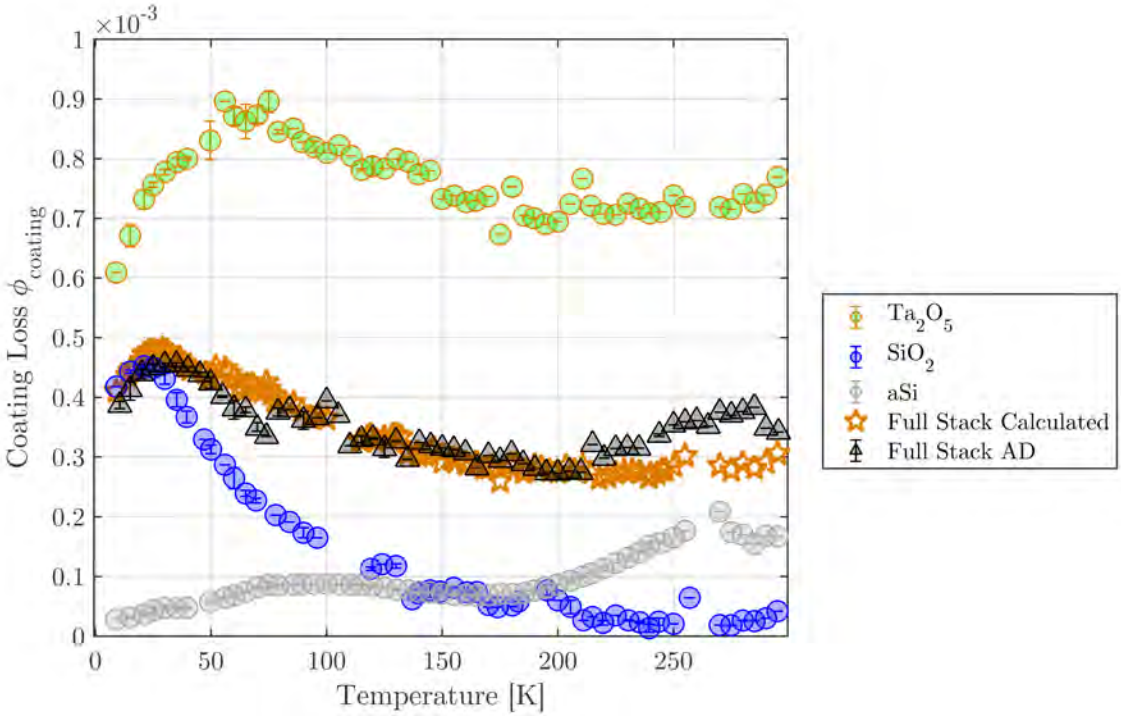

Figure 8.1: Coating loss of single layer RLVIP SiO_2, aSi, Ta_2O_5 and the *Full Stack* coatings in their as-deposited states, on cSi cantilevers measured as a function of temperature.

Between $210\,K$ and $300\,K$ the loss of the *Full Stack* is somewhat higher than predicted from the single-layer measurements, but excellent agreement is found in the rest of the temperature range. Below $200\,K$, the full stack coating loss increases

to a peak - arising from the silica layers - at around 20 K [204]. This cantilever data, taken in parallel with the results in this chapter, will be referred to throughout.

8.0.2 Losses of uncoated cSi Disks

Mechanical loss measurements of cSi disks were carried out using two separate apparatus, the cryogenic nodal support described in section 7.4 and a similar system contained within a closed-cycle, pulse cooled cryostat [251]. Due to its smaller and more accessible experimental space, a much simpler GeNS apparatus is used with the Si balancing lens, held atop an steel cylinder of comparable thickness. The temperature of a balanced sample is taken as the temperature measured on the steel post after the system has been left to stabilise for ≥ 6 hours.

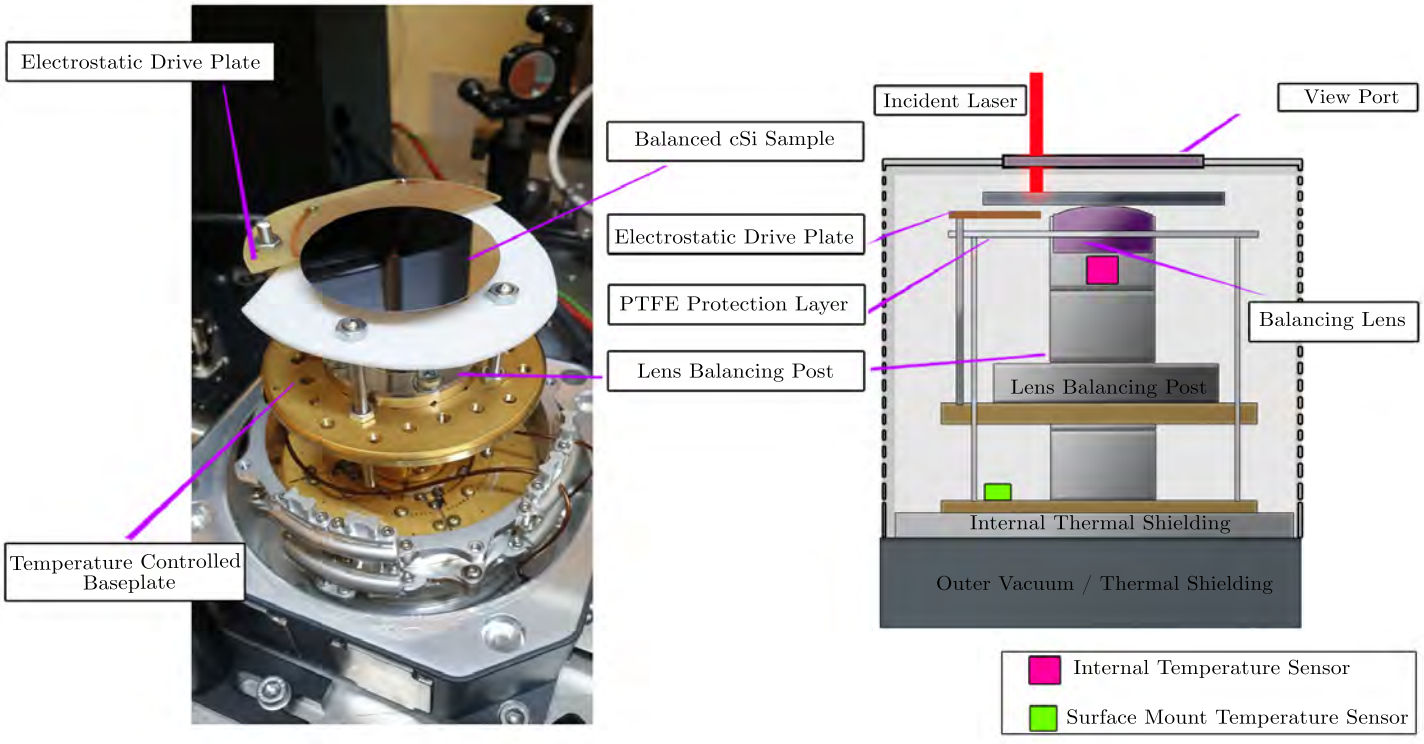

Figure 8.2: (LHS) Image of experimental space of the Montana Instruments Pulse Tube cooled cryostat. (RHS) Illustration of the experimental chamber, temperature sensor locations are labelled.

The ring downs of each mode of the disk was measured using a low power (≤ 5 mW) 633 nm laser reflected on the edge of the disk and onto an external

split-photodiode, this method is described in more detail in [93]. A rotatable polariser was used to attenuate the laser intensity by ≈70 %, thus reducing laser heating of the sample which are particularly prevalent at temperatures below 50 K. To test that the level of mechanical losses of this system produces values comparable to those previously measured using the CryoGeNS system, the ⌀ = 2" (50.8 mm), t ≈360 µm disk was measured in both systems - the results are shown in Figures 8.3(a-g)[1]. The Pulse Tube cryostat can produce data at temperatures <80 K, to cool the sample to an equivalent temperature using the CryoGeNS system would require liquid helium for which the system had not yet been tested.

Between 175 K and 300 K the magnitude of the loss for all measured modes should be limited by the thermoelastic loss of cSi however the loss in the majority of cases is below these predictions. This effect was discussed earlier in more detail in Chapter 5. At temperatures where thermoelastic loss is reduced (110 K - 140 K) the minimum mechanical loss measured differs between both systems. The discrepancy between the temperature at which the thermoelastic minimum is measured on the CryoGeNS system (discussed previously in section 7.2) is prevalent when compared to losses measured on the Pulse Tube system (see Fig 8.3(c)); where there is a difference of $\Delta T \approx 15$ K. At 120 K the lowest mechanical loss measured using the Pulse Tube cryostat was 53% lower (1.3×10^{-7}, $f = 7.4$ kHz) than the lowest measurements produced on the CryoGeNS system (2.8×10^{-7} $f = 7.1$ kHz). From the studies carried out in section 7.2 the differences between measured losses at each temperature on both systems are dependant on the how well the temperature of the system is calibrated.

The thermoelastic minimum is seen to occur at the expected temperature (∼123 K) in the Pulse Tube system, and at about 15 K lower in the CryoGeNS system, due to the previously discussed temperature calibration issue in this system (see section 7.2). Excellent agreement in the magnitude of the loss is obtained for around half of the modes: for some modes, the CryoGeNS yields significantly higher loss. This could indicate that there is additional frictional damping of these modes during this measurement.

[1]Throughout this chapter any stated resonant frequencies will be stated from their values measured/modelled at 293 K.

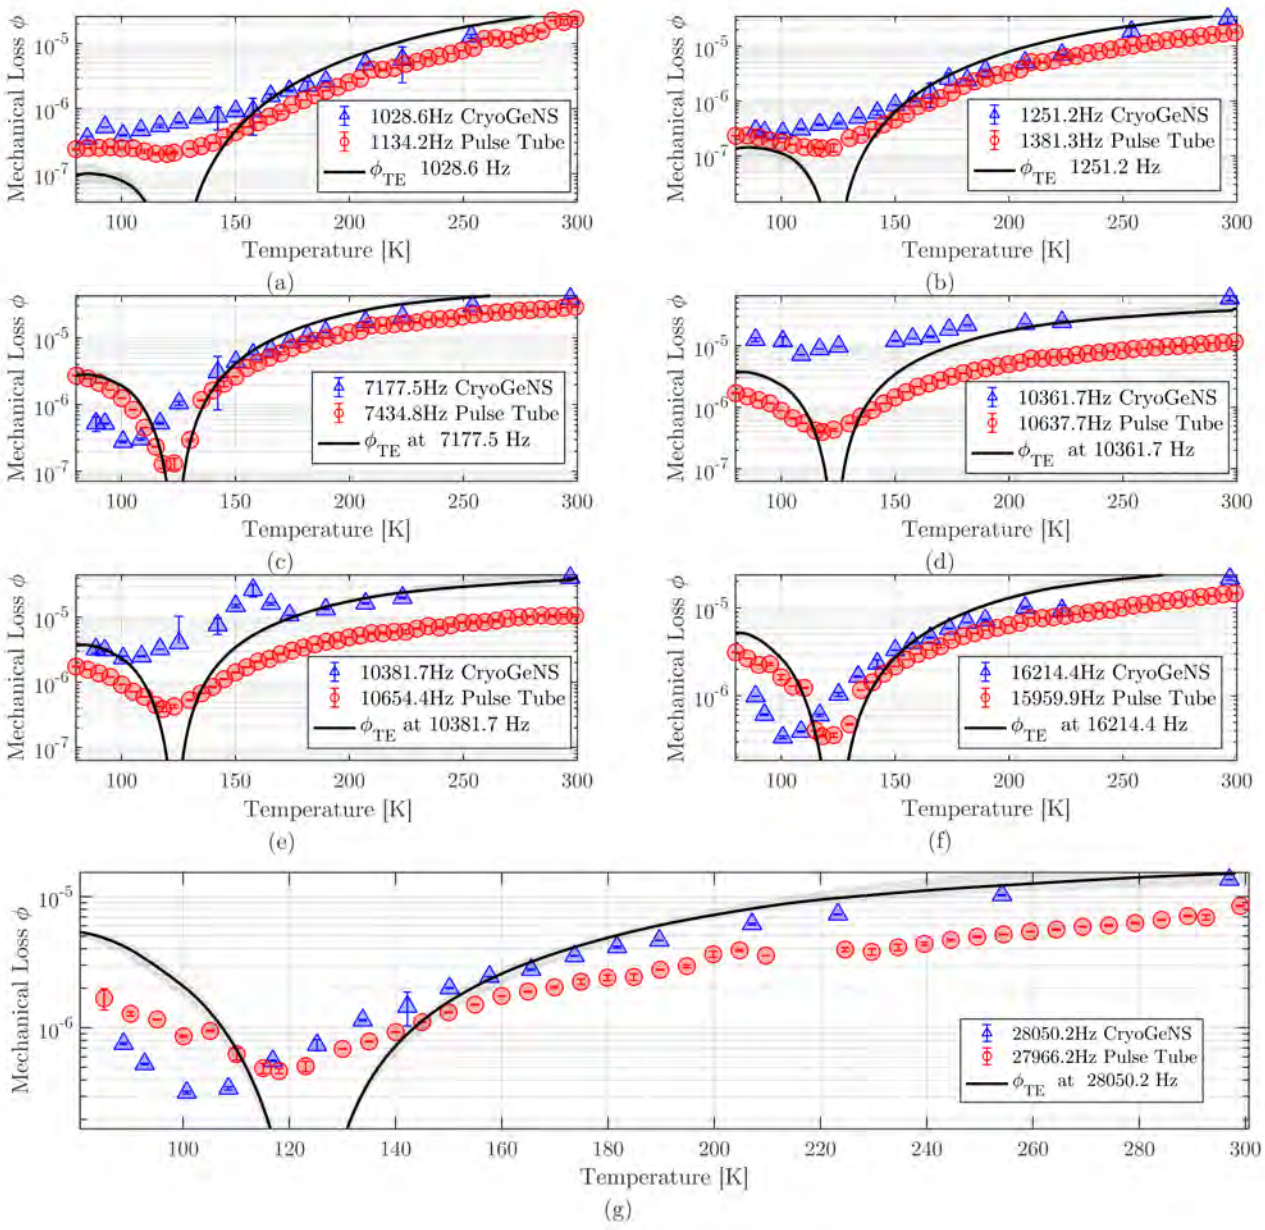

Figure 8.3: Comparison of mechanical loss from like modes on two uncoated $\varnothing = 2$" (50.8 mm), $t = \approx 360\,\mu$m cSi disks measured using the CryoGeNS and Pulse Tube cryostats. Sub-figures (a-g) are compared with the predicted thermoelastic for the disk measured in the CryoGeNS system.

8.1 Measurements of RLVIP aSi-SiO-$_2$Ta$_2$O$_5$

This section presents the mechanical loss of a $\varnothing = 2$" (50.8 mm), $t \approx 360\,\mu$m cSi disk coated with the *Full Stack* material (SiO_2-Ta_2O_5-aSi) as a function of sample temperature using two cryogenic systems. As the uncoated disk substrate was not measured before coating deposition, the loss of nominally uncoated substrates which were measured in each system are used.

Figure 8.4 shows the loss of the same (2,0) mode of the uncoated and coated samples. This figure also details the calculated loss of the *Full Stack* coating using Eq 5.11 with respect to the uncoated disk which was measured in each cryostat. The coating loss from the *Full Stack* can be resolved throughout the temperature range, except very close to room temperature where there is no significant difference in loss between the coated and uncoated samples. The coating loss increases quite steadily with decreasing temperature, and is dominated by the loss of its Ta_2O_5 layers as is shown above in Fig 8.1.

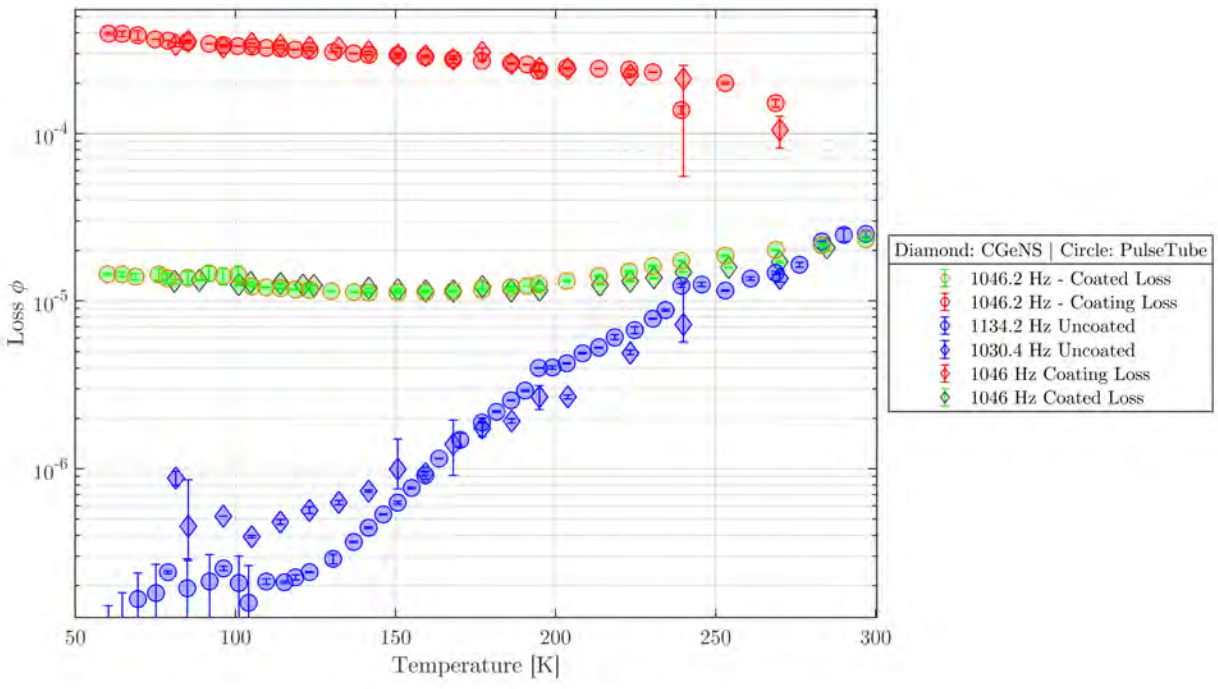

Figure 8.4: Mechanical loss of a (2,0) mode measured on an uncoated and a *Full Stack* multimaterial coated disk CryoGeNS system compared with an identical (2,0) mode on a second uncoated disk measured using the Pulse Tube cooled cryostat.

In the CryoGeNS system the lowest mechanical losses are recorded at T≈106 K (more details in Chapter 7). The uncoated disk measured in the Pulse Tube system reduced to a minimum at 121 K, coincident with a reduction in the thermoelastic loss of silicon expected at this temperature. At temperatures >160 K the uncoated mechanical loses measured on both systems increases, but at a rate below the predicted levels of thermoelastic loss.

Since this is observed in both systems, it seems unlikely that inaccuracies in temperature measurement are responsible for this. In the previous section, it was also discussed that variations in literature values of material properties for silicon could contribute to large variations in $\phi_{\text{thermoelastic}}$ at lower temperatures <123 K, but do not significantly change predictions. This effect will be discussed in more detail in section 8.3.

8.1.0.1 Losses of Coated Substrates

Figures 8.5 and 8.6 show a comparison of the loss of the coated disk measured in each system. Between 80 K and 200 K the first three mode frequencies measured in both systems show less than 8% variation in ϕ_{coated} between both systems. Above 200 K, the two systems begin to deviate, but the same trend in loss is evident. Given that the levels of loss measured on coated samples in the two systems show considerable agreement, this adds more confidence to the measurements taken in each system. The temperature stability in each system differs, as can be seen by viewing the variation in loss between each temperature step. Losses measured on the Pulse Tube cryostat shows a smooth transition in loss between each temperature step, with relatively small variations in loss at each step.

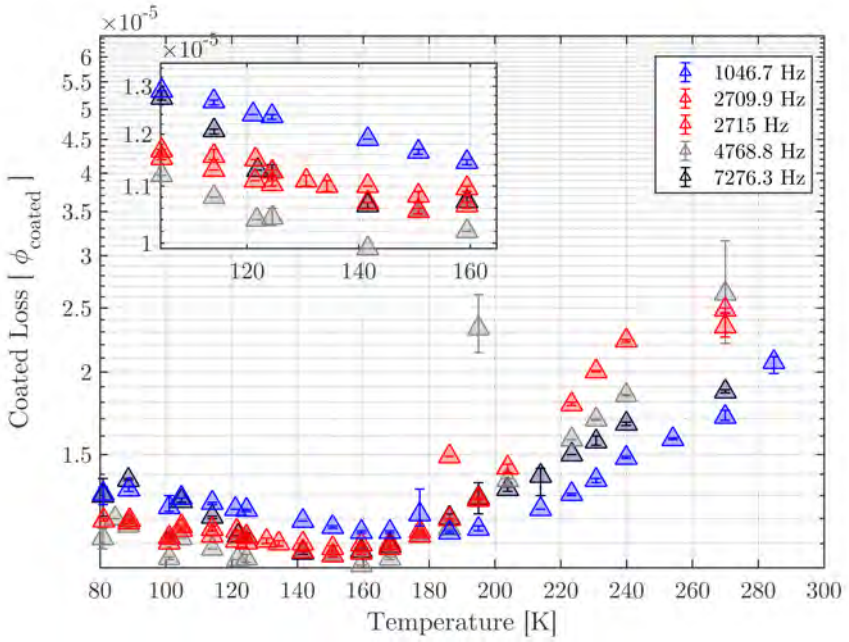

Figure 8.5: Measured losses for the *Full Stack* HR coated sample measured on the CryoGeNS system.

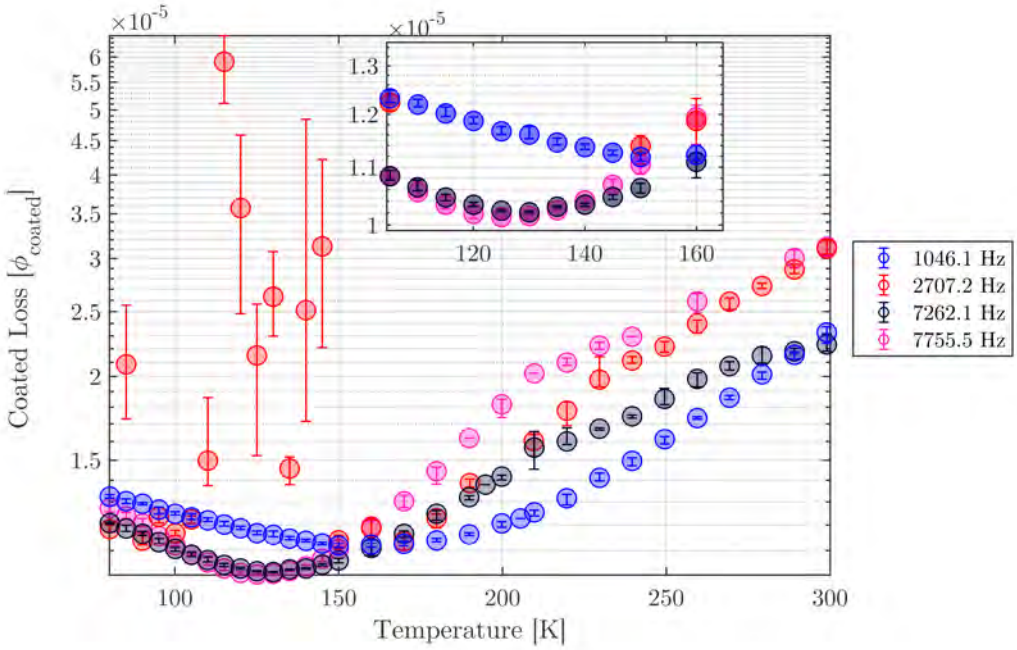

Figure 8.6: Measured losses for the *Full Stack* HR coated sample measured in the Pulse Tube cooled cryostat.

8.1.0.2 Coating Loss

Using the uncoated losses from both the CryoGeNS and Pulse Tube cryostat systems, the loss of the coating layer was calculated, using strain energy ratios ($\frac{E_{sub}}{E_{coat}}$) calculated using COMSOL, and are detailed in Tab. 8.1. The ratios were assumed to be independent of temperature. The *Full Stack* coating is analysed as a composite single layer coating with the average properties of each material (see Equations 6.1 and 6.2). Results of these calculations for measurements on the CryoGeNS system are summarised below in Fig 8.7.

f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$
1046.6	27.66	7302	34.45	17481	31.31	25344	34.41
1290.2	40.27	10256	33.87	17649	34.2	25354	34.41
2709	33.07	10269	33.98	17661	34.55	25470	31.44
2715	33.07	11179	31.61	18404	34.79	25514	31.53
4438	32.74	11200	31.61	20200	31.63	27115	35.36
4439	32.74	12802	33.62	20211	31.63	27122	35.42
4752	33.11	12806	33.62	22099	34.87	29970	33.1
4769	33.92	13713	34.22	22116	34.87	-	-
7269	33.71	15466	31.34	23521	32.06	-	-
7276	33.71	15516	32.19	23584	32.06	-	-

Table 8.1: Modelled Frequencies and calculated ratio of elastic energy strain densities for a $\varnothing = 2$" cSi disk coated with the *Full Stack*, calculated using COMSOL.

The average coating loss calculated from the measurements in the CryoGeNS system decreases steadily with increasing temperature between 95 K and 200 K. Above 200 K the coating loss shows significantly more spread and no clear trend is apparent, probably due to thermoelastic loss of the substrate starting to dominate, making it difficult to disentangle the loss of the coating. This effect is also apparent for some modes between 140 K and 180 K, leading to a larger spread in the calculated coating loss.

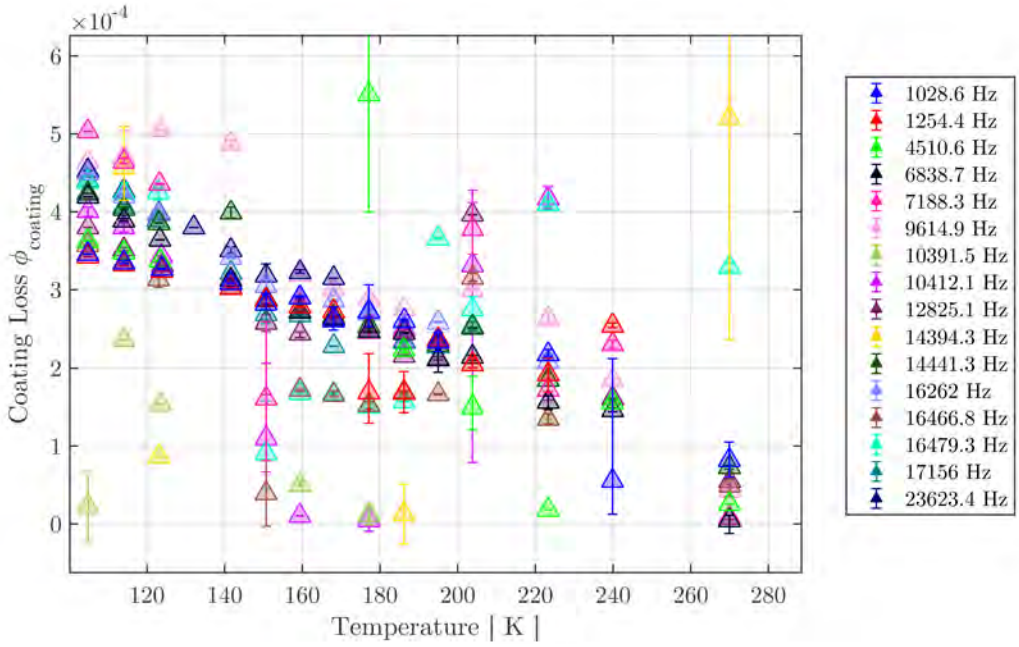

Figure 8.7: Comparison of coating loss values of the *Full Stack* measured at different mode frequencies as a function of temperature using the CryoGeNS system.

The coating loss calculated from measurements in the Pulse Tube system follows a very similar trend to CryoGeNS measurements below $T = 150\,K$ (see Fig 8.8). At 80 K, the CryoGeNS system yields an average coating loss of $(3.3 \pm 0.092) \times 10^{-4}$, showing excellent agreement with the Pulse Tube measurements of the same sample of $(3.5 \pm 0.016) \times 10^{-4}$at 1 kHz. It is interesting to note that the apparent dip in losses measured on the CryoGeNS system which occurs at $\approx 160\,K$ (see Fig 8.5) has little effect on the calculated coating loss, between $(150\,K < T < 200\,K)$, where $\phi_{coating}$ continues to decrease with increasing temperature. Modes such as those at 7.75 kHz, however, shows an increase in loss of 15% between 180 K and 200 K. The same trend is not observed in the CryoGeNS system.

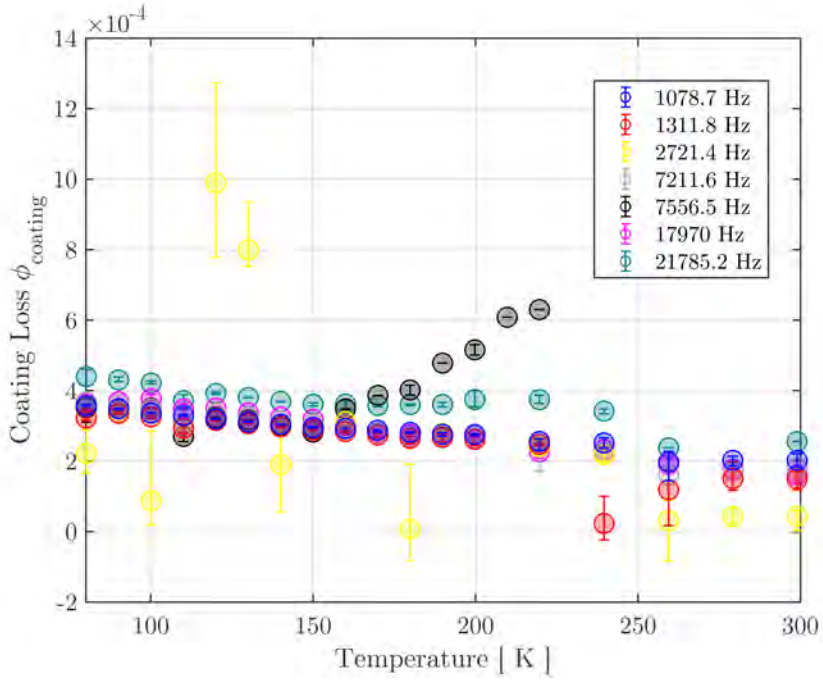

Figure 8.8: Comparison of coating loss values of the *Full Stack* measured at different mode frequencies as a function of temperature using the Pulse Tube cooled cryostat.

Above 200 K in both measurements, similar divergence between coating losses measured for different modes is observed.This is likely due to reduced sensitivity to coating loss due to the higher substrate loss in this range. For both cryostats the coating loss above 200 K, from some modes shows values close to or below zero - this could be related to stress-induced changes to the substrate thermoelastic loss and distribution of $\frac{E_{sub}}{E_{coat}}$, discussed further in section 8.3.0.2.

In Fig 8.9 measurements of *Full Stack* coating loss using a cantilever substrate are compared at similar frequencies.

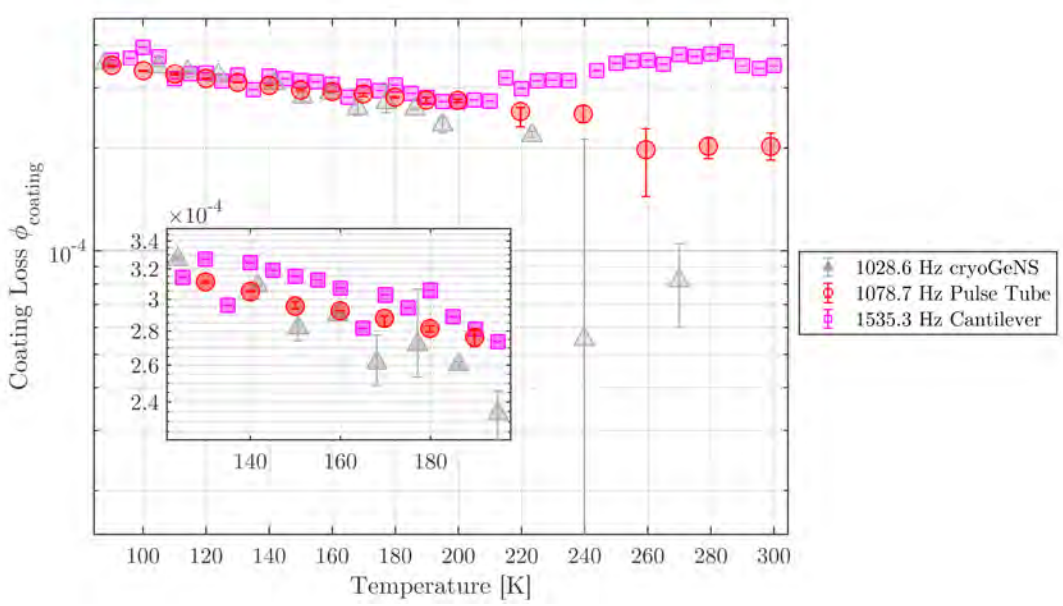

Figure 8.9: Comparison of *Full Stack* coating loss measurements on cSi cantilever and $\varnothing$=2" × ∼360 µm cSi disk.

Between 85 K and 200 K, the coating loss decreases steadily with increasing temperature. All three measurement methods show excellent agreement - although it is interesting to note to note that the Pulse Tube data is significantly less noisy than the other measurements. Above 200 K, the coating loss measured on the disks continues to decrease, while the cantilever coating loss begins to increase with some indication of a small peak or plateau close to room temperature. This could indicate that the clamping pressure applied to the cantilever, which varies with temperature, decreased above 200 K therefore increasing the measured loss.

If the afore-mentioned stress-related changes to substrate thermoelastic loss occur, these may be have a larger effect on the thinner, more deformed cantilever, possibly explaining the different coating loss results above 200 K (see section 8.3.0.2).

8.2 Measurements of RLVIP Ta_2O_5-SiO_2

The mechanical loss of the *Upper Stack*, five bi-layers of RLVIP Ta_2O_5 and SiO_2 measured on two different substrates, a $\varnothing$=2" × ∼360 µm cSi disk and a $t \approx 68$ µm thick cSi cantilever was measured. Each substrate was measured as a function of

temperature, with the cSi disk measured using the Pulse Tube cooled cryostat (see Fig 8.2) and the cantilever measured using the cantilever cryostat apparatus [93].

From the Pulse Tube cooled cryostat two First Family modes (2,0) and (2,0) at 1102.6 Hz and 1340.5 Hz were measured between 70 K and 300 K and the coating loss calculated using the same measurements of uncoated disks as before (see section 8.1.0.2). Due to issues with system stability, modes of higher frequencies could not be measured. Fig 8.10 shows the measured losses and the calculated ϕ_{coating}. The ratio of elastic strain energy $\frac{E_{\text{sub}}}{E_{\text{coat}}}$ between the coating and substrate for resonant modes at 1102.6Hz and 1340.6 Hz are 51.78 and 75.6 respectively.

The loss of the sample coated with the *Upper Stack* shows a steady linear increase decreasing temperature, as would be expected from the increase in the loss of its constituent tantala and silica layers below room temperature (see Fig 8.1). The maximum loss was measured as $(1.7 \pm 0.013) \times 10^{-5}$ at T ≈ 70 K, decreasing to $(2.5 \pm 0.11) \times 10^{-5}$ at 290 K. The coating loss of these two modes was identical below 200 K increasing from $(4.1 \pm 0.012) \times 10^{-4}$ at 200 K to $(5.4 \pm 0.45) \times 10^{-4}$ at 70 K . Some differences above this temperature are likely due to the high loss of the substrate in this regime, which reduces the sensitivity/resolution with which coating loss can be determined as is shown in Fig 8.12.

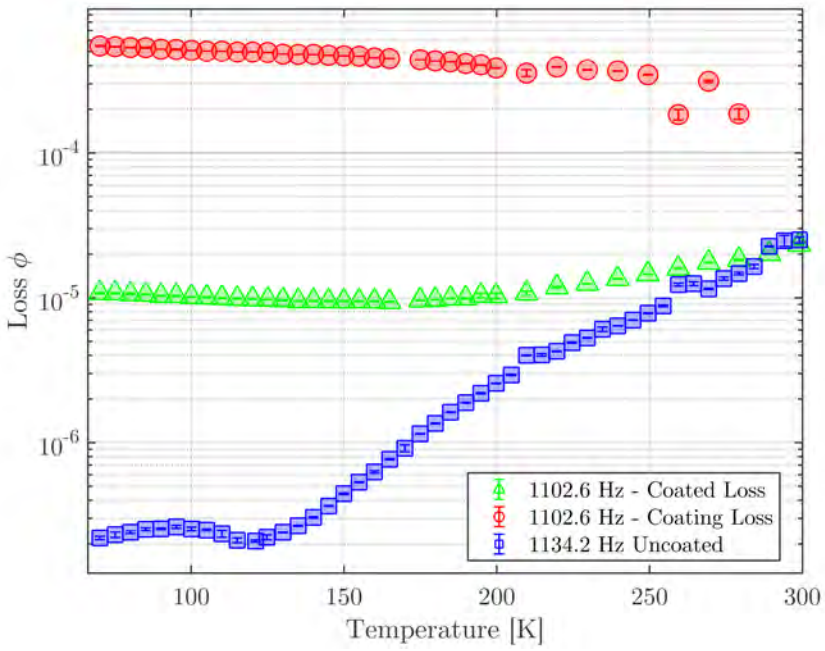

Figure 8.10: Mechanical loss of a (2,0) mode measured on an uncoated and a *Upper Stack* multimaterial coated disk using the Pulse Tube cooled cryostat.

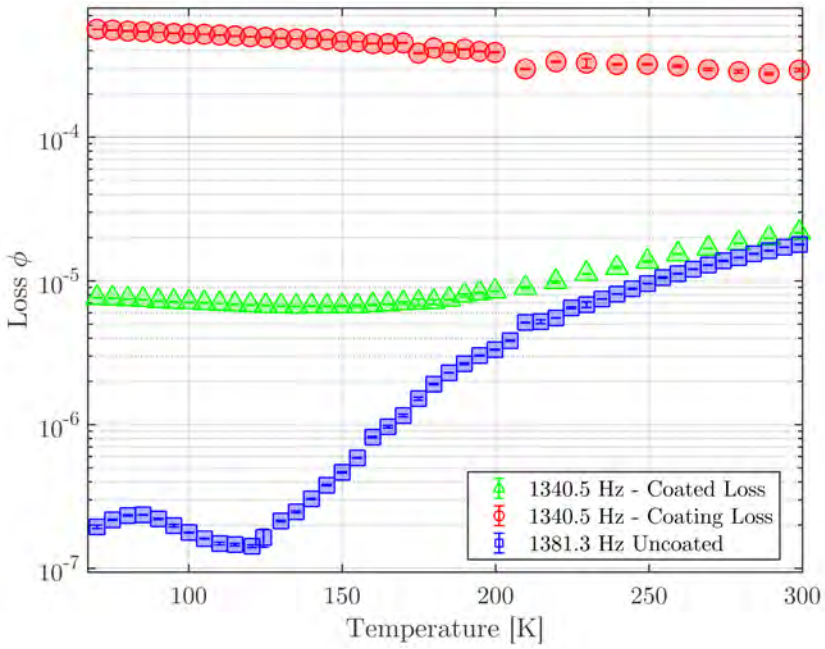

Figure 8.11: Mechanical loss of the second (2,0) mode measured on an uncoated and a *Upper Stack* multimaterial coated disk using the Pulse Tube cooled cryostat.

Using measurements of the *Upper Stack* coated cantilever substrates a comparison between the coating loss produced from modes of the disk substrate can be produced. Figure 8.12 compares the coating loss values shown in Figures 8.10 and 8.11 with those produced from a coated cantilever substrate.

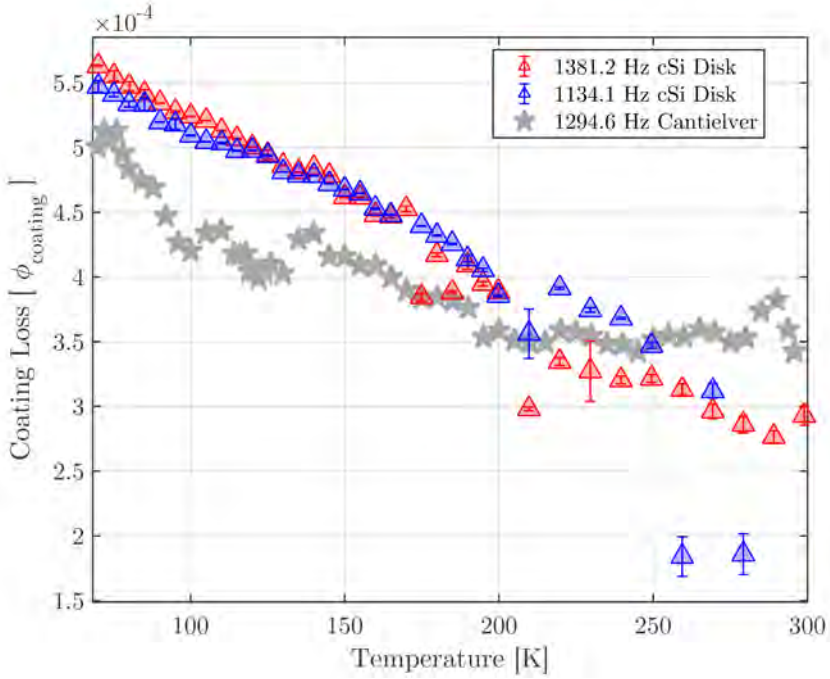

Figure 8.12: Coating Losses of the *Upper Stack* material measured on cSi disk using the Pulse Tube cryostat and cantilever substrates at similar frequencies.

Below 100 K, the coating loss measured from the disk substrate was 5% - 10% higher than that measured using the cantilever. The difference in loss between substrates diverges as temperature of the sample approaches 123 K. Above 123 K, ϕ_{coating} continues to decrease on both substrates where the trends from the cantilever and 1.1 kHz disk mode intersect at 210 K where the cantilever coating loss plateaus and remains nominally constant at 3.5×10^{-4} between 210 K and 300 K. The disk coating loss continues to decrease through this temperature range showing a minimum loss of $(2.7\pm0.05)\times10^{-4}$ at 290 K.

8.3 Measurements of RLVIP aSi-SiO_2

Measurements of a nominally identical cSi disk, coated with the *Lower Stack* coating material due to experimental restrictions occurring in March 2020 was only

measured using the CryoGeNS system. The measured losses for all modes which were successfully measured (and showed no signs of MMI) are summarised by mode family in Figs 8.13, 8.14 and 8.15.

First Family modes measured on the coated sample between 80 K and 130 K decreases as the temperature of the sample increases, with a reduction of ∼48% measured on average. Above 140 K the measured loss of each mode increases as the thermoelastic loss of the substrate and aSi layers increases. It is interesting to note that at temperatures above 180 K the measured mechanical loss of the coated sample is measured to be below the predicted thermoelastic loss of the substrate. This difference suggests that there either is an inconsistency in the control data used, the temperature of the sample or that the coating has altered the thermoelastic loss of the substrate. For the Second and Third Family, the mechanical losses are more consistent in magnitude between different measured frequencies.

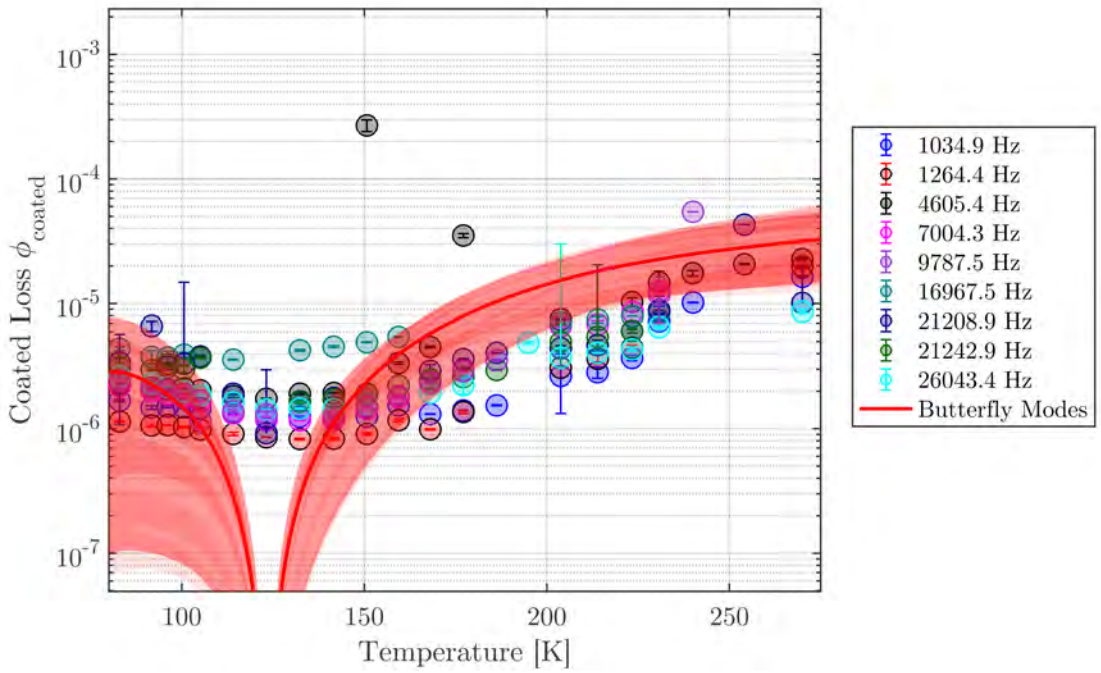

Figure 8.13: Comparison of losses from First Family modes measured on a sample, coated with the *Lower Stack* material. as a function of temperature using the CryoGeNS system. Values are compared with the predicted uncoated substrate thermoelastic loss for each mode family.

At temperatures above 180 K 45% of frequencies were measured below the predicted thermoelastic limits, with this trend continuing up to 240 K. Temperatures where the predicted substrate thermoelastic is reduced ($110 < T < 130$) losses from the coated sample associated with Second and Third Family modes reaches a minimum of 1.0×10^{-6} 15% higher than the lowest loss measured for First family modes. The differences in loss measured between different mode families could suggest that an additional surface loss term is present, which does not affect First and Second/Third mode families in the same manner. In section 7.4 (see Fig 7.16) inhomogeneities in the samples surface thickness were measured on uncoated ⌀ =2" cSi disks nominally identical to the one used in this study which could justify the differences in measured losses.

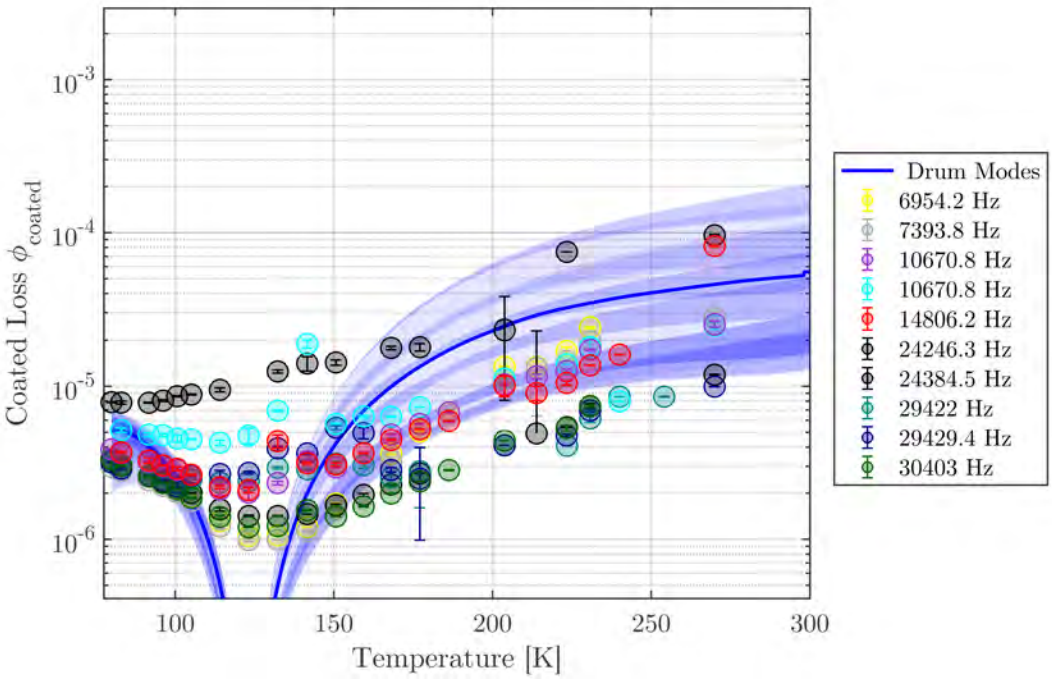

Figure 8.14: Comparison of losses from Second Family modes measured on a sample, coated with the *Lower Stack* material. as a function of temperature using the CryoGeNS system. Values are compared the predicted uncoated substrate thermoelastic for each mode family.

As the inhomogeneities on the surface of the sample show a radial thickness gradient with the thickest portion at the centre of the sample, modes shapes which interact more with this region should therefore have different surface loss contribution. The mechanical losses from each family, between 80 K and 180 K are shown in Fig 8.16. At the thermoelastic minimum, the loss of each mode reduces to

a different magnitude, the lowest losses found on the 1.03 kHz first family mode. Comparing this with the measured losses of the 16.7 kHz - Third, and 7.39 kHz - Second Family modes show a correlation between mechanical loss, frequency and radial motion of each mode shape.

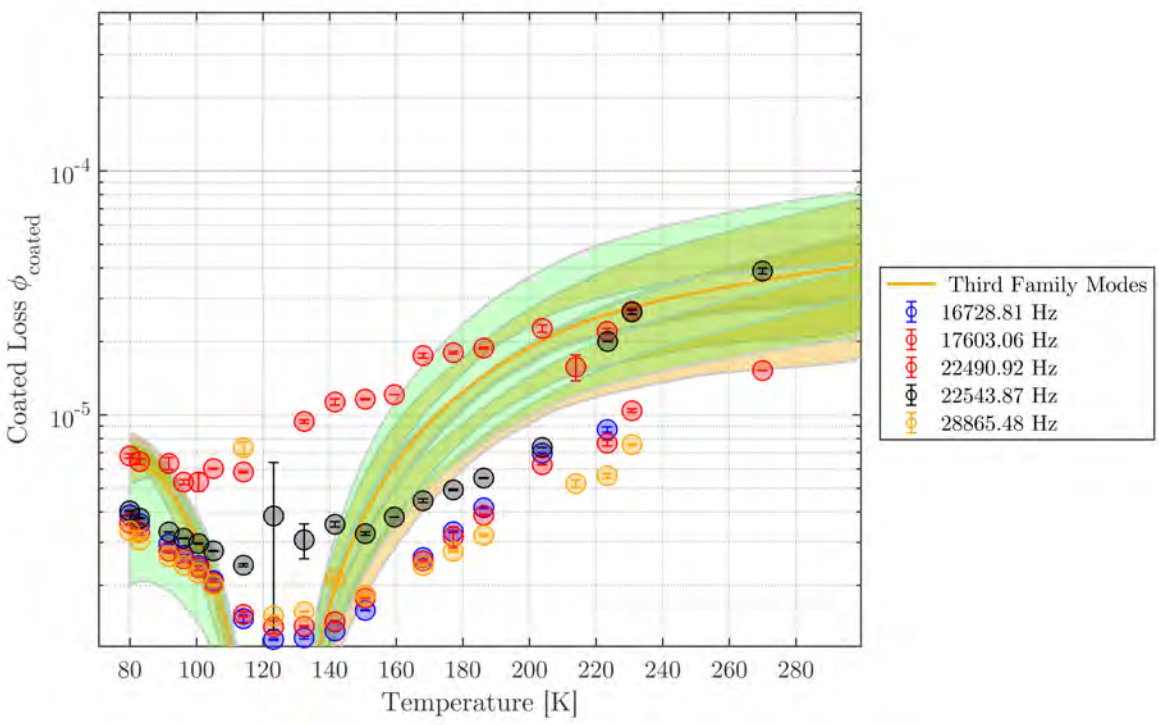

Figure 8.15: Comparison of losses from Third Family modes measured on a sample, coated with the *Lower Stack* material. as a function of temperature using the CryoGeNS system. Values are compared with the predicted uncoated substrate thermoelastic for each mode family.

Implications of Surface Losses

Work has been published which discusses the correlation between surface losses in cSi cantilever substrates [252]. This showed that the level of surface loss in silicon varies with the surface roughness, with lower frequency modes (with larger surface deformation) being more susceptible. The effect of the surface loss on the total loss of the sample depends on the surface-to-volume ratio. The results discussed in [252] could corroborate the difference in losses measured between the 7.39 kHz (Second Family) and 16.7 kHz (Third Family) modes at 123 K, due to the larger deformations produced by the 16.7 kHz mode overlapping with the thickest region of the uncoated sample. By comparing the surface deformation of this mode to the

16.7 kHz (Second Family) and 1.03 kHz (first family) modes a decreasing concentration of motion close to the centre of the sample.

It could therefore be speculated that these regions are contributing a higher component to the surface loss if the variations in thickness also correspond to an increase in roughness.

8.3.0.1 Coating Loss

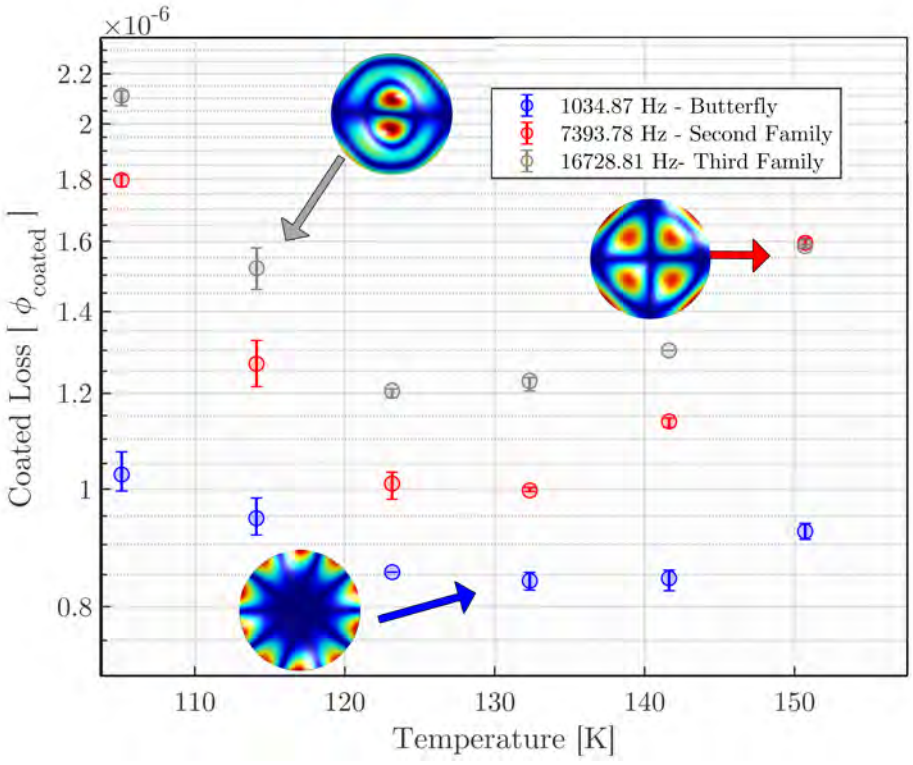

Figure 8.16: Comparison of measured losses from modes from three different mode families. First Family : 1.03 kHz, Second Family : 7.39 kHz, Third Family : 16.7 kHz.

After examining the possible correlations between surface loss and mode shape the loss of the coating layer was calculated using energy ratios produced from COMSOL. Due to the low mechanical loss of aSi and SiO_2 at cryogenic temperatures (T>100 K), these values are comparable to those of an uncoated substrate, making the calculated coating loss susceptible to small fluctuations in the substrate loss. In order to best take the potential variability in substrate losses into account, two different uncoated cSi disks were measured, one in the CryoGeNS system and the other using the Pulse

Tube cooled cryostat. The ratio of elastic strain energy density stored in the coating and substrate of a for a $\varnothing = 2$" $t = 341\mu$m cSi disk coated with the *Lower Stack* are detailed in Tab. 8.2.

f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$	f (Hz)	$\frac{E_{sub}}{E_{coat}}$
1031.8	63.14	9868.8	77.52	17518	77.87	28346	75.62
1249	92.36	9873.6	77.79	19138	72.3	29458	71.66
1722.9	73.23	10480	72.41	19138	72.3	29458	71.66
2615.6	75.67	10480	72.41	21282	79.8	29733	72.67
2615.6	75.67	11923	77.04	21282	79.8	31245	80.09
4037.4	75.17	11923	77.04	22063	73.39	31438	78.68
4037.4	75.17	13206	78.29	22063	73.39	31490	81.68
4562.8	75.79	13206	78.29	23860	78.63	31490	81.68
4589.5	77.66	14523	71.75	23860	78.63	34803	74.14
6774.8	69.79	14628	73.64	24116	71.74	34803	74.14
6995.4	77.18	16319	71.7	24126	71.92	35114	72
6995.4	77.18	16982	78.24	26095	80.91	35114	72.01
7215.9	79.01	17004	79.05	26099	81.04	-	-
7606.9	75.89	17203	79.65	28017	72.56	-	-

Table 8.2: Modelled Frequencies and calculated ratio of elastic energy strain densities for a $\varnothing = 2$" $t = 341\mu$m cSi disk coated with the *Lower Stack*, calculated using COMSOL.

Comparing these measurements requires a caveat that the *Lower Stack* losses have been shifted in temperature following the method described in section 7.3 allowing the relative temperatures of both control samples to be compared. As this correction does not apply to measurements taken on the Pulse Tube cryostat, an error on the estimated sample temperature and the resultant values of ϕ_{coating} from CryoGeNS measurements is expected. However as the resultant change in the calculated sample temperature does not change for different modes the observed effect of negative losses shown in Fig 8.17 remains. Below 150 K the loss of each uncoated substrate is low to resolve the contributions of the coating. At 85 K, the coating loss was calculated as $(7.1 \pm 0.3) \times 10^{-5}$ using the CryoGeNS control data and $(1.2 \pm 0.0025) \times 10^{-4}$ from

the Pulse Tube control data, the large difference in calculated coating loss between systems stemming from the level of loss for each uncoated substrate. Between 100 K and 130 K the thermoelastic loss of the substrate decreases, allowing the coating loss to be calculated with these effects reduced. In this temperature range measurements of ϕ_{coating} continues to decrease in both systems which can be attributed to a decrease in loss of the SiO_2 layers (see Fig 8.1).

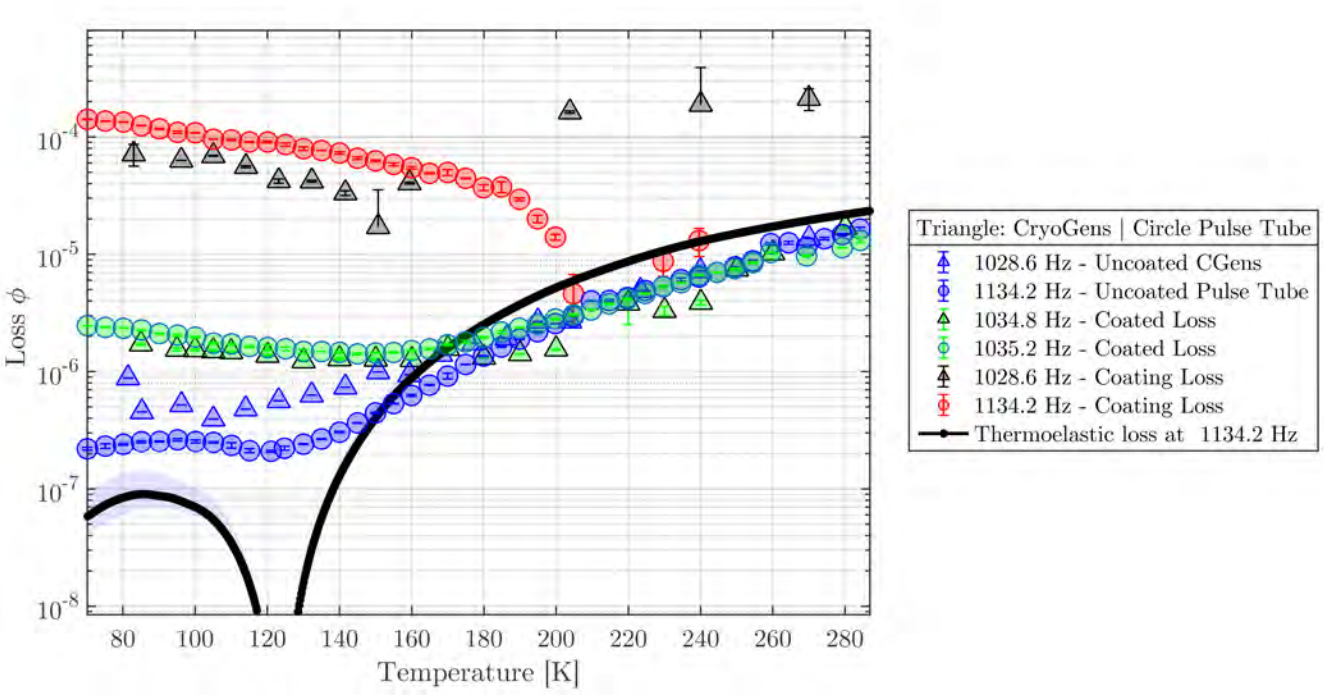

Figure 8.17: Calculated coating loss of the *Lower Stack* deposited on a $\varnothing = 2" \times 360\,\mu$m cSi disk as a function of temperature compared with uncoated and coated loss measurements and thermoelastic loss predictions.

At temperatures above 150 K the losses of the uncoated substrate measured in the CryoGeNS system become comparable to the loss of the coated sample, producing a nonphysical, negative, value of ϕ_{coating}. The uncoated disk measured in the Pulse Tube cryostat shows 51% lower levels of loss a 121 K than the disk measured using the CryoGeNS system. Using these values to predict ϕ_{coating} can resolve the loss of the coating layer up to a sample temperature 170 K before the coated loss again decreases below the loss of the uncoated substrate and the predicted thermoelastic limit. These results - that above 180 K the loss of the coated disk is lower than the loss of the uncoated substrate - is a strong indication that the coating stress may be altering the loss of the substrate and or coating layers [203].

Fig 8.18 shows the loss of a cSi cantilever coated with the *Lower Stack* compared to the coating loss measured using a disk substrate. Above 200 K the measured loss of the coating layers is influenced by the thermoelastic loss of the cSi substrate, causing the loss to increase. Due to the layers of aSi inside the coating, these layers will also contribute to the magnitude of $\phi(\mathrm{T})_{\mathrm{therm}}$.

Between 150 K and 200 K the thermoelastic loss of the substrate falls below the level of ϕ_{coating}, producing the lowest measurements of $\phi(\mathrm{T})_{\mathrm{coating}}$ at this frequency. In the 8 K - 150 K range, the loss of the cantilever increases, to a peak value at 21 K. Such an increase in loss over these temperatures is consistent with the increase in mechanical loss for SiO_2 [93, 105, 150, 228].

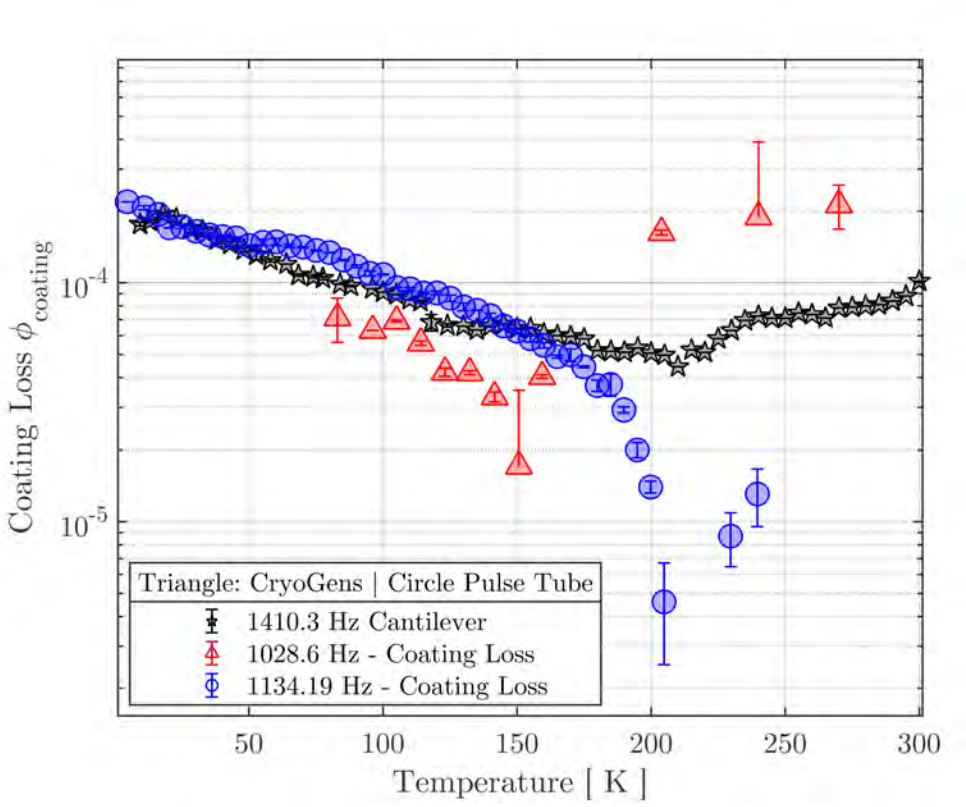

Figure 8.18: Cryogenic coating loss of the *Lower Stack* coatings on cSi cantilever compared with coating losses measured on cSi disks from the CryoGeNS and Pulse Tube cryo systems as a function of sample temperature.

Below 150 K, coating losses from cSi disk and cSi cantilever show good agreement. At higher temperatures the differences in coating loss between cantilever and disk substrates increases, with the cantilever coating loss being 64% higher at

150 K. The cantilever substrate again continues to show relatively little change in the measured decreasing trend of coating loss between 150 K and 210 K, above this range the measured loss is influenced by $\phi_{\text{thermoelastic}}$.

8.3.0.2 Discussion of Results

The discrepancy between the levels of coating loss for different substrates suggests that the level of loss of each disk substrate is also different from the other. As the comparison between different disk substrates produces broadly similar values of coating loss below 150 K with the same trend as a function of temperature, this difference in loss has little effect on the calculated values of ϕ_{coating} agreeing with the loss produced by the coated cantilever. Above 150 K the difference in substrate loss becomes of higher consequence to the calculation of ϕ_{coating}.
Assuming that above 200 K the measured loss of both cSi cantilever and disk substrates are dominated by $\phi_{\text{thermoelastic}}$, when using Eq 7.6 there is a large discrepancy between predicted $\phi_{\text{thermoelastic}}$ and the level of experimentally measured loss on disk substrates (see Fig 8.17). Comparing the methods used to calculate the thermoelastic loss for these geometries [241] and [135] the key difference is the use of volumetric thermal expansion for disk resonators, which considers the expansion of the disk through its width and radius, as well as its thickness, raising $\phi_{\text{thermoelastic}}$ compared to $\phi_{\text{thermoelastic}}$ for a cantilever substrate.

Given that there is a discrepancy between uncoated measurements of ϕ for disk substrates, the disagreement between theory and measurement could be explained under this premise. However, the fact remains that even by lowering the level of $\phi_{\text{thermoelastic}}$, this does not affect the measured difference in loss between an uncoated and coated substrate.

This result can be explained by either one of three possibilities :-

- The loss of the *Lower Stack* coating measured in both cryostats was not optimal.

  In the case where measurements of a sample's uncoated mechanical loss has been affected by external damping, the ringdown duration is reduced increasing ϕ. If measurements of the *Lower Stack* coating have been subject to

such effects, this would suggest that the loss of the coated substrate could be lower, exacerbating the difference between coated and uncoated substrate losses. While this argument also applies to measurements of the uncoated substrate, the same levels of loss from different disk substrates makes this an unlikely result.

- Estimation of sample temperature in all experiments is inaccurate and disk temperature is over-estimated at $\mathrm{T} > 150\,\mathrm{K}$.
- The loss of the *Lower Stack* coating is accurate and the stress imparted by the coating layers is effecting the $\phi_{\mathrm{thermoelastic}}$ loss of the now coated substrate.

The possibility of using the change in the disk's mode frequencies as a cross-check of the disk temperature was investigated. The mode frequencies change with temperature due to the temperature dependence of the Young's modulus, density and dimensions of the disk. This relationship has been previously studied by S.Reid [253] for cSi cantilevers, finding a quadratic relationship between $\phi(f, \mathrm{T})$ and frequency. The measured frequency of each sample should then follow this relationship as a function of temperature, where any deviations may suggest that external factors have influenced the sample.

The measured change in frequency for the (2,0) First Family mode, and a cantilever mode of similar frequency, are compared in Fig 8.19.

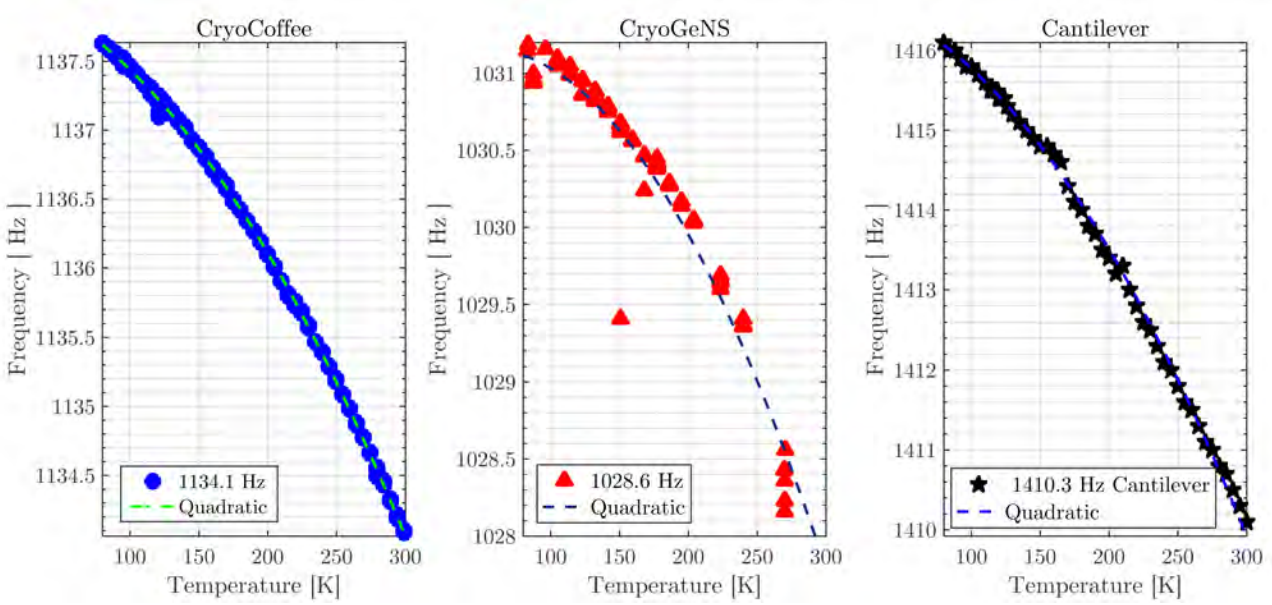

Figure 8.19: Comparison of the measured frequency shifts for two uncoated disks measured using (a) Pulse Tube cooled cryostat (b) CryoGeNS and a cantilever coated with the *Lower Stack* HR coating (c).

This comparison between three different cryogenic systems shows a clear correlation between resonant frequency and measured sample temperature. As the cantilever measurement system involves direct thermal contact to the upper and lower surface of the sample, its temperature can be most accurately gauged from its' supporting structure. Comparing the change in frequencies between the cantilever (c) and disk, the Pulse Tube system (a) shows better agreement in measured trend. The change in frequency of each mechanical mode shows a clear quadratic trend as their temperature is increased from 80 K to 300 K :-

$$\text{PulseTube} : -3.6 \times 10^{-5}\text{T}^2 - 0.0026\text{T} + 1.14 \times 10^3, \tag{8.1}$$

$$\text{CryoGeNS} : -5.4 \times 10^{-5}\text{T}^2 + 0.0055\text{T} + 1.03 \times 10^3, \text{and} \tag{8.2}$$

$$\text{Cantilever} : -6.6 \times 10^{-5}\text{T}^2 - 0.0029\text{T} + 1.42 \times 10^3 \tag{8.3}$$

Comparing $f(\text{T})$ from the cantilever sample to the two cSi disks shows a much larger change in frequency over the plotted temperature range, which is assumed to be an effect defined by the sample geometry.

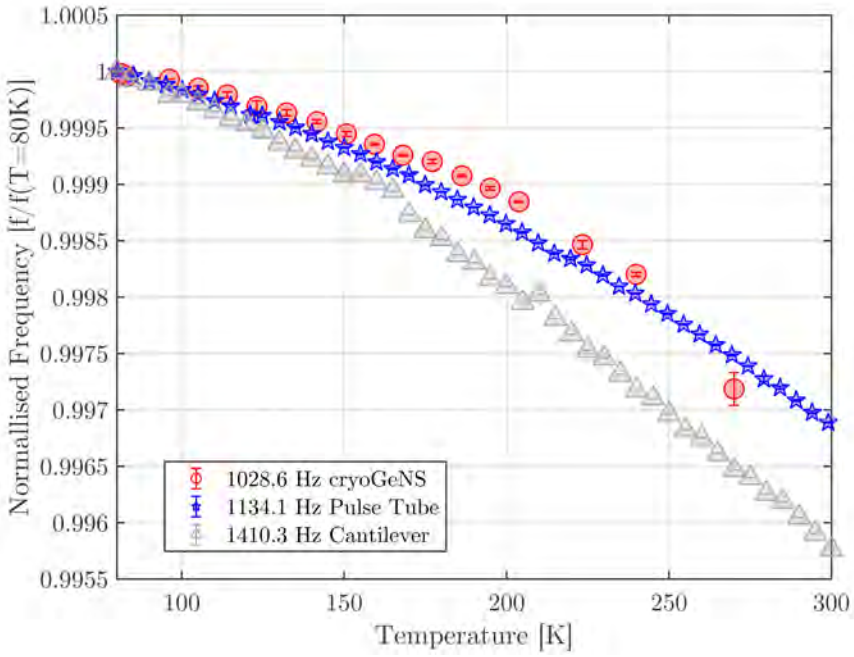

Figure 8.20: Comparison of measured frequency shifts shown in Fig 8.19 where the frequencies of each data-set have been normalised to 1 at 80 K.

As the temperature of both disk samples is increased, the measured frequencies from each substrate follow different trends. Above 250 K the functional form of $f(\mathrm{T})$ exhibits a change which is not observed on the other substrates, this possibly suggests that the temperature of these points in the CryoGeNS system has been adversely affected in some way. Between 160 K and 200 K a comparable $f(\mathrm{T})$ trend is observed between the two disk samples suggesting that these measurements were taken at comparable temperature stability. However, below 150 K the loss measured in the CryoGeNS system begins to decrease with respect to the Pulse Tube cryostat, indicating another difference between in the behaviour of the samples measured in each cryostat. The clear difference in frequency shifts measured on similar geometries above 160 K potentially places a degree of uncertainty on any measurements taken in the CryoGeNS system above this range. By comparing the measured frequency shift of the cantilever to those produced by the Pulse Tube cryostat the similarities in the measured change in frequency also validates these measurements.

However, as was shown in Fig 8.18, the coating loss measured using the Pulse Tube

system does not produce physical values between 150 K and 293 K. Unless the cSi disk that was coated in the *Lower Stack* material has different thermoelastic behaviour to samples from the same manufacturing run, then this could then suggest that due to the differing thermal expansion properties of the coating and substrate that a curvature is induced in the sample. It is possible that this changes the thermoelastic properties of the substrate, and distribution of strain energies in the sample [203]. The effect of coating thermoelastic loss [178] could also affect these predictions, but are not considered in this chapter.

8.4 Comparison of Coating Loss from Different Substrates

In Chapter 6, the mechanical loss of the *Full Stack* and its component stacks were measured at room temperature (≈ 293 K) using a much thicker substrate, t$\approx$2.6 mm, $\varnothing$ $\approx$3" (76.2 mm) Corning 7980-0A grade SiO_2 disks. In a previous study by M. Fletcher, the thickness of a substrate compared to an applied coating layer can determine if ϕ_{coating} will be influenced by the increased stress/curvature due to the coating layer [203]. They surmised that the change in curvature of a sample was the dominant source of the change in ϕ_{coating}, and this would therefore have a larger effect on thinner substrates. The room temperature coating losses measured on these disks were therefore compared to cSi disk and cSi cantilever measurements.

For reference, the radius of curvature of cSi cantilevers coated in the *Full Stack* and its two component coatings was measured by Maya M. Kinley-Hanlon in the as-deposited state to ascertain the total change in curvature with respect to an uncoated sample (using the apparatus described in [93]) allowing the total stress imparted by the coating to be calculated using Stoney's equation [254]. The average stress of the *Full Stack* was calculated as $(4.4 \pm 1.1) \times 10^8$ Pa compared to the $(2.9 \pm 0.48) \times 10^8$ Pa from the *Upper Stack* and $(7.8 \pm 2.8) \times 10^8$ Pa for the *Lower Stack*. The measured difference in stress measured on cantilever substrates then suggests that a similar change in stress could be assumed for each cSi disk substrate. It is thought that the higher levels of stress produced from the *Lower Stack* coating would then produce a larger change in the thermoelastic loss of the

uncoated disk substrate and in the ratio of elastic strain energy stored in the coating and substrate. Changes in both of these components therefore have a larger effect on samples where t_{coat} is comparable to t_{sub}.

Upper Stack

The loss of the *Upper Stack* coating, is dominated by the loss of its Ta_2O_5 layers, which increases with decreasing temperature and shows a clear peak at $\approx 50\,\mathrm{K}$ [93,204] when measured using cSi cantilevers. However, this is not observed in measurements from the coated cSi disk. At $50\,\mathrm{K}$ the magnitude of ϕ_{coating} from the cSi cantilever and disk samples are comparable, but outwith this temperature, ϕ_{coating} from the disk decreases linearly with decreasing temperature, even when the cantilever coating loss has started to fall on the other side of the loss peak. It can be speculated from this comparison that the mechanisms responsible for the low-temperature loss peak of Ta_2O_5 observed in the cSi cantilever data is somehow suppressed. However, the reverse can also be argued, stating that as the cSi disk is approximately four times thicker than the cSi cantilever, the change in loss is somehow masked by the substrate.

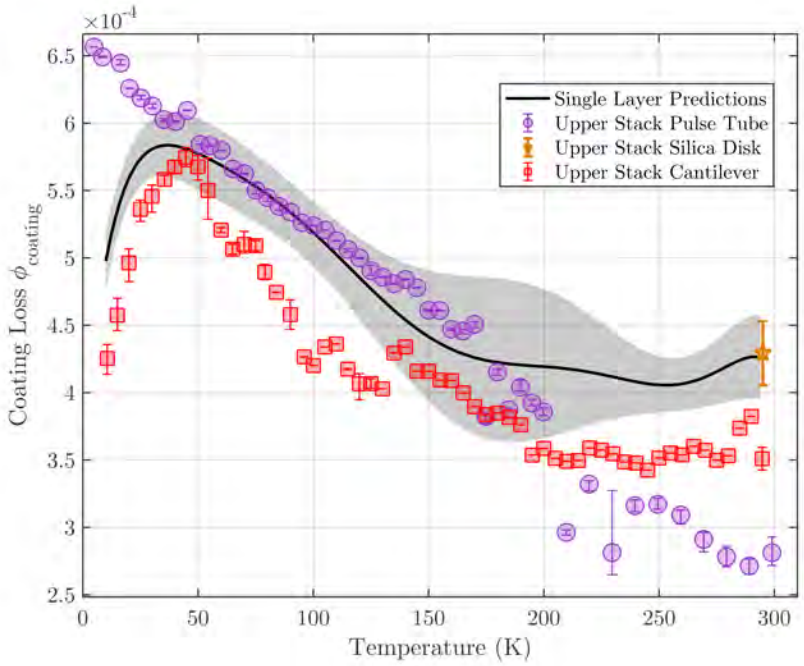

Figure 8.21: Comparison of ϕ_{coating} of the *Upper Stack* on cSi cantilever (t≈68 µm), cSi disk (t≈360 µm), and SiO_2 disk (t≈2.6 cm) as a function of temperature. Predictions of the coatings loss from cSi cantilevers coated in single layers of SiO_2 and Ta_2O_5 are also shown.

Using measurements of single layer SiO_2 and Ta_2O_5 coatings (see Fig 8.1) the expected loss of the *Upper Stack* was calculated using Eq 6.5 over the temperature range. These predictions (shown in Fig 8.21) align much closer with measurements made using a cSi disk substrate, showing excellent agreement between 100 K and 150 K. Outwith this range, the single-layer predictions are much higher than ϕ_{coating} measured on either cSi substrate. At room temperature, the predictions of the *Upper Stack* is more than 20% higher than ϕ_{coating} from the cSi cantilever and more than 35% than from the cSi disk. However, these predictions encapsulate the average loss of the *Upper Stack* from a SiO_2 disk (see Chapter 6) at room temperature.

It was discussed in Chapter 6 that during deposition of the multimaterial stacks studied in this chapter, the deposition chamber reached temperatures above 200 °C. While the same coating technique was used to deposit every layer, the time taken to deposit a single layer coating will be radically different from an HR stack's deposition. Thus it can be speculated that the temperature that each coated substrate reached is proportional to the deposition duration with coating materials at the bottom of

each stack receiving the longest exposure, effectively heat treating these layers and reducing their mechanical loss.

Lower Stack

Fig 8.22 shows the results produced by performing the same comparison of coating losses from single layer RLVIP aSi and SiO_2 to measurements of the *Lower Stack* coating as a function of temperature. Below 80 K, the loss peak associated with layers of SiO_2 is observed at ≈ 20 K in measurements from the *Lower Stack* cSi cantilever, and in the predictions from single-layer materials.

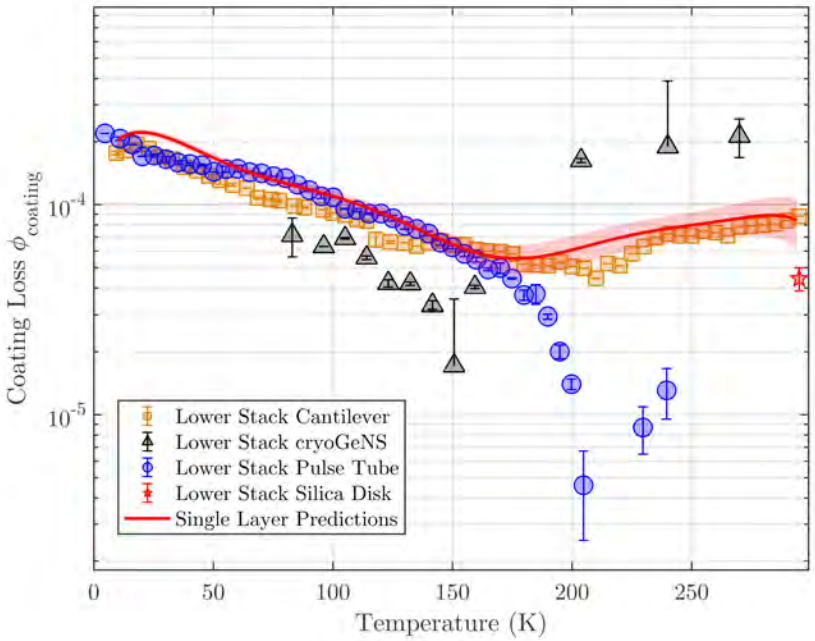

Figure 8.22: Comparison of $\phi_{coating}$ of the *Lower Stack* on cSi cantilever ($t \approx 68\,\mu m$), cSi disk ($t \approx 360\,\mu m$), and SiO_2 disk ($t \approx 2.6$ cm) as a function of temperature. Predictions of the coatings loss from cSi cantilevers coated in single layers of SiO_2 and aSi are also shown.

Measurements of the *Lower Stack* on a cSi disk show the same magnitude of loss as measured on the *Lower Stack* coated cantilever, but do not resolve the loss peak. Between 80 K and 180 K, predictions of loss from each substrate and coating show good agreement. Above 180 K, the coating losses from both *Lower Stack* coated disks rapidly decrease towards zero, where the resolution between the loss of the coated substrate and control data is reduced. This difference is thought to be due to stress-induced changes in energies stored in the coating and substrate, as well as

effects on the substrates thermoelastic loss (see section 8.3.0.2) and inadequate control data.

As the *Lower Stack* coating was found to produce a cantilever with higher stress than samples coated with the *Upper* and *Full* stacks (see section 8.3.0.2) it follows that this effect would be more pronounced for these coating layers. Therefore, it can be speculated that as the cSi cantilever is thinner and can curve in response to the coating layer, some stress is relieved in the process. The same stress applied to the thicker silica disk does not cause a noticeable change in its curvature, and its loss is, therefore, less susceptible to this effect.

Full Stack

Measurements from substrates coated with the *Full Stack* coatings show good agreement with single layer predictions between 100 K and 200 K. This same agreement level of agreement has also been observed in the *Upper* and *Lower* Stacks. Below 70 K, the loss of the *Full Stack* coated cantilever begins to increase, reaching a maximum at 32 K.

This increase in loss is also reflected in single layer predictions. Once again the loss peak is not observed in the cSi disk measurements. However, contrary to the comparisons of the *Upper* and *Lower* stacks, the magnitude of ϕ_{coating} is consistently of lower magnitude than the coating loss peak. Failure to resolve the coating loss peak using cSi disks coated with each coating stack could indicate that the disk's temperature is higher than expected between 4 K and 50 K, suggesting that the cooling of the disk is limited by an external factor, such as the heat imparted by the 633 nm laser, which central to this experiment's ringdown measurements.

Above 150 K, a divergence in ϕ_{coating} from different substrates is observed. This follows previous observations on the *Lower* and *Upper* stack coatings. At these temperatures, ϕ_{coating} from the *Full Stack* coated cantilever increases above the levels predicted by its single-layer materials. This increase could be due to an abrupt change in the sample's loss, but as an increase of such magnitude is not observed in any other measurement, this seems unlikely. It is more probable that an environmental factor, such as vacuum tank pressure decreased as the cantilever

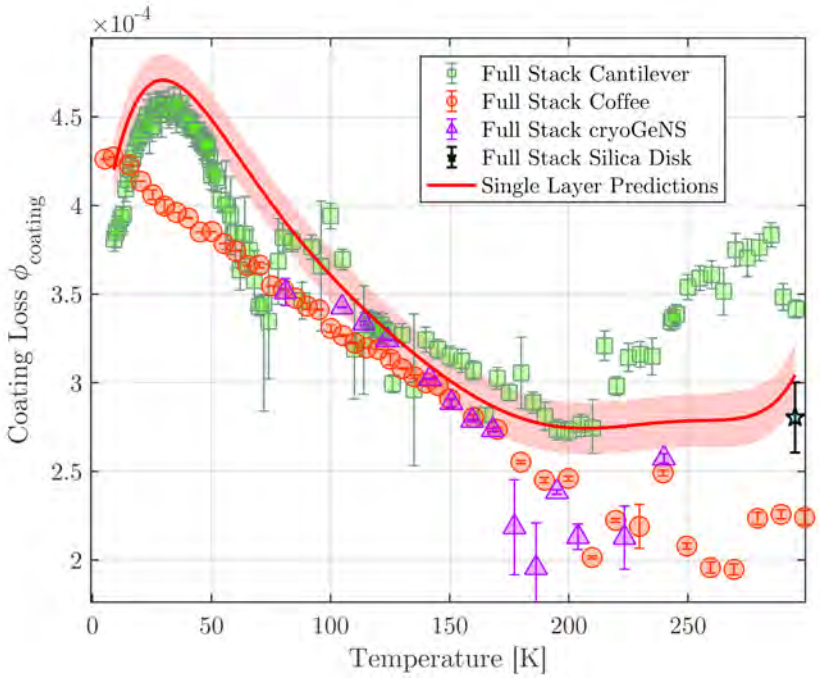

Figure 8.23: Comparison of ϕ_{coating} of the *Full Stack* on cSi cantilever (t≈68 µm), cSi disk (t≈360 µm), and SiO_2 disk (t≈2.6 cm) as a function of temperature. Predictions of the coatings loss from cSi cantilevers coated in single layers of Ta_2O_5, SiO_2 and aSi are also shown.

cryostats temperature increased, causing the abrupt loss change.

At 293 K average measurements of ϕ_{coating} from a *Full Stack* coated 3" silica disk is shown. Its' losses are comparable to the single-layer cantilever predictions but are ≈ 20% higher than values measured on cSi disks, and ≈15% lower than values from the coated cSi cantilever. The agreement between single layer predictions and silica disk measurements was also observed in *Upper Stack* coatings in Fig 8.21. As this was not observed in the analysis of the *Lower Stack* coatings, this suggests that the higher loss layers of Ta_2O_5 dominate over any other effects in this temperature range.

Discussion

By comparing the levels of coating loss of the *Upper*, *Lower* and *Full* stack HR coatings from measurements of cSi cantilever, cSi disk and silica disks between 4 K and 293 K several differences in ϕ_{coating} extracted from each substrate have been found.

Between 4 K and 50 K measurements of cSi cantilevers can resolve the mechanical loss peaks from the layers of Ta_2O_5 and SiO_2, while cSi disks coated with the same HR coating stacks cannot. The absence of a coating loss peak at these temperatures suggests one of two possibilities :-

- the loss peak measured on cSi cantilevers is an artefact of the deformation of the cantilever substrate due to the bi-metallic strip effect, or
- the disk does not reach the low temperatures required to produce a measurement of the loss peak.

As measurements of the uncoated cSi disks between 4 K and 50 K are more than an order of magnitude lower than the loss of each coated sample, the loss of each uncoated substrate would not be detrimental to measurements of coating loss for any of the materials studied in this chapter, within this temperature range.

Several works investigate the effect of stress-induced curvature in the TE loss of a silicon substrate but [203, 233, 234], none focus on these effects under cryogenic conditions. Therefore it is difficult to advocate the bi-metallic strip effect responsible for the observed loss peak in cryogenic cantilever measurements. Given recent work published by L. Kuo *et al.* [255] which shows that the mechanical loss peak in silica / titania-based nanolayer coatings can be suppressed on cSi cantilevers, suggests that the mechanical loss peaks in these materials are produced by intrinsic properties of the coating layers. It then seems unlikely that this explanation can be correct; however, further experimentation is needed. This suggests that the disk temperature in this region is the limiting factor in these measurements with data from samples coated with each stack suggesting that a coated disk's temperature does not cool below ≈50 K, where the loss peak would be observed[2]. As measurements of uncoated samples have shown that lower substrate temperatures can be achieved, further study is needed to explain this discrepancy in loss between substrates.

Between 60 K and 150 K measurements of coated cSi disk and cantilever samples show excellent agreement with single layer predictions and measurements of their

[2]This same effect is not observed in measurements of uncoated cSi disks, where temperature of the disks are shown to be lower than 50 K

respective HR coatings measured in this chapter. These results give a high degree of confidence to the levels of ϕ_{coating} measured on each substrate within this temperature range.

Above 150 K, coating losses from cSi cantilevers, cSi disks and silica disks all estimate values of ϕ_{coating} which differ between samples, and in the majority of cases, from single-layer predictions. The level of disagreement between samples seems to correlate with the mechanical loss of that sample, with coating losses extracted from silica disk measurements showing much better agreement with the single-layer predictions for HR stacks which contain layers of higher loss Ta_2O_5, than with those which do not.

This result then suggests that values of ϕ_{coating} from thin cSi substrates can be very sensitive to changes in the thermoelastic loss of the uncoated substrate, which is expected to dominate the loss of cSi disks and cantilevers in this temperature range. However, as measurements of cSi cantilevers coated with the *Lower Stack* show good agreement with single layer predictions, this suggests that this effect is greater for cSi disk substrates. It is speculated that the differences observed in ϕ_{coating} from different substrates can be attributed in part to the bi-metallic strip effect between 200 K and 300 K. The total stress of each coated sample will differ depending on the thicknesses of the coating and substrate, as well as their respective Young's moduli. As work carried out by M. Fletcher [203] has shown that the curvature of a coated sample can change the loss of a coated substrate and the thermoelastic loss of an uncoated cSi cantilever (at 293 K) it follows that the same effect can also apply to sufficiently thin cSi disks with the magnitude of the effect following as a function of coating and substrate thicknesses. The same work also concluded that due to the thickness of the silica disks used in these measurements, the effect of curvature on coating loss was negligible, giving further confidence to these measurements. To correctly estimate the loss of coated, thin cSi samples, the change in curvature due to the bi-metallic strip effect must be taken into consideration, with this analysis falling beyond the scope of this thesis.

8.5 Thermal Noise Implications

Here, the coating thermal noise of a HR stack of the *Full Stack* design, used in a cryogenic GW detector, will be analysed based on the measurements presented in the previous section. The mechanical loss from each coating used in this section are detailed at the key temperatures of 295 K, 123 K and 10 K in Tab. 8.3 and 8.4, each of the relevant material properties are detailed in Tab. 6.1. While there are relatively large changes in the mechanical losses of the single layer coating materials, less variations in $\phi_{\text{coating}}(\text{T})$ are observed in Tab. 8.3 for the *Full Stack*. This could suggest that the some methods of mechanical dissipation of these materials are suppressed when deposited in this configuration.

Temperature [K]	*Upper Stack*	*Lower Stack*	*Full Stack*	aSi	Ta_2O_5	SiO_2
295	3.4×10^{-4}	8.8×10^{-5}	3.4×10^{-4}	1.7×10^{-4}	7.7×10^{-4}	4.1×10^{-4}
123	4.0×10^{-4}	6.8×10^{-5}	3.3×10^{-4}	8.7×10^{-5}	7.7×10^{-4}	6.0×10^{-5}
20	5.0×10^{-4}	2.0×10^{-4}	4.4×10^{-4}	3.4×10^{-5}	7.1×10^{-4}	4.5×10^{-4}
10	4.1×10^{-4}	1.8×10^{-4}	3.9×10^{-4}	2.8×10^{-5}	6.1×10^{-4}	3.0×10^{-4}

Table 8.3: Levels of mechanical loss of the *Full Stack*, *Lower Stack*, *Upper Stack* and its constituent single layer materials measured on cantilevers at 295 K, 123 K and 10 K.

Temperature [K]	Upper Stack (CryoGeNS)	Lower Stack (CryoGeNS)	Upper Stack (Pulse Tube)	Lower Stack (Pulse Tube)
295	-	-	3.7×10^{-4}	-
123	4.9×10^{-4}	4.3×10^{-5}	4.9×10^{-4}	8.5×10^{-5}
20	-	-	6.3×10^{-4}	1.7×10^{-4}
10	-	-	6.5×10^{-4}	2.1×10^{-4}

Table 8.4: Levels of coating loss from the *Lower Stack Upper Stack* materials calculated from measurements of coated 2" cSi disk substrates, measured using the CryoGeNS and Pulse Tube cryostats at 295 K, 123 K and 10 K.

The levels of $\phi(\text{T})$ from coatings measured on different substrates will be determined from single layer components in their as-deposited states in the case of cantilever substrates, and from the *Upper Stack* and *Lower Stack* coatings for disk

substrates. As third generation gravitational wave detectors such as the Einstein Telescope - LF [78] and LIGO [104] plan to operate at 10 K and 123 K respectively, the thermal noise produced by the *Full Stack* coating will be evaluated at these temperatures. As other third generation detectors may also plan to operate at 20 K this temperature is also of interest.

8.5.1 Coating Performance in LIGO Voyager

Coating thermal noise can be calculated using the approach derived by Hong et al. [132] and Yam et al. [134]. As the method for calculating coating thermal noise described by Yam [134], requires that the mechanical loss of each coating layer linked to its position in the coating, this requires a small intervention. As it is not easily possible to extract the loss of an individual layer in a HR coating stack after its' deposition, the measured losses of the *Lower Stack* and the *Upper Stack* can instead be used. This simplification states that each coating layer contained within the *Lower Stack* has an identical mechanical loss, but each layer retains its optical properties. This is also true for coating positions in the *Upper Stack*, allowing the light propagation through the coating to be correctly calculated. For cantilever substrates these calculations utilise the losses of each single layer coating, using Eq. 6.5 while values of Brownian thermal noise use the measured losses from the *Upper* and *Lower* stacks.

From the recent work published by Adhikari et al. [104] the proposed LIGO Voyager HR coating design is composed of aSi and SiO_2 layers where their respective thicknesses have been optimised for $\lambda_{\mathrm{IFO}} = 2000\,\mathrm{nm}$. In the recent works published on LIGO Voyager [104], the mechanical losses used to calculate the coating thermal noise of the Voyager coating design are $\phi_{\mathrm{silica}} \leq 1\times10^{-5}$ and $\phi_{\mathrm{aSi}} = 1\times10^{-4}$. While it is reasonable to assume that a high energy coating deposition method such as IBS is able to produce low loss silica, the technology required to produce aSi with this magnitude of loss has not currently been proven. As such these values will be used as the design requirement for LIGO Voyager and will be compared to a quarter wave stack of equivalent reflectivity to the LIGO Voyager design (R = 99.9998%), but composed of RLVIP aSi and SiO_2.

The equivalent RLVIP Voyager coating will be compared to the thermal noise

contributions of the *Full Stack* material described in the previous section. It should also be noted that in these calculations the values of ϕ_{bulk} and ϕ_{shear} are considered equal ($\phi_{\text{shear}} = \phi_{\text{bulk}}$).

Figure 8.25 then compares the thermal noise of the as-deposited *Full Stack* coating from (a) measurements of its component stacks on cSi disk and (b) measurements of single layers on cSi cantilever substrates as a function of temperature. These values are pitted against the equivalent RLVIP Voyager coating and an equivalent RLVIP 'aLIGO' coating, which contains 19 bi-layers of quarter wave thick RLVIP SiO_2-Ta_2O_5 (R= 99.9999%). The thermal noise of both equivalent coatings will be calculated from the measurements of single layer coatings shown in section 8.0.1. An λ_{IFO} =2000 nm of laser beam radius of 8.4 cm on the front surface of the optic is assumed in each case, and thermal noise will be evaluated at $f = 100$ Hz.

In order for a coating stack to be viable for use as a HR ETM / ITM coating in a gravitational wave detector, the level of thermal noise produced by the coating and its absorption of the main interferometer laser light must be reduced. We first consider the thermal noise behaviour of the RLVIP aLIGO equivalent. - expected to absorb ≈ 3 ppm from measurements of the *Upper Stack* at 2000 nm (see Chapter 4). The thermal noise it produces decreases due to the reduction in temperature. However, at 123 K the coating thermal noise is a factor of 5 above the Voyager requirement.

Out of the three HR coating designs in this study the Brownian thermal noise produced by the RLVIP aSi-SiO_2 stack is the lowest, ≥ 20 % lower than the Voyager thermal noise requirements. However, the optical absorption of this coating is 181 ppm, much higher than tolerable (≤ 1 ppm) in a cryogenic detector.

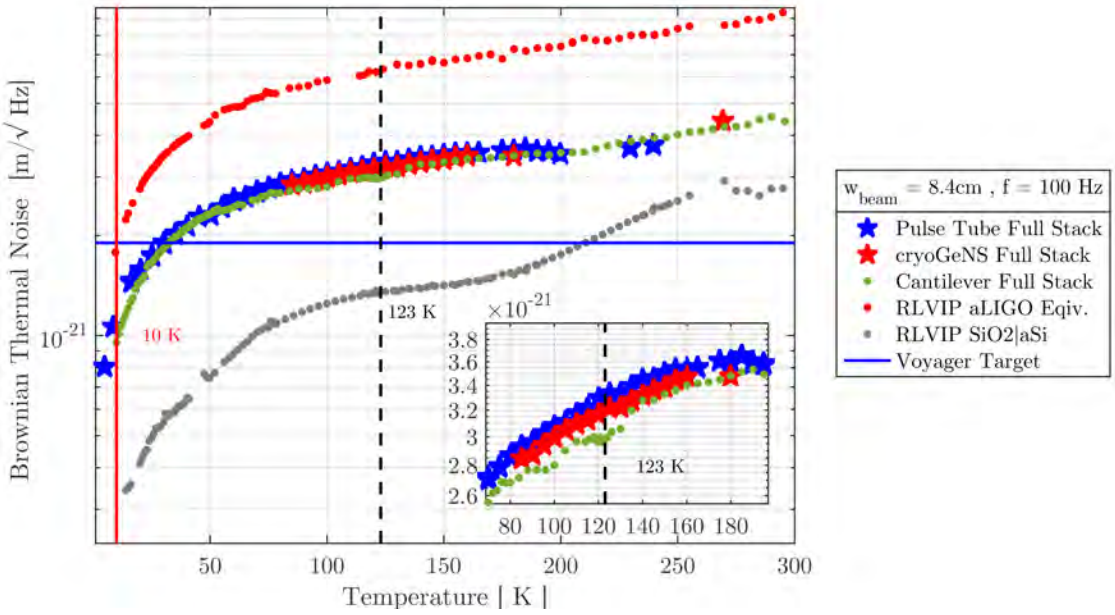

Figure 8.24: LIGO Voyager: Calculated levels of Brownian coating thermal noise from the as-deposited *Full Stack* HR stack from cSi cantilever and disk measurements. This is compared with the values LIGO Voyager design coating design (blue line), and equivalent RLVIP Voyager and RLVIP aLIGO coatings.

The level of Brownian thermal noise produced by the *Full Stack* coating lies directly between the RLVIP aSi-SiO_2 and aLIGO equivalent coatings. The reduction in Brownian thermal noise with decreasing temperature is evident in measurements taken from both cSi disk and cSi cantilever samples which show excellent agreement a each measured temperature step. Between 80 K and 160 K there are continuous predictions from single layer cantilevers and disk substrates measured using the CryoGeNS and Pulse Tube cryostats with each substrates which reflecting one and other. At the temperatures of interest for third generation detectors (20 K and 123 K) a reduction in CTN by a factor of 1.3 and 3, respectively, compared to the coating performance at room temperature, is obtained. Comparing the level of thermal noise of the *Full Stack* to the RLVIP aLIGO coating at room temperature and at 123 K a reduction of 52% and 65% respectively are shown. It can also be shown that the thermal noise improvements by changing the laser beam dimensions from the current aLIGO beam radius ($w_{beam} = 6.4$ cm) to the Voyager design $w_{beam} = 8.4$ cm gives a 23% reduction in CTN at 293 K, and by more than 40% at 123 K.

On comparing the *Full Stack* to the aSi coatings, a large difference between the thermal noise levels is exhibited by both coating stacks, with the aSi-SiO_2 stack out-performing that of the *Full Stack* by 59% at 123 K. This is as expected due to the high cryogenic loss of the Ta_2O_5 layers in the *Full Stack*. Comparing the aSi-SiO_2 absorption of 181ppm to the Full Stack value of 10ppm (see Chapter 4), the effects of test mass heating would be greatly reduced.

By comparing the Voyager coating design (shown as a blue line in Fig 8.25) to the equivalent RLVIP material, there is also a difference in thermal noise performance. If the HR stack were to be deposited by IBS, using the values of ϕ at 123 K cited in [104] a value of 1.21×10^{-21} was calculated. This value is then 5 % lower than if the coating were composed of the as-deposited RLVIP materials. As the values of IBS materials quoted in [104] are for the Voyager design coating losses for both SiO_2 and aSi, their heat treatment after deposition is inferred. A comparison then to the *Full Stack* which has not been heat treated to those which have is then unwarranted. By heat treating the RLVIP coatings the mechanical loss of their layers can be reduced (see Chapter 5) and therefore can produce lower levels of Brownian thermal noise.

8.5.2 Coating Performance in the Einstein Telescope

The Einstein Telescope (ET) is a proof of concept design study, comprised of six independent 10 km interferometers planning to be built between 100 and 300 m below the Earths' surface. Each of the nested interferometers aiming to increase on current second-generation sensitivities by a factor 100. Such a configuration of interferometers, previously suggested for aLIGO, allows for two separate detectors whose strain sensitivities are optimised for high frequency (HF) between 100 Hz and 10 kHz and low frequency (LF), between 2 Hz and 100 Hz (see Chapter 1).

Using a larger $\lambda_{\mathrm{IFO}} =$ 1500 nm laser beam ($w_{\mathrm{beam}} = 9$ cm) and a system temperature of 10 K, ET-LF hopes to reach a strain sensitivity of $\approx 3.6 \times 10^{-21}$ m/$\sqrt{\mathrm{Hz}}$ at 10 Hz. In order to reach this level of sensitivity HR coating stack with low optical absorption and mechanical losses at this temperature are required. Works published by Craig *et al.* [190] found that silica doped hafnia ($SiO_2 : HfO_2$) could be used as a low index material in a multi-material coating design to meet these requirements. For an ETM coating calculations by Craig *et al.*

showed that a coating comprised of 2 layers of SiO_2/Ta_2O_5 +10 layers of SiO_2:HfO_2/aSi would have produced a coating thermal noise of $1.9\times10^{-21}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. Scaling this to account for 2×ITMs and 2×ETMs produces a $3.5\times10^{-21}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$.

As the level of coating thermal noise produced by this coating design has been evaluated using $\lambda_{\mathrm{IFO}} = 1500\,\mathrm{nm}$ and is a viable layer design for ET-LF, comparing the level of thermal noise produced by the *Full Stack* under these conditions is practical. To provide fair comparison between the two coating stacks, the physical thickness' of each layer in the *Full Stack* will be decreased to $t = \frac{\lambda}{4}$ to produce 99.999% reflectivity at 1550 nm. The Brownian thermal noise of the equivalent RLVIP aLIGO and LIGO Voyager coatings are also shown.

Throughout the calculated temperature range the coating thermal noise of the *Full Stack* material compared to the equivalent RLVIP coatings follows the same format as shown in Fig 8.25. The blue horizontal line denotes the calculated coating thermal noise produced by the Ta_2O_5-SiO_2/HfO_2-aSi ETM coating design by Craig *et al.*. At an operating temperature of 10 K the level of coating thermal noise produced by the *Full Stack* coating approaches the predictions of the hypothetical coating produced by Craig *et al.*.

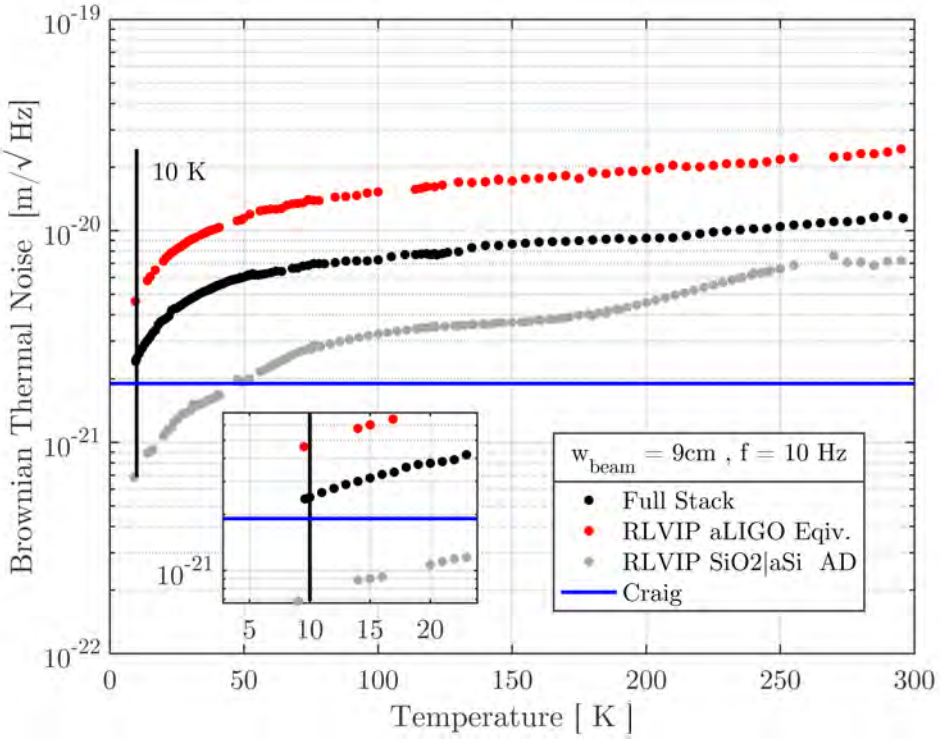

Figure 8.25: Einstein Telescope LF: Calculated levels of Brownian coating thermal noise from the as-deposited *Full Stack* HR stack from cSi cantilever measurements. This is compared with the values of the silica hafnia coating design by Craig *et al.* coating design (blue line).

The *Full Stack* reaching $2.47\times10^{-21}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$, 23% higher the SiO_2:HfO_2/aSi coating at 10 K. The work by Craig *et al.* also assumes a ∼40% lower loss of aSi, and Ta_2O_5 with ∼20% lower mechanical loss. If the *Full Stack* coating were to be produced using these values, its level of thermal noise would reduce by a further ∼6%. It can also be postulated that by doping the Ta_2O_5 layers with TiO_2 or ZrO_2 its loss could potentially reduce the CTN of a *Full Stack* design to the levels required by ET-LF[3].

To compare the optical absorption of both coatings, the calculated values of k (see Tab. 4.5) are used by assuming the absorption of the high refractive index layers in each stack dominate over SiO_2 ($k_{\mathrm{aSi}}=4.2\times10^{-4}$, $k_{\mathrm{Ta_2O_5}}=7.5\times10^{-7}$). This allowed an optical absorption of 12.5 ppm to be estimated. Compared to the 3.4 ppm calculated by Craig *et al.* for their design, the absorption is noticeably

[3]This conclusion also relies on the same materials being used for an equivalent ITM coating, following the design following the conclusions of Martin and Steinlechner [256].

higher. However, in this work the optical absorption of aSi is based on the works of Birney *et al.* [7], which produced an $k_{\mathrm{aSi}} = 1.2\times10^{-5}$ and lower absorption of Ta_2O_5 $k_{\mathrm{Tantala}} = 8\times10^{-8}$) [164]. If the *Full Stack* could be produced by IBS with similar values of optical absorption, a stack with 4.7 ppm could be achieved meeting the current ET-LF requirements.

It should be noted that while the current ET-LF absorption requirements state that the absorption of HR ETM coating is ≤ 5 ppm, this would introduce an unacceptable level of heating into the optic. As such it would be expected that a limit of ≤ 0.5 ppm, the levels currently used in aLIGO, would be more realistic.

8.6 Conclusions

In this chapter, an apparatus designed by the author to measure the mechanical loss of cSi disk substrates was used to characterise the loss of a novel multimaterial coating design as a function of temperature. Values of coating loss for the *Upper*, *Lower* and *Full Stack* coatings were compared with measurements produced in other cryogenic systems and substrate geometries. Comparisons of the different levels of loss observed in each case shows remarkable agreement between the CryoGeNS and Pulse Tube nodal supports between 80 K and 300 K. Distinct trends in the levels of mechanical loss measured from different mode shapes were observed a possible correlation between mode shape and excess surface loss was observed.

Comparing measurements of uncoated disk substrates from both cryostats shows disagreement with the predicted level of thermoelastic loss at high temperatures above 150 K. This effect is particularly prevalent in measurements of the *Lower Stack* coating, which shows comparable levels of uncoated and coated loss in this temperature range. The indistinct difference between both sets of measurements hinders calculation of the coating loss above 150 K, as this produces nonphysical, negative values. To better understand this discrepancy the change in resonant mode frequency from both disk and cantilever substrates were compared, with results suggesting that the sample temperature in the CryoGeNS system could be higher than previously estimated above 150 K. However, the second cryostat with higher temperature stability also reproduces the non-physical values of coating loss through

150 K and 293 K in some cases. This leads to the hypothesis the differences in thermal expansion coefficients between the coating and substrate causes differential stresses in the sample, inducing curvature. The induced curvature has been previously shown to alter the thermoelastic losses of cantilever substrate through the same mechanism [203]. Measurements of the relative stress on cantilever substrates then corroborate this hypothesis.

Measurements of the *Upper*, *Lower* Stacks were then used to calculate the expected level of thermal noise from the *Full Stack* coating in a third generation gravitational wave detector. Comparing the values calculated from disk and cantilever substrates to a LIGO Voyager and equivalent aLIGO coating design validates the working principle of the multi-material coating stack at 123 K. The performance of the *Full Stack* coating was also compared to a theoretical SiO_2:HfO_2/aSi coating for the ET-LF at 10 K. These results have shown the ability to reduce both the thermal noise and optical absorption of an HR coating stack by using buried layers of aSi in place of Ta_2O_5 at 123 K and 10 K.

Chapter 9

Conclusions

Since the first confirmed observation of a gravitational wave signal in 2015, there have been detections of more than 50 inspiral events from black hole pairs and binary neutron stars. These detections were made by monitoring the position of highly isolated test mass mirrors inside an interferometric gravitational wave detector. A passing gravitational wave causes shifts in the position of a test mass mirror by $<1\times10^{-18}$ m, meaning any source of noise that would inhibit such measurement must be reduced far below this level. At the most sensitive frequency band ($\sim$ 50 Hz-150 Hz) a gravitational wave detectors sensitivity is limited by coating Brownian noise, stemming from thermally driven vibrations in the coating layers.

The research presented in this thesis has involved developing measurement techniques that can characterise the key coating properties of mechanical loss and optical absorption. As the next generation of gravitational wave detectors have a wide range of desired operation temperatures, wavelengths and laser powers, coating materials that complement each set of conditions are required.

A room temperature nodal support to measure mechanical losses of uncoated silica disks at room temperature was developed, allowing multiple consecutive suspensions of samples to be carried out without opening the vacuum system. The losses measured using this apparatus were compared to existing wire suspension methods, where improved repeatability of mechanical loss measurements was observed, as a result of the reduced interactions between the sample and its

suspension.

A cryogenic nodal support to measure coated and uncoated cSi disks was developed and compared to the losses measured in a second commercial cryostat between 4 K and 300 K. These measurements highlight the challenges involved in the temperature estimation of a balanced sample. The calibration method used in the cryoGeNS system proved inadequate at low temperatures as these results were affected by the use of thermal varnish. Parallel measurements using a pulse tube cooled cryostat highlights that temperature stability over more than 5 hrs is required to allow the sample to reach thermal equilibrium at the desired temperature. Samples coated with a highly-stressed aSi/SiO_2 coating also show signs that stress and curvature induced by the bi-metallic strip effect is maybe altering the thermoelastic loss of the substrate and ratio of strain energy stored in the sample, producing un-physical losses in this temperature range.

To facilitate optical absorption studies, changes to the photothermal common-path interferometry set up were made. These changes decreased fluctuations in laser power, increasing system stability and allowing the development of the 2D absorption mapping by tracking and correction of pump laser power fluctuations.

The main focus of this thesis was the experimental verification of a novel multi-material coating design. This coating stack, known as the *Full Stack* takes advantage of the low absorption of some SiO_2-Ti:Ta_2O_5 layers to reduce the transmission of laser light to lower layers consisting of higher absorption, but lower mechanical loss, SiO_2-aSi. This allows some of the benefits of the low mechanical loss and high refractive index of aSi to be exploited without the material's high optical absorption damaging the coating's performance.

As Ta_2O_5 is one of the key materials in the HR *Full Stack* coating, its optical absorption was studied in response to different heat treatment temperatures. Incremental heat treatments between 200 °C, and 700 °C, showed the absorption of RLVIP Ta_2O_5 could be reduced by 80% at 1064 nm, showing signs of crystallisation above this temperature. Heat treatment in different environments showed that the absorption of Ta_2O_5 increased by more than 80% after multiple heat treatments

under vacuum at 400 °C. In contrast, heat treatments at the same temperature in air reduced its absorption by more than 85%. Large changes in the absorption of these samples were observed over several years when stored in air. These results suggest that Ta_2O_5 deposited by RLVIP is 'oxygen-poor', with changes in oxygen concentration dictating changes in optical absorption. This study is planned to continue through collaboration with Stanford University.

The changes in absorption of Ta_2O_5 with increasing heat treatment temperature were used as a basis for the heat treatments to be carried out on the novel multi-material coating design. This work was the first confirmation of the multi-material concept, demonstrating that layers of low mechanical loss, high optical absorption, aSi can reduce Brownian thermal noise in HR coating mirrors while minimising the impact of the absorption.

The mechanical loss and optical absorption (at wavelengths of 1064 nm, 1550 nm and 2000 nm) of the multi-material coating and it's two sub-stacks were studied as a function of heat treatment temperature. In each case, an optimum heat treatment temperature of 500°C was observed with no visible signs of delamination or defect. The optical absorption of the *Full Stack* coating was reduced to 8 ppm, 95%, lower than the absorption of the *Lower Stack* (SiO_2-aSi) coating (181 ppm) at $\lambda = 2000$ nm. The mechanical loss of the *Full Stack* was reduced by 26% ($\phi_{\text{coating}} = 2\times10^{-4}$) after heat treatment at this temperature, which due to the additional layers of aSi showed more than a 40% lower loss than 5 bi-layers of SiO_2-Ta_2O_5.

Measurements of these coatings at cryogenic temperatures produced the first demonstration that a multi-material coating design can be used to reduced coating Brownian thermal noise of a SiO_2-Ta_2O_5 stack at low temperatures. Calculations from measurements of the *Full Stack* coating design predicted a factor of 1.3 reduction in CTN at 123 K, compared to a SiO_2-Ta_2O_5 HR stack with the same reflectivity. Further cooling widens the gap with a factor 1.9 reduction gained by operating at 10 K. While the optical absorption of the *Full Stack* is too high for direct implementation into next-generation gravitational wave detectors ($\approx$8 ppm at 2000 nm), this absorption is not intrinsic to the multi-material approach. It could be reduced further by adding additional layers of SiO_2-Ta_2O_5 to the HR stack, at the cost of a slight decrease in thermal noise performance. Experimental verification of

the multi-material coating approach has opened the door to other coating designs, which could prove fundamental to third-generation gravitational wave detectors.

By monitoring the effects of optical absorption of LIGO Livingston test mass coatings, a method which allows possible changes in coating absorption over long periods to be calculated from resonant frequency shifts of each optic. Comparing the frequency shifts of three mechanical modes of each test mass, the average trend in absorption decreased for each optic between January and September in 2017 with a distinct increase coincident with maintenance inside the vacuums system, in May 2017. Throughout the measurement, one mirror (ITM-X) showed statically different trends in absorption compared to all other mechanical modes measured on each mirror, which could suggest optic misalignment or a region of degrading coating material. This technique will continue to be developed and used for long-term monitoring of mirror absorption.

Coating thermal noise is one of the major obstacles to be overcome in the near and longer-term future of gravitational astronomy. In the absence, currently, of materials with sufficiently low mechanical loss and optical absorption, multimaterial coating designs which have been experimentally verified in this thesis appear likely to play an important role in the future of the field, enabling these thermal noise challenges to be met, and thus helping to enable the next generations of gravitational wave detectors to reach their full potential.

Bibliography

[1] H. Wang, C. Blair, M. Dovale Álvarez, A. Brooks, M. F. Kasprzack, J. Ramette, P. M. Meyers, S. Kaufer, B. O'Reilly, C. M. Mow-Lowry, A. Freise, Thermal modelling of Advanced LIGO test masses, Classical and Quantum Gravity 34 (11) (2017) 115001. http://arxiv.org/abs/1612.02866v2

[2] D. Coyne (ed.), LIGO Vacuum Compatible Materials List (2014). https://dcc.ligo.org/LIGO-E960050/public

[3] D. Shoemaker, Advanced LIGO reference design v2, LIGO Document (March) (2011) 1–49. https://dcc-llo.ligo.org/public/0001/M060056/000/M060056-00.doc

[4] The Virgo Collaboration, Aerial view of virgo (2019). http://public.virgo-gw.eu/virgo-in-a-nutshell/

[5] C. Blair, Parametric Instability in Gravitational Wave Detectors, Ph.D. thesis, University Of Western Australia (2017). https://dcc.ligo.org/LIGO-P1700242

[6] Tafelmaier, Tafelmaier Dünnschicht Technik GmbH Rosenheim Ion-Plating Coating (2018). http://www.tafelmaier.de/index.php?lng=en

[7] R. Birney, J. Steinlechner, Z. Tornasi, I. W. Martin, S. Macfoy, et al., Exploring the fundamental limits of NIR absorption in amorphous silicon 1 1–6. https://dcc.ligo.org/LIGO-P1800148

[8] B. P. Abbott, Prospects for observing and localizing gravitational-wave transients with advanced LIGO and advanced Virgo, Living Reviews in Relativity 19 (2016) 1–39. https://arxiv.org/abs/1304.0670

[9] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, Allemandou, others., Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical and Quantum Gravity 32 (2) (2015) 024001. http://arxiv.org/abs/1408.3978v3

[10] K. Somiya, Detector configuration of KAGRA-the Japanese cryogenic gravitational-wave detector, Classical and Quantum Gravity 29 (12) (2012) 124007. http://arxiv.org/abs/1111.7185v2

[11] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, others., GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, Physical Review X 9 (3) (2019) 31040. https://doi.org/10.1103/PhysRevX.9.031040

[12] The LIGO Scientific Collaboration, The Virgo Collaboration, Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run (2021). https://arxiv.org/abs/2010.14527

[13] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, Properties of the Binary Black Hole Merger GW150914, Physical Review Letters 116 (2016) 241102. https://link.aps.org/doi/10.1103/PhysRevLett.116.241102

[14] S. W. Hawking, W. Israel, Three Hundred Years of Gravitation, 1989.

[15] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, D. H. Hartmann, How massive single stars end their life, The Astrophysical Journal 591 (1) (2003) 288–300. http://dx.doi.org/10.1086/375341

[16] B. Haskell, R-modes in neutron stars: Theory and observations, International Journal of Modern Physics E 24 (09) (2015) 1541007. http://dx.doi.org/10.1142/S0218301315410074

[17] G. Ushomirsky, R-modes in accreting and young neutron stars (2003) 284–295. https://core.ac.uk/download/pdf/216153483.pdf

[18] B. F. Schutz, Gravitational-wave sources, Classical and Quantum Gravity 13 (11A) (1996) A219–A238. https://doi.org/10.1088/0264-9381/13/11a/031

[19] J. Bayley, C. Messenger, G. Woan, Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals, Physical Review D 100 (2). http://dx.doi.org/10.1103/PhysRevD.100.023006

[20] R. J. Dupuis, G. Woan, Bayesian estimation of pulsar parameters from gravitational wave data, Physical Review D 72 (2005) 102002. https://link.aps.org/doi/10.1103/PhysRevD.72.102002

[21] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, others., Searches for periodic gravitational waves from unknown isolated sources and scorpius x-1: Results from the second LIGO science run, Physical Review D 76 (2007) 082001. https://link.aps.org/doi/10.1103/PhysRevD.76.082001

[22] A. Das, S. S. Dave, O. Ganguly, A. M. Srivastava, Pulsars as Weber gravitational wave detectors, Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics 791 (2019) 167–171. https://doi.org/10.1016/j.physletb.2019.02.031

[23] LIGO Burst Source Working Group, Gravitational Wave Burst Search, Tech. rep., LIGO Lab (2004). https://labcit.ligo.caltech.edu/{~}ajw/bursts/BurstPropFinal.pdf

[24] The LIGO Scientific Collaboration, the Virgo Collaboration, Search for Gravitational Waves Associated with Gamma-Ray Bursts Detected by Fermi and Swift During the LIGO-Virgo Run O3a (2020) 1–24. http://arxiv.org/abs/2010.14550

[25] N. Christensen, Stochastic gravitational wave backgrounds, Reports on Progress in Physics 82 (1) (2018) 016903. http://dx.doi.org/10.1088/1361-6633/aae6b5

[26] B. Schutz, Determining the hubble constant from gravitational wave observations, Nature 323 (6086) (1986) 310–311. https://ui.adsabs.harvard.edu/abs/1986Natur.323..310S

[27] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, Observation of gravitational waves from a binary black hole merger, Physical Review Letters 116 (2016) 061102. https://link.aps.org/doi/10.1103/PhysRevLett.116.061102

[28] B. P. Abbott, et al., GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, Physical Review Letters 116 (24) (2016) 1–14. http://arxiv.org/abs/1602.03837v1

[29] B. P. Abbott, et al., GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2, Physical Review Letters 118 (22) (2017) 1–17. http://arxiv.org/abs/1706.01812v2

[30] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., Gw170608: Observation of a 19 solar-mass binary black hole coalescence, The Astrophysical Journal 851 (2) (2017) L35. http://dx.doi.org/10.3847/2041-8213/aa9f0c

[31] The LIGO Scientific Collaboration, The Virgo Collaboration, GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence, Physical Review Letters 119 (14) (2017) 1–16. http://arxiv.org/abs/1709.09660v3

[32] J. Driggers, S. Vitale, A. Lundgren, M. Evans, K. Kawabe, S. Dwyer, K. Izumi, R. Schofield, A. Effler, D. Sigg, et al., Improving astrophysical parameter estimation via offline noise subtraction for advanced ligo, Physical Review D 99 (4). http://dx.doi.org/10.1103/PhysRevD.99.042001

[33] H. Y. G. Billingsley, L. Zhang, Characterisation of advanced LIGO core optics, Vol. 66, American Society for Precision Engineering (ASPE), 2017. http://dx.doi.org/10.1103/PhysRevD.99.042001

[34] B. J. Meers, Recycling in laser-interferometric gravitational-wave detectors, Physical Review D 38 (1988) 2317–2326. https://link.aps.org/doi/10.1103/PhysRevD.38.2317

[35] M. Evans, S. Gras, P. Fritschel, J. Miller, L. Barsotti, D. Martynov, A. Brooks, D. Coyne, R. Abbott, R. X. Adhikari, et al., Observation of parametric instability in advanced ligo, Physical Review Letters 114 (16). http://dx.doi.org/10.1103/PhysRevLett.114.161102

[36] S. Biscans, S. Gras, C. D. Blair, J. Driggers, M. Evans, P. Fritschel, T. Hardwick, G. Mansell, Suppressing parametric instabilities in LIGO using

low-noise acoustic mode dampers, Physical Review D 100 (2019) 122003. https://link.aps.org/doi/10.1103/PhysRevD.100.122003

[37] M. Tse, H. Yu, N. Kijbunchoo, A. Fernandez-Galiana, P. Dupej, et al., Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy, Physical Review Letters 123 (2019) 231107. https://link.aps.org/doi/10.1103/PhysRevLett.123.231107

[38] D. J. Ottaway, P. Fritschel, S. J. Waldman, Impact of upconverted scattered light on advanced interferometric gravitational wave detectors, Optics Express 20 (8) (2012) 8329–8336. http://www.opticsexpress.org/abstract.cfm?URI=oe-20-8-8329

[39] D. Aisa, S. Aisa, C. Campeggi, M. Colombini, A. Conte, L. Farnesini, et al., The advanced virgo monolithic fused silica suspension, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 824 (2016) 644 – 645, frontier Detectors for Frontier Physics: Proceedings of the 13$^{\text{th}}$ Pisa Meeting on Advanced Detectors. http://www.sciencedirect.com/science/article/pii/S016890021501092X

[40] L. Naticchioni, The payloads of advanced Virgo : current status and upgrades, Journal of Physics: Conference Series 957 (2018) 012002. https://doi.org/10.1088/1742-6596/957/1/012002

[41] F. Acernese, M. Agathos, L. Aiello, A. Allocca, A. Amato, et al., Increasing the astrophysical reach of the advanced Virgo detector via the application of squeezed vacuum states of light, Physical Review Letters 123 (2019) 231108. https://link.aps.org/doi/10.1103/PhysRevLett.123.231108

[42] H. Vahlbruch, M. Mehmet, K. Danzmann, R. Schnabel, Detection of 15 dB Squeezed States of Light and their Application for the Absolute Calibration of Photoelectric Quantum Efficiency, Physical Review Letters 117 (2016) 110801. https://link.aps.org/doi/10.1103/PhysRevLett.117.110801

[43] J. Oppenheimer, G. Volkoff, On Massive Neutron Cores, Physical Review Journals 55 (4) (1939) 374–381. http://www.insa.nic.in/writereaddata/UpLoadedFiles/IJPAM/20005a87_195.pdf

[44] LIGO Scientific Collaboration, Virgo Collaboration, GBM, Fermi and INTEGRAL Collaboration, IceCube Team, AstroSat Cadmium Zinc Telluride Imager Collaboration, IPN Collaboration, The Insight-Hxmt Collaboration, ANTARES Collaboration, The Swift Team, AGILE Team, The 1M2H Collaboration, The Dark Energy Camera GW-EM Collaboration, the DES Collaboration, The DLT40 and GRAWITA Team, GRAvitational Wave Inaf Collaboration, The Fermi Large Area Telescope and ATCA Array, Australia Telescope Compact and ASKAP Pathfinder, Australian SKA Group, Las Cumbres Observatory and OzGrav and DWF and AST3 Collaborations, CAASTRO Collaboration, The VINROUGE Collaboration, MASTER and J-GEM and GROWTH and JAGWAR and NRAO, Caltech - and TTU-NRAO Collaborations, NuSTAR and Pan-STARRS Team, The MAXI Consortium, TZAC Collaboration, KU Telescope, Nordic Optical and EPESSTO and GROND University, Texas Tech Group, SALT and TOROS Collaboration, Transient Robotic Observatory of the South Collaboration, The BOOTES and MWA Array, Murchison Widefield Collaboration, The CALET Collaboration, IKI-GW Follow-up Collaboration, H. E. S. S. Collaboration, LOFAR and LWA Array, Long Wavelength Collaboration, HAWC Collaboration, The Pierre Auger Collaboration, ALMA Team, Euro VLBI Collaboration, Pi of the Sky University, The Chandra Team at McGill and DFN Network, Desert Fireball and ATLAS and Survey, High Time Resolution Universe and RIMAS and RATIR and Africa/MeerKAT, SKA South, Multi-messenger observations of a binary neutron star merger, The Astrophysical Journal 848 (2) (2017) L12. https://doi.org/10.3847/2041-8213/aa91c9

[45] S. L. Shapiro, S. A. Teukolsky, A. P. Lightman, Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects, Vol. 36, 1983. http://physicstoday.scitation.org/doi/10.1063/1.2915325

[46] S. Chandrasekhar, The Dynamical Instability of Gaseous Masses Approaching The Schwarzchild Limit in General Relativity, Astrophysical Journal 140 (3) (1964) 417. http://adsabs.harvard.edu/full/1964ApJ...140..417C

[47] C. Meegan, G. Lichti, P. N. Bhat, E. Bissaldi, M. S. Briggs, V. Connaughton, R. Diehl, G. Fishman, J. Greiner, A. S. Hoover, A. J. Van Der Horst, A. Von Kienlin, R. M. Kippen, C. Kouveliotou, S. McBreen, W. S. Paciesas,

R. Preece, H. Steinle, M. S. Wallace, R. B. Wilson, C. Wilson-Hodge, The fermi gamma-ray burst monitor, Astrophysical Journal 702 (1) (2009) 791–804. http://arxiv.org/abs/0908.0450v1

[48] F. Observatory, GCN/FERMI NOTICE: 524666471, Tech. rep. (2018). https://gcn.gsfc.nasa.gov/other/524666471.fermi

[49] E. P. J. van den Heuvel, Gamma-ray observatory integral reloaded. https://arxiv.org/abs/1709.04034

[50] M. M. Kasliwal, et al., Illuminating gravitational waves: A concordant picture of photons from a neutron star merger, Science 358 (6370) (2017) 1559–1565. http://arxiv.org/abs/1710.05436v1

[51] N. Gehrels, G. Chincarini, P. Giommi, K. O. Mason, J. A. Nousek, A. A. Wells, et al., The Swift Gamma-Ray Burst Mission, The Astrophysical Journal 611 (2) (2004) 1005–1020. https://doi.org/10.1086/422091

[52] P. S. Cowperthwaite, E. Berger, V. A. Villar, B. D. Metzger, et al., The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. II. UV, Optical, and Near-IR Light Curves and Comparison to Kilonova Models 17. http://arxiv.org/abs/1710.05840$%$0Ahttp://dx.doi.org/10.3847/2041-8213/aa8fc7

[53] C. McCully, D. Hiramatsu, D. A. Howell, G. Hosseinzadeh, I. Arcavi, et al., The Rapid Reddening and Featureless Optical Spectra of the optical counterpart of GW170817, AT 2017gfo, During the First Four Days 32. http://arxiv.org/abs/1710.05853$%$0Ahttp://dx.doi.org/10.3847/2041-8213/aa9111

[54] M. Nicholl, E. Berger, D. Kasen, B. D. Metzger, J. Elias, et al., The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. III. Optical and UV Spectra of a Blue Kilonova From Fast Polar Ejecta, The Astrophysical Journal Letters 848 (2) (2017) L18. http://arxiv.org/abs/1710.05456$%$0Ahttp://dx.doi.org/10.3847/2041-8213/aa9029

[55] E. Pian, P. D'Avanzo, S. Benetti, M. Branchesi, E. Brocato, S. Campana, E. Cappellaro, S. Covino, V. e. a. D'Elia, Spectroscopic identification of r-process nucleosynthesis in a double neutron-star merger, Nature 551 (7678) (2017) 67–70. http://dx.doi.org/10.1038/nature24298

[56] A. J. Levan, J. D. Lyman, N. R. Tanvir, J. Hjorth, I. Mandel, Stanway, et al., The environment of the binary neutron star merger GW170817 28. http://arxiv.org/abs/1710.05444$%$0Ahttp://dx.doi.org/10.3847/2041-8213/aa905f

[57] N. R. Tanvir, A. J. Levan, C. González-Fernández, O. Korobkin, et al., The emergence of a lanthanide-rich kilonova following the merger of two neutron stars, The Astrophysical Journal 848 (2) (2017) L27. https://doi.org/10.3847/2041-8213/aa90b6

[58] J. Weber, Detection and generation of gravitational waves, Physical Review 117 (1960) 306–313. https://link.aps.org/doi/10.1103/PhysRev.117.306

[59] P. R. Saulson, Fundamentals of Interferometric Gravitational Wave Detectors, 2nd Edition, WORLD SCIENTIFIC, 2017. https://www.worldscientific.com/doi/abs/10.1142/10116

[60] J. Weber, Gravitational Radiation from the Pulsars, Physical Review Letters 21 (1968) 395–396. https://ui.adsabs.harvard.edu/abs/1968PhRvL..21..395W/abstract

[61] J. Weber, Evidence for Discovery of Gravitational Radiation, Physical Review Letters 22 (24) (1969) 1320–1324. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.22.1320

[62] R. W. P. Drever, J. Hough, R. Bland, G. W. Lessnoff, Search for Short Bursts of Gravitational Radiation, Nature 246 (5432) (1973) 340–344. https://doi.org/10.1038/246340a0

[63] V. T. Gertsenshtein M. E., Pustovoit, On the detection of low frequency gravitational waves, Journal of Experimental and Theoretical Physics U.S.S.R 16 (1962) p605. http://www.jetp.ac.ru/cgi-bin/e/index/e/16/2/p433?a=list

[64] A. Vinante, R. Mezzena, G. Prodi, S. Vitale, M. Cerdonio, M. Bonaldi, P. Falferi, Thermal noise in a high q ultracryogenic resonator, Review of Scientific Instruments 76. https://aip.scitation.org/doi/10.1063/1.1947927

[65] A. A. Michelson, E. W. Morley, On the relative motion of the earth and the luminiferous ether, American Journal of Science s3-34 (203) (1887) 333–345. https://www.ajsonline.org/content/s3-34/203/333

[66] F. A. E. Pirani, Republication of: On the physical significance of the Riemann tensor, General Relativity and Gravitation 41 (5) (2009) 1215–1232. https://doi.org/10.1007/s10714-009-0787-9

[67] The Nobel Committee for Physics, The Laser Interferometer Gravitational-Wave Observatory and the First Direct Observation of Gravitational Waves, Tech. rep. (2017). https://www.nobelprize.org/uploads/2018/06/advanced-physicsprize2017-1.pdf

[68] H. Ward, J. Hough, G. P. Newton, B. J. Meers, N. A. Robertson, S. Hoggan, G. A. Kerr, J. B. Mangan, R. W. Drever, Laser Interferometric Sensing Techniques for Very Small Displacements—With Applications to Gravitational Radiation Detectors, IEEE Transactions on Instrumentation and Measurement 34 (2) (1985) 261–265. https://ieeexplore.ieee.org/document/4315319

[69] H. Billing, K. Maischberger, A. Rüdiger, R. Schilling, L. Schnupp, W. Winkler, An argon laser interferometer for the detection of gravitational radiation, Journal of Physics E: Scientific Instruments 12 (11) (1979) 1043–1050. https://iopscience.iop.org/article/10.1088/0022-3735/12/11/010/pdf

[70] J. Hough, R. Hutchins, J. Logan, A. McLaren, M. Plissi, N. A. Robertson, S. Rowan, K. A. Strain, S. Twyford, Developments in Isolation, Suspension and Thermal Noise issues for GEO 600, Tech. rep., University of Glasgow (1996). https://dcc-llo.ligo.org/DocDB/0073/P960019/000/P960019-00.pdf

[71] The LIGO Scientific Collaboration, LIGO II Suspension: Reference Design, Tech. rep. (2000). https://dcc.ligo.org/DocDB/0026/T000012/000/T000012-00.pdf

[72] N. A. Robertson, M. Barton, G. Cagnoli, C. A. Cantley, D. Coyne, D. Crooks, et al., Advanced LIGO Suspension System Conceptual Design, Tech. rep. (2002). https://labcit.ligo.caltech.edu/{~}janeen/T010103-04{_}Concpt{_}Design.pdf

[73] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, Acernese, others., Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, Living Reviews in Relativity 21 (1) (2018) 3. https://arxiv.org/abs/1304.0670

[74] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, et al., Advanced LIGO, Classical and Quantum Gravity 32 (7). http://arxiv.org/abs/1411.4547v1

[75] J. G. Rollins, E. Hall, C. Wipf, L. McCuller, K. Kuns, R. Adhikari, et al., Gravitational wave interferometer noise calculator. https://git.ligo.org/gwinc/pygwinc

[76] P. R. Saulson, Terrestrial gravitational noise on a gravitational wave antenna, Physical Review D 30 (1984) 732–736. https://link.aps.org/doi/10.1103/PhysRevD.30.732

[77] S. A. Hughes, K. S. Thorne, Seismic gravity-gradient noise in interferometric gravitational-wave detectors, Physical Review D 58 (12). http://dx.doi.org/10.1103/PhysRevD.58.122002

[78] Einstein Telescope Collaboration, Einstein gravitational wave Telescope Conceptual Design Study (2011) 1–454. http://www.et-gw.eu

[79] F. Matichard, B. Lantz, R. Mittleman, K. Mason, J. Kissel, J. Mciver, B. Abbott, Seismic isolation of Advanced LIGO gravitational waves detectors : Review of strategy, instrumentation, and performance, Classical and Quantum Gravity (2015) 1–114. http://arxiv.org/abs/1502.06300v4

[80] A. V. Cumming, L. Cunningham, G. D. Hammond, K. Haughian, J. Hough, S. Kroker, I. W. Martin, R. Nawrodt, S. Rowan, C. Schwarz, A. A. Van Veggel, Silicon mirror suspensions for gravitational wave detectors, Classical and Quantum Gravity 31 (2) (2014) 025017. https://iopscience.iop.org/article/10.1088/0264-9381/31/2/025017/meta

[81] J. Hough, B. J. Meers, G. P. Newton, N. A. Robertson, H. Ward, et al., Proposal for a Joint German-British Interferometric Gravitational Wave Detector, MPQ Report 147 (September). http://eprints.gla.ac.uk/114852/

[82] C. M. Caves, Quantum-mechanical radiation-pressure fluctuations in an interferometer, Physical Review Letters 45 (1980) 75–79. https://link.aps.org/doi/10.1103/PhysRevLett.45.75

[83] S. Reid, Studies of materials for future ground-based and space-based interferometric gravitational wave detectors, Ph.D. thesis, University of Glasgow (2006). http://www.physics.gla.ac.uk/~stuartr/talks/SReid_thesis_very_small.pdf

[84] J. Südbeck, S. Steinlechner, M. Korobko, R. Schnabel, Demonstration of interferometer enhancement through Einstein–Podolsky–Rosen entanglement, Nature Photonics 14 (4) (2020) 240–244. https://arxiv.org/abs/1908.09602

[85] C. M. Caves, Quantum-mechanical noise in an interferometer, Physical Review D 23 (1981) 1693–1708. https://link.aps.org/doi/10.1103/PhysRevD.23.1693

[86] S. D. Penn, A. Ageev, D. Busby, G. M. Harry, A. M. Gretarsson, K. Numata, P. Willems, Frequency and surface dependence of the mechanical loss in fused silica, Physics Letters, Section A: General, Atomic and Solid State Physics 352 (1-2) (2006) 3–6. http://arxiv.org/abs/gr-qc/0507097v1

[87] B. J. Meers, Recycling in laser-interferometric gravitational-wave detectors, Physical Review D 38 (1988) 2317–2326. https://link.aps.org/doi/10.1103/PhysRevD.38.2317

[88] D. Sigg, LUNGO – Update Large Ultra - low Noise Gravitational - wave Observatory, Tech. rep. (2015). https://dcc.ligo.org/DocDB/0121/G1501058/002/G1501058-v2.pdf

[89] B. W. Barr, Experimental Investigations into Advanced Configurations and Optical Techniques for Laser Interferometric Gravitational Wave Detectors, Phd thesis, University of Glasgow (2003). https://ui.adsabs.harvard.edu/abs/2003PhDT.......258B

[90] H. Grote, K. Danzmann, K. L. Dooley, R. Schnabel, J. Slutsky, H. Vahlbruch, First long-term application of squeezed states of light in a gravitational-wave observatory, Physical Review Letters 110 (18) (2013) 1–4. http://arxiv.org/abs/1302.2188v3

[91] R. Loudon, P. L. Knight, Squeezed light, Journal of Modern Optics 34 (6-7) (1987) 709–759. https://www.tandfonline.com/doi/abs/10.1080/09500348714550721

[92] J. Veitch, I. Mandel, B. Aylott, B. Farr, V. Raymond, C. Rodriguez, M. van der Sluys, V. Kalogera, A. Vecchio, Estimating parameters of coalescing compact binaries with proposed advanced detector networks, Physical Review D 85 (2012) 104045. https://link.aps.org/doi/10.1103/PhysRevD.85.104045

[93] R. Robie, Characterisation of the mechanical properties of thin-film mirror coating materials for use in future interferometric gravitational wave detectors, Ph.D. thesis, University of Glasgow (2018). http://theses.gla.ac.uk/30645/

[94] A. F. Brooks, B. Abbott, M. A. Arain, G. Ciani, et al., Overview of advanced LIGO adaptive optics, Applied Optics 55 (29) (2016) 8256–8265. http://ao.osa.org/abstract.cfm?URI=ao-55-29-8256

[95] K. H. Lee, G. Hammond, J. Hough, R. Jones, S. Rowan, A. Cumming, Improved fused silica fibres for the advanced LIGO monolithic suspensions, Classical and Quantum Gravity 36 (18) (2019) 0–15. https://iopscience.iop.org/article/10.1088/1361-6382/ab28bd/meta

[96] A. V. Cumming, Aspects of mirrors and suspensions for advanced gravitational wave detectors, Ph.D. thesis, University of Glasgow (2008). http://theses.gla.ac.uk/468/

[97] K. Toland, A. Conway, L. Cunningham, G. Hammond, J. Hennig, S. Hild, J. Hough, R. Jones, E. Momany, L. Perri, S. Rowan, A. Cumming, Development of a pulling machine to produce micron diameter fused silica fibres for use in prototype advanced gravitational wave detectors, Classical and Quantum Gravity 35 (16) (2018) 165004. https://iopscience.iop.org/article/10.1088/1361-6382/aacf81

[98] Datrier, L, Image of LIGO Hanford (2019). https://ldatrier.github.io/

[99] Caltech, Massachusetts Institute of Technology , LIGO Lab, Ligo livingston laboratory (2015). https://www.ligo.caltech.edu/image/ligo20150731c

[100] G. Losurdo, G. Calamai, E. Cuoco, L. Fabbroni, G. Guidi, M. Mazzoni, R. Stanga, et al., Inertial control of the mirror suspensions of the Virgo interferometer for gravitational wave detection, Review of Scientific Instruments 72 (9) (2001) 3653–3661. https://doi.org/10.1063/1.1394189

[101] K. L. Dooley, Status of GEO 600, Journal of Physics: Conference Series 610 (1) (2015) 4–7. https://arxiv.org/abs/1411.6588

[102] The LIGO Scientific Collaboration, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, Adams, et al., A gravitational wave observatory operating beyond the quantum shot-noise limit, Nature Physics 7 (12) (2011) 962–965. https://doi.org/10.1038/nphys2083

[103] H. Grote, The LIGO Scientific Collaboration, The geo 600 status, Classical and Quantum Gravity 27 (8) (2010) 084003. https://doi.org/10.1088/0264-9381/27/8/084003

[104] R. A. Adhikari, A. K, A. F. Brooks, C. Wipf, O. Agular, P. Altin, B. Barr, L. Barsotti, R. Bassiri, A. Bell, et al., A cryogenic silicon interferometer for gravitational- wave detection A cryogenic silicon interferometer for gravitational-wave detection, Classical and Quantum Gravity 37 (2020) 40. https://iopscience.iop.org/article/10.1088/1361-6382/ab9143/pdf

[105] I. W. Martin, R. Nawrodt, K. Craig, C. Schwarz, R. Bassiri, G. Harry, et al., Low temperature mechanical dissipation of an ion-beam sputtered silica film, Classical and Quantum Gravity 31 (3) (2014) 035019. http://eprints.gla.ac.uk/92858/

[106] B. P. Abbott, et al., Exploring the sensitivity of next generation gravitational wave detectors, Classical and Quantum Gravity 34 (4) (2017) 1–11. http://arxiv.org/abs/1607.08697v3

[107] S. Dwyer, D. Sigg, S. W. Ballmer, L. Barsotti, N. Mavalvala, M. Evans, Gravitational wave detector with cosmological reach, Physical Review D 91 (8). http://dx.doi.org/10.1103/PhysRevD.91.082001

[108] K. Ackley, V. B. Adya, P. Agrawal, P. Altin, others., Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global

network, Publications of the Astronomical Society of Australia 37 (2020) 1–13. http://arxiv.org/abs/2007.03128v2

[109] P. A. Seoane, et al., The Gravitational Universe (5 2013). https://arxiv.org/abs/1305.5720

[110] J. Luo, L. S. Chen, H. Z. Duan, Y. G. Gong, S. Hu, J. Ji, et al., TianQin: A space-borne gravitational wave detector, Classical and Quantum Gravity 33 (3) (2016) 1–32. https://arxiv.org/abs/1512.02076

[111] W.-H. Ruan, Z.-K. Guo, R.-G. Cai, Y.-Z. Zhang, Taiji program: Gravitational-wave sources (2020). https://arxiv.org/abs/1807.09495

[112] R. Brown, XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies , The Philosophical Magazine 4 (21) (1828) 161–173.

[113] A. Einstein, P. Drude, M. Plank, F. Kohlrausch, G. Quincke, W. C. Röntgen, E. Warburg, Contributions to Annalen der Physik, Annalen der Physik Band 4 (1901) 513–523.

[114] Compte rendu (Association française pour l'avancement des sciences), Vol. 3e Session, 1908. https://www.biodiversitylibrary.org/item/27906#page/14/mode/1up

[115] H. B. Callen, Thermodynamics And An Introduction To Thermostatistics (1985). https://www.wiley.com/en-gb/Thermodynamics+and+an+Introduction+to+Thermostatistics$%$2C+2nd+Edition-p-9780471862567

[116] R. Than, Pumpdown and Ultimate Pressures (Internal Documentation), Tech. rep. (2014). https://dcc.ligo.org/LIGO-E1400012

[117] R. E. Hooke, Lectures de Potentia Restitutiva, Or of Spring Explaining the Power of Springing Bodies, Vol. 6, John Martyn, 1678. https://books.google.co.uk/books?id=LAtPAAAAcAAJ

[118] A. S. Nowick, B. Berry, Anelastic Relaxation in Crystalline Solids, Academic Press Inc : 111 Fifth Avenue, New York , New York 10003, 1972. https:

//www.sciencedirect.com/book/9780125226509/anelastic-relaxation-in-crystalline-solids

[119] D. Mari, The Contribution of Mechanical Spectroscopy to Understanding Grain Boundary Sliding, Materials Research 21 (suppl 2) (2018) 1–11. https://www.scielo.br/scielo.php?script=sci_arttext&pid=S1516-14392018000800220&lng=en&tlng=en

[120] M. Granata, L. Balzarini, J. Degallaix, V. Dolique, R. Flaminio, D. Forest, D. Hofman, et al., Internal friction and young's modulus measurements on SiO2 and Ta2O5 films done with an ultra-high Q silicon-wafer suspension, Archives of Metallurgy and Materials 60 (1) (2015) 365–370. https://search.proquest.com/openview/7bb6a5840234aa12669c0f02f606c7d3/1?pq-origsite=gscholar&cbl=2026345

[121] G. Vajente, R. Birney, A. Ananyeva, S. Angelova, R. Asselin, et al., Effect of elevated substrate temperature deposition on the mechanical losses in tantala thin film coatings, Classical and Quantum Gravity 35 (7) (2018) 075001. http://eprints.gla.ac.uk/156965/

[122] G. Parisi, F. Sciortino, Structural glasses: Flying to the bottom, Nature Materials 12 (2) (2013) 94–95. http://dx.doi.org/10.1038/nmat3540

[123] P. R. Saulson, Thermal noise in mechanical experiments, British Journal of Urology 42 (8) (1990) 2438–2445. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.42.2437

[124] S. Mukherjee, L. S. Finn, Removing Instrumental Artifacts : Suspension Violin Modes, AIP Publishing 523 (1) (2000) 362. http://dx.doi.org/10.1063/1.1291880

[125] F. J. Raab, A. Gillespie, Thermally excited vibrations of the mirrors of laser interferometer gravitational-wave detectors, Physical Review D 52 (2) (1995) 577–585. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.52.577

[126] Y. Levin, Theory of thermal noise in optical mirrors, Optical Coatings and Thermal Noise in Precision Measurement 9781107003 (2012) 1–5. https://as

sets.cambridge.org/97811070/03385/excerpt/9781107003385_excerpt.pdf

[127] S. Rowan, J. Hough, D. R. Crooks, Thermal noise and material issues for gravitational wave detectors, Physics Letters, Section A: General, Atomic and Solid State Physics 347 (1-3) (2005) 25–32. https://www.sciencedirect.com/science/article/pii/S0375960105009412

[128] A. Amato, Low Thermal Noise Coating for New Generation Gravitational-Wave Detectors, Ph.D. thesis, Université de Lyon (2020). https://tel.archives-ouvertes.fr/tel-02475821/file/TH2019AmatoAlex2.pdf

[129] N. Nakagawa, A. M. Gretarsson, E. K. Gustafson, M. M. Fejer, Thermal noise in half-infinite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror, Physical Review D 65 (10). http://dx.doi.org/10.1103/PhysRevD.65.102001

[130] Y. T. Liu, K. S. Thorne, Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses, Physical Review D - Particles, Fields, Gravitation and Cosmology 62 (12) (2000) 10. https://arxiv.org/abs/gr-qc/0002055

[131] G. M. Harry, A. M. Gretarsson, P. R. Saulson, S. E. Kittelberger, S. D. Penn, W. J. Startin, S. Rowan, M. M. Fejer, D. R. M. Crooks, G. Cagnoli, J. Hough, N. Nakagawa, Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings, Classical and Quantum Gravity 19 (5) (2002) 897–917. https://doi.org/10.1088/0264-9381/19/5/305

[132] T. Hong, H. Yang, E. K. Gustafson, R. X. Adhikari, Y. Chen, Brownian thermal noise in multilayer coated mirrors, Physical Review D - Particles, Fields, Gravitation and Cosmology 87 (8) (2013) 1–27. http://arxiv.org/abs/1207.6145v1

[133] J. D. Ferry, H. S. Myers, Viscoelastic Properties of Polymers, Vol. 108, The Electrochemical Society, 1961. https://doi.org/10.1149/1.2428174

[134] W. Yam, S. Gras, M. Evans, Multimaterial coatings with reduced thermal noise, Physical Review D - Particles, Fields, Gravitation and Cosmology 91 (4) (2015) 1–6. https://arxiv.org/abs/1411.3234

[135] G. Cagnoli, M. Lorenzini, E. Cesarini, F. Piergiovanni, M. Granata, D. Heinert, F. Martelli, R. Nawrodt, A. Amato, Q. Cassar, J. Dickmann, S. Kroker, D. Lumaca, C. Malhaire, C. B. Rojas Hurtado, Mode-dependent mechanical losses in disc resonators, Physics Letters, Section A: General, Atomic and Solid State Physics 382 (33) (2018) 2165–2173. http://dx.doi.org/10.1016/j.physleta.2017.05.065

[136] M. R. Abernathy, G. M. Harry, F. Travasso, I. Martin, S. Reid, et al., The effects of heating on mechanical loss in tantala/silica optical coatings, Physics Letters, Section A: General, Atomic and Solid State Physics 372 (2) (2008) 87–90. https://www.sciencedirect.com/science/article/abs/pii/S0375960107010304

[137] G. Vajente, M. Fazio, L. Yang, A. Gupta, A. Ananyeva, G. Billinsley, C. S. Menoni, Method for the experimental measurement of bulk and shear loss angles in amorphous thin films, Physical Review D 101 (4) (2020) 42004. https://doi.org/10.1103/PhysRevD.101.042004

[138] M. M. Fejer, S. Rowan, G. Cagnoli, D. R. Crooks, A. Gretarsson, G. M. Harry, J. Hough, S. D. Penn, P. H. Sneddon, S. P. Vyatchanin, Thermoelastic dissipation in inhomogeneous media: Loss measurements and displacement noise in coated test masses for interferometric gravitational wave detectors, Physical Review D - Particles, Fields, Gravitation and Cosmology 70 (8). http://arxiv.org/abs/gr-qc/0402034v1

[139] V. B. Braginsky, M. L. Gorodetsky, S. P. Vyatchanin, Thermo-refractive noise in gravitational wave antennae, Physics Letters, Section A: General, Atomic and Solid State Physics 271 (5-6) (2000) 303–307. http://arxiv.org/abs/cond-mat/0008109v1

[140] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, G. Ogin, Thermo-optic noise in coated mirrors for high-precision optical measurements, Physical Review D 78 (10). http://dx.doi.org/10.1103/PhysRevD.78.102003

[141] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, et al., GW150914: The Advanced LIGO Detectors in the Era of First Discoveries, Phys. Rev. Lett. 116 (2016) 131103. https://link.aps.org/doi/10.1103/PhysRevLett.116.131103

[142] M. Granata, Coating research and development at LMA, Labratorie des Materiaux Avances, Geneva, 2017, p. 12. https://dcc.ligo.org/LIGO-G1700854

[143] G. Parisi, F. Zamponi, The ideal glass transition of hard spheres, Journal of Chemical Physics 123 (14) (2005) 144501. http://arxiv.org/abs/cond-mat/0506445v1

[144] X. Liu, D. R. Queen, T. H. Metcalf, J. E. Karel, F. Hellman, Hydrogen-free amorphous silicon with no tunneling states, Physical Review Letters 113 (2) (2014) 1–5. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.113.025503

[145] G. M. Harry, M. R. Abernathy, et al., Titania-doped tantala/silica coatings for gravitational-wave detection, Classical and Quantum Gravity 24 (2) (2006) 405–415. https://doi.org/10.1088/0264-9381/24/2/008

[146] H.-W. Pan, S.-J. Wang, L.-C. Kuo, S. Chao, M. Principe, I. M. Pinto, R. DeSalvo, Thickness-dependent crystallization on thermal anneal for titania/silica nm-layer composites deposited by ion beam sputter method, Optics Express 22 (24) (2014) 29847–29854. http://www.opticsexpress.org/abstract.cfm?URI=oe-22-24-29847

[147] L.-C. Kuo, H.-W. Pan, C.-L. Chang, S. Chao, Low cryogenic mechanical loss composite silica thin film for low thermal noise dielectric mirror coatings, Optics Letters 44 (2) (2019) 247. https://www.osapublishing.org/ol/abstract.cfm?uri=ol-44-2-247

[148] H.-W. Pan, L.-C. Kuo, S.-Y. Huang, M.-Y. Wu, Y.-H. Juang, C.-W. Lee, H.-C. Chen, T. T. Wen, S. Chao, Silicon nitride films fabricated by a plasma-enhanced chemical vapor deposition method for coatings of the laser interferometer gravitational wave detector, Physical Review D 97 (2018) 022004. https://link.aps.org/doi/10.1103/PhysRevD.97.022004

[149] J. Steinlechner, I. W. Martin, R. Bassiri, A. Bell, M. M. Fejer, J. Hough, A. Markosyan, R. K. Route, S. Rowan, Z. Tornasi, Optical absorption of ion-beam sputtered amorphous silicon coatings, Physical Review Letters D 93 (2016) 062005. https://link.aps.org/doi/10.1103/PhysRevD.93.062005

[150] M. E. Fine, H. Van Duyne, N. T. Kenney, Low-temperature internal friction and elasticity effects in vitreous silica, Journal of Applied Physics 25 (3) (1954) 402–405. https://aip.scitation.org/doi/10.1063/1.1721649

[151] V. Braginsky, V. Mitrofanov, V. Panov, R. Krotkov, Systems with Small Dissipation, Vol. 55, 1987. https://press.uchicago.edu/ucp/books/book/chicago/S/bo5973099.html

[152] J. Jiang, A. S. Mishkin, K. Prasai, M. Yazback, R. Bassiri, M. M. Fejer, H.-p. Cheng, Atomic structures modelling and mechanical loss study of amorphous ZrO_2 doped Ta_2O_5 (2020) 1–12. https://dcc.ligo.org/LIGO-P1900371

[153] S. Rowan, I. W. Martin, Substrate thermal noise, Cambridge University Press, 2012. https://www.cambridge.org/core/books/optical-coatings-and-thermal-noise-in-precision-measurement/substrate-thermal-noise/402E230A7E9C525DA2B0A6A467DB7350

[154] R. Nawrodt, A. Zimmer, T. Koettig, C. Schwarz, D. Heinert, M. Hudl, R. Neubert, M. Thürk, S. Nietzsche, W. Vodel, P. Seidel, A. Tünnermann, High mechanical Q-factor measurements on silicon bulk samples, Journal of Physics: Conference Series 122 (2008) 012008. https://iopscience.iop.org/article/10.1088/1742-6596/122/1/012008

[155] E. D. Black, An introduction to Pound–Drever–Hall laser frequency stabilization, American Journal of Physics 69 (1) (2001) 79–87. http://aapt.scitation.org/doi/10.1119/1.1286663

[156] Heraeus Quarzglas GmbH & Co., SUPRASIL®3001 and 3002 Data Sheet (2006). http://www-eng.lbl.gov/~shuman/NEXT/MATERIALS&COMPONENTS/Quartz/Suprasil_3001_and_3002.pdf

[157] The LIGO Scientific Collaboration, Advanced LIGO, Classical and Quantum Gravity 32 (7) (2015) 074001. http://arxiv.org/abs/1411.4547$%$5Cnhttp://stacks.iop.org/0264-9381/32/i=7/a=074001?key=crossref.20895763c84bce3f8929251031b2475c

[158] W. Winkler, K. Danzmann, A. Rüdiger, R. Schilling, Heating by optical absorption and the performance of interferometric gravitational-wave detectors,

Physical Review B 3 (8) (1971) 2604–2607. https://pubmed.ncbi.nlm.nih.gov/9905843/

[159] G. Billingsley, L. Zhang, ITM03; LHO ITMx Post O2 inspection, Sonoma, 2017, pp. 1–16. https://dcc.ligo.org/LIGO-G1800345

[160] C. Chree, Longitudinal Vibrations of a circular bar, The quarterly journal of pure and applied mathematics, London, 1886.

[161] H. Wang, C. Blair, M. Kasprzack, A. Brooks, P. Meyers, S. Kaufer, O'Reilly, Coating absorption estimation for Advanced LIGO (617) (2015) 883–895. https://dcc.ligo.org/DocDB/0121/T1500469/002/Coating_absorption_ETM.pdf

[162] P. Hello, J.-Y. Vinet, Analytical models of transient thermoelastic deformations of mirrors heated by high power cw laser beams, Journal de Physique 51 (20) (1990) 2243–2261. http://www.edpsciences.org/10.1051/jphys:0199000510200224300

[163] Spriricon Laser Beam Diagnositcs, Hartmann Wavefront Analyzer Tutorial, Spriricon 10885-001, Rev. D (2004) 17. https://www.researchgate.net/profile/Mohamed_Mourad_Lafifi/post/Looking_for_any_in-depth_guidelines_available_on_how_to_use_Shack-Hartman_sensor_for_thermal_lens_measurement/attachment/59d644a979197b807799fe13/AS$%$3A449637107671041$%$401484213248465/download/Hartmann+Wavefront+Analyzer+Tutorial+.pdf

[164] L. Pinard, C. Michel, B. Sassolas, L. Balzarini, J. Degallaix, V. Dolique, R. Flaminio, D. Forest, M. Granata, B. Lagrange, N. Straniero, J. Teillon, G. Cagnoli, Mirrors used in the LIGO interferometers for first detection of gravitational waves, Applied Optics 56 (4) (2017) C11–C15. http://ao.osa.org/abstract.cfm?URI=ao-56-4-C11

[165] A. Mullavey, Guardian Implementation at LLO (2013). https://dcc.ligo.org/LIGO-G1300346

[166] R. Tomassini, Blade Tip Timing and Blade Tip Clearance Measurement System Based on Magnetoresistive Sensors, Ph.D. thesis, Universitá Deglie STudi De

Padova (2017). http://paduaresearch.cab.unipd.it/10025/1/Tomassini{_}Roberto{_}Tesi.pdf

[167] S. McCormick, aLIGO Logbook Entry, Tech. rep., LIGO LLO, Interal Documentation (2017). https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=33843

[168] V. Sannibale, B. Abbott, Y. Aso, V. Boschi, D. Coyne, R. Desalvo, et al., Recent results of a seismically isolated optical table prototype designed for advanced LIGO, Journal of Physics: Conference Series 122 (2008) 1–6. https://authors.library.caltech.edu/14250/1/SANjpcs08.pdf

[169] G. Billingsley, A. Brooks, L. Zhang, HWS absorption map and Transmission map Comparison O3 Test masses, LIGO Lab, 2019, pp. 1–21. https://dcc.ligo.org/LIGO-G1900693

[170] A. Alexandrovski, M. Fejer, A. Markosian, R. Route, Photothermal common-path interferometry (PCI): new developments, Solid State Lasers XVIII: Technology and Devices 7193 (November) (2009) 71930D. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7193/71930D/Photothermal-common-path-interferometry-PCI-new-developments/10.1117/12.814813.short?SSO=1

[171] E. Hecht, Optics, Pearson Education, Incorporated, 2017. https://books.google.co.uk/books?id=ZarLoQEACAAJ

[172] O. Arnon, P. Baumeister, Electric field distribution and the reduction of laser damage in multilayers, Applied Optics 19 (11) (1980) 1853–1855. http://ao.osa.org/abstract.cfm?URI=ao-19-11-1853

[173] C. Desai, N. Chourasia, S. Dhar, Effect of heat treatment on the optical absorption of bisbte3 thin films, Materials Letters 45 (2) (2000) 116–119. https://www.sciencedirect.com/science/article/pii/S0167577X00000872

[174] D. M. Mattox, Handbook of Physical Vapor Deposition (PVD) Processing , 2nd Edition, 2009. http://www.elsevier.com/wps/find/bookdescription.cws_home/717814/description#description

[175] Tafelmaier, Ion-Plating - Tafelmaier Dünnschicht-Technik GmbH, Rosenheim, Germany (2020). https://www.tafelmaier.de/index.php?page=beschichtung#ionplating

[176] E. Cesarini, M. Lorenzini, E. Campagna, F. Martelli, F. Piergiovanni, F. Vetrano, G. Losurdo, G. Cagnoli, A gentle nodal suspension for measurements of the acoustic attenuation in materials, Review of Scientific Instruments 80 (5) (2009) 053904–053904–6. https://aip.scitation.org/doi/10.1063/1.3124800

[177] G. Vajente, A. Ananyeva, G. Billingsley, E. Gustafson, A. Heptonstall, E. Sanchez, C. Torrie, A high throughput instrument to measure mechanical losses in thin film coatings, Review of Scientific Instruments 88 (7) (2017) 073901. https://aip.scitation.org/doi/10.1063/1.4990036

[178] M. M. Fejer, S. Rowan, G. Cagnoli, D. R. M. Crooks, A. Gretarsson, G. M. Harry, J. Hough, S. D. Penn, P. H. Sneddon, S. P. Vyatchanin, Thermoelastic dissipation in inhomogeneous media: loss measurements and displacement noise in coated test masses for interferometric gravitational wave detectors, Physical Review D 70 (2004) 082003. https://link.aps.org/doi/10.1103/PhysRevD.70.082003

[179] J. Lin, T. Suzuki, D. Matsunaga, K. Hieda, Low Crystallization Temperature for Ta_2O_5 Thin Films, Japanese Journal of Applied Physics, Part 1: Regular Papers and Short Notes and Review Papers 42 (11) (2003) 7023–7024. https://iopscience.iop.org/article/10.1143/JJAP.42.7023/meta

[180] J. Y. Tewg, Y. Kuo, J. Lu, Suppression of crystallization of tantalum oxide thin film by doping with zirconium, Electrochemical and Solid-State Letters 8 (1) (2005) 27–30. https://iopscience.iop.org/article/10.1149/1.1836122

[181] R. Bassiri, K. B. Borisenko, D. J. H. Cockayne, J. Hough, I. MacLaren, S. Rowan, Probing the atomic structure of amorphous Ta_2O_5 mirror coatings for advanced gravitational wave detectors using transmission electron microscopy , Journal of Physics: Conference Series 241 (2010) 012070. https://iopscience.iop.org/article/10.1088/1742-6596/241/1/012070/pdf

[182] J. Jin, J. Liu, X. Wang, J. Guo, N. Song, Effect of color center absorption on temperature dependence of radiation-induced attenuation in optical fibers

at near infrared wavelengths, Journal of Lightwave Technology 31 (6) (2013) 839–845. https://www.osapublishing.org/jlt/abstract.cfm?uri=jlt-31-6-839

[183] H. de Waal, Internal friction of sodium disilicate-glass after ion exchanged, Physics of Chemistry of Glasses 10 (3) (1969) 108–116.

[184] P. Jussieu, Backscattering spectrometry, Ion Beam Analysis: Fundamentals and Applications 37 (2014) 57–78. https://tinyurl.com/3gnktrra

[185] P. R. Poudel, B. Rout, D. R. Diercks, J. A. Paramo, Y. M. Strzhemechny, F. D. McDaniel, Effects of thermal annealing on the formation of buried β-SiC by ion implantation, Journal of Electronic Materials 40 (9) (2011) 1998–2003. https://link.springer.com/article/10.1007/s11664-011-1695-9

[186] A. Dahshan, K. A. Aly, Determination of the thickness and optical constants of amorphous Ge-Se-Bi thin films, Philosophical Magazine 89 (12) (2009) 1005–1016. https://www.tandfonline.com/doi/abs/10.1080/14786430902835644

[187] P. G. Murray, I. W. Martin, K. Craig, J. Hough, R. Robie, S. Rowan, M. R. Abernathy, T. Pershing, S. Penn, Ion-beam sputtered amorphous silicon films for cryogenic precision measurement systems, Physical Review D - Particles, Fields, Gravitation and Cosmology 92 (6) (2015) 1–11. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.92.062001

[188] J. Steinlechner, I. W. Martin, A. S. Bell, J. Hough, M. Fletcher, P. G. Murray, R. Robie, S. Rowan, R. Schnabel, Silicon-Based Optical Mirror Coatings for Ultrahigh Precision Metrology and Sensing, Physical Review Letters 120 (26) (2018) 263602. https://doi.org/10.1103/PhysRevLett.120.263602

[189] L. Terkowski, I. W. Martin, D. Axmann, Behrendsen, et al., Influence of deposition parameters on the optical absorption of amorphous silicon thin films, Phys. Rev. Research 2 (2020) 033308. https://link.aps.org/doi/10.1103/PhysRevResearch.2.033308

[190] K. Craig, J. Steinlechner, P. G. Murray, A. S. Bell, R. o. Birney, Mirror Coating Solution for the Cryogenic Einstein Telescope, Physical Review Letters 122 (23) (2019) 231102. https://doi.org/10.1103/PhysRevLett.122.231102

[191] J. Steinlechner, I. W. Martin, J. Hough, C. Krueger, S. Rowan, R. Schnabel, Thermal Noise Reduction and Absorption Optimisation via Multi-Material Coatings (2014) 1–11. http://arxiv.org/abs/1411.3150

[192] S. C. Tait, J. Steinlechner, M. M. Kinley-Hanlon, P. G. Murray, J. Hough, G. McGhee, F. Pein, S. Rowan, R. Schnabel, C. Smith, L. Terkowski, I. W. Martin, Demonstration of the multimaterial coating concept to reduce thermal noise in gravitational-wave detectors, Physical Review Letters 125 (2020) 011102. https://link.aps.org/doi/10.1103/PhysRevLett.125.011102

[193] Corning Inc., Corning HPFS 7979, 7980, 8652, 8655 Fused Silica Optical Materials Product Information (2014) 1–8. https://www.corning.com/media/worldwide/csm/documents/5bf092438c5546dfa9b08e423348317b.pdf

[194] P. Sortais, T. André, J. Angot, S. Bouat, J. Jacob, T. Lamy, P. Sole, Frequency scaling with miniature COmpact MIcrowave and Coaxial ion sources, Review of Scientific Instruments 85 (2) (2014) 3–6. https://aip.scitation.org/doi/10.1063/1.4852275

[195] J. Li, W. Yang, J. Su, C. Yang, Effects of Deposition Temperature on Structural, Optical Properties and Laser Damage of $LaTiO_3$ Thin Films, Advances in Condensed Matter Physics 2018 (100). https://www.hindawi.com/journals/acmp/2018/7328429/

[196] K. Evans, Computational Modelling of Amorphous mirror coatings for use in Advanced Gravitational wave detectors, Ph.D. thesis, University of Glasgow (2012). http://encore.lib.gla.ac.uk/iii/encore/record/C__Rb2984182

[197] A. Amato, G. Cagnoli, M. Canepa, E. Coillet, J. Degallaix, Dolique, et al., High-reflection coatings for gravitational-wave detectors: State of the art and future developments, Journal of Physics: Conference Series 957 (1) (2018) 1–9.

[198] LIGO, Instrument Science White Paper 2018 (Feb 2018). https://dcc.ligo.org/LIGO-T1800133/public

[199] E. Çetinörgö-Goldenberg, J. E. Klemberg-Sapieha, L. Martinu, Effect of postdeposition annealing on the structure, composition, and the mechanical and Optical characteristics of niobium and tantalum oxide films, Applied Optics 51 (27) (2012) 6498–6507. https://pubmed.ncbi.nlm.nih.gov/23033019/

[200] J. Miller, L. Barsotti, S. Vitale, P. Fritschel, M. Evans, D. Sigg, Prospects for doubling the range of Advanced LIGO, Physical Review D - Particles, Fields, Gravitation and Cosmology 91 (6) (2015) 1–6.

[201] Einstein Telescope Collaboration, ET – Einstein gravitational wave Telescope Design Study (3) (2011) 1–454. http://www.et-gw.eu

[202] A. W. Leissa, Vibration of Plates, Vol. 27, 1969. https://ntrs.nasa.gov/search.jsp?R=19700009156

[203] M. Fletcher, Investigating methods and materials which can more accurately model and reduce a gravitational wave detector's coating thermal noise., Ph.D. thesis, University of Glasgow (2018). http://theses.gla.ac.uk/72473/

[204] I. W. Martin, Studies of materials for use in future interferometric gravitational wave detectors, Ph.D. thesis, University of Glasgow (2009). http://theses.gla.ac.uk/1517/

[205] P. G. Murray, Measurement of the Mechanical Loss of Test Mass Materials for Advanced Gravitational Wave Detectors, Ph.D. thesis, University of Glasgow (2008). http://theses.gla.ac.uk/565/

[206] K. Numata, Intrinsic losses in various kinds of fused silica, Classical and Quantum Gravity 19 (1697) (2002) 1697. https://iopscience.iop.org/article/10.1088/0264-9381/19/7/363

[207] S. GmbH, Precision in Measurement (2020). https://sios.de/

[208] M. Granata, K. Craig, G. Cagnoli, C. Carcy, W. Cunningham, J. Degallaix, R. Flaminio, D. Forest, M. Hart, J.-S. Hennig, J. Hough, I. MacLaren, I. W. Martin, C. Michel, N. Morgado, S. Otmani, L. Pinard, S. Rowan, Cryogenic measurements of mechanical loss of high-reflectivity coating and estimation of thermal noise, Optics Letters 38 (24) (2013) 5268. https://www.osapublishing.org/ol/abstract.cfm?uri=ol-38-24-5268

[209] D. Fuller, Coefficient of Friction (2011). https://web.mit.edu/8.13/8.13c/references-fall/aip/aip-handbook-section2d.pdf

[210] R. E. Hummel, R. E. Hummel, Thermal Expansion, Electronic Properties of Materials (2011) 439–441. https://link.springer.com/book/10.1007/978-1-4419-8164-6

[211] A. Technologies, Laser and optics Users Manual, 2002. https://www.uzimex.cz/soubory/20080403_laser_optika_cast_1.pdf

[212] K. A. Laux, A. Jean-Fulcrand, H. J. Sue, T. Bremner, J. S. Wong, The influence of surface properties on sliding contact temperature and friction for polyetheretherketone (PEEK) (2016). https://www.sciencedirect.com/science/article/pii/S0032386116308643

[213] J. Lind, P. Lindholm, J. Qin, Kassman Rudolphi, Friction and wear studies of some peek materials, Tribologia 33 (2) (2015) 20–28. https://journal.fi/tribologia/article/download/69255/30721/

[214] Curbell Plastics, Strong, stiff plastic with outstanding chemical resistance; performs over a wide range of temps (2016). https://www.curbellplastics.com/Research-Solutions/Materials/PEEK#:~:text=Strong$%$2C$%$20stiff$%$20plastic$%$20with$%2$0outstanding,mechanical$%$20strength$%$20and$%$20dimensional$%$20stability.

[215] DuPont Fluoroproducts, Teflon. http://www.rjchase.com/ptfe_handbook.pdf

[216] R. Takahashi, F. Kuwahara, E. Majorana, M. A. Barton, T. Uchiyama, et al., Vacuum-compatible vibration isolation stack for an interferometric gravitational wave detector TAMA300, Review of Scientific Instruments 73 (6) (2002) 2428. https://aip.scitation.org/doi/10.1063/1.1473225

[217] A. Amato, G. Cagnoli, M. Canepa, E. Coillet, J. Degallaix, V. Dolique, D. Forest, M. Granata, V. R. Martinez, C. Michel, L. Pinard, B. Sassolas, J. Teillon, High-reflection coatings for gravitational-wave detectors: State of the art and future developments, Journal of Physics: Conference Series 957 (1) (2018) 1–9. http://arxiv.org/abs/1712.05701v1

[218] A. Gurkovsky, S. Vyatchanin, The thermal noise in multilayer coating, Physics Letters, Section A: General, Atomic and Solid State Physics 374 (33) (2010) 3267–3274. http://dx.doi.org/10.1016/j.physleta.2010.06.012

[219] N. M. Kondratiev, A. G. Gurkovsky, M. L. Gorodetsky, Thermal noise and coating optimization in multilayer dielectric mirrors, Physical Review D - Particles, Fields, Gravitation and Cosmology 84 (2) (2011) 1–9. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.84.022001

[220] M. R. Abernathy, Mechanical properties of coating materials for use in the mirrors of interferometric gravitational wave detectors. http://theses.gla.ac.uk/3671/

[221] V. B. Braginsky, A. A. Samoilenko, Measurements of the optical mirror coating properties, Physics Letters, Section A: General, Atomic and Solid State Physics 315 (3-4) (2003) 175–177. http://arxiv.org/abs/gr-qc/0304100v1

[222] R. M. Jones, Mechanics of Composite Materials, 2nd Edition, Vol. 44, Taylor & Francis, Inc., 2018. https://www.google.co.uk/books/edition/Mechanics_Of_Composite_Materials/K05ZDwAAQBAJ?hl=en&gbpv=0

[223] Gooch and Housego, Gooch and Housego (2021). https://gandh.com/

[224] V. Kaajakari, Silicon as an anisotropic mechanical material-a tutorial (1998) 1–5. http://www.kaajakari.net/{~}ville/research/tutorials/elasticity{_}tutorial.pdf

[225] P. G. Murray, J. Steinlechner, M. I, S. C. Tait, M. Kinley-Hanlon, G. McGhee, Multimaterial coatings for 3rd generation gravitational wave detectors, 2019. https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?.submit=Identifier&docid=G1900549&version=1

[226] M. J. Hart, Amorphous Mirror Coatings for Ultra-High Precision Interferometry, Ph.D. thesis, University of Glasgow (2017). http://theses.gla.ac.uk/8407/

[227] H.-W. Pan, S.-J. Wang, L.-C. Kuo, S. Chao, M. Principe, I. M. Pinto, R. DeSalvo, Thickness-dependent crystallization on thermal anneal for titania/silica nm-layer composites deposited by ion beam sputter method, Optics Express 22 (24) (2014) 29847. https://www.osapublishing.org/oe/fulltext.cfm?uri=oe-22-24-29847&id=304844

[228] A. Schroeter, Mechanical losses in materials for future cryogenic gravitational wave detectors, Ph.D. thesis, Council of the Faculty of Physics and Astronomy from Friedrich Schiller University Jena (2008). https://d-nb.info/988296454/34

[229] E. J. Boyd, L. Li, R. Blue, D. Uttamchandani, Measurement of the temperature coefficient of Young's modulus of single crystal silicon and 3C silicon carbide below 273 K using micro-cantilevers (2013). https://www.sciencedirect.com/science/article/pii/S0924424713001969

[230] M. A. Hopcroft, W. D. Nix, T. W. Kenny, What is the young's modulus of silicon?, Journal of Microelectromechanical Systems 19 (2) (2010) 229–238. https://ieeexplore.ieee.org/document/5430873

[231] M. Levinshtein, S. Rumyantsev, M. Shur, Handbook Series on Semiconductor Parameters, WORLD SCIENTIFIC, 1996. https://www.worldscientific.com/doi/abs/10.1142/2046

[232] K. G. Lyon, G. L. Salinger, C. A. Swenson, G. K. White, Linear thermal expansion measurements on silicon from 6 to 340 K, Journal of Applied Physics 48 (3) (1977) 865–868. https://aip.scitation.org/doi/10.1063/1.323747

[233] J. E. Bishop, V. K. Kinra, Elastothermodynamic damping in laminated composites, International Journal of Solids and Structures 34 (9) (1997) 1075–1092. http://dx.doi.org/10.1016/S0020-7683(96)00085-6

[234] S. Vengallatore, Analysis of thermoelastic damping in laminated composite micromechanical beam resonators, Journal of Micromechanics and Microengineering 15 (12) (2005) 2398–2404. https://iopscience.iop.org/article/10.1088/0960-1317/15/12/023

[235] Infrared Laboratories, Liquid Cryogen Dewars Dewar Product Comparison (2021). https://www.irlabs.com/products/cryostats/liquid-cryogen-dewars/

[236] Lake Shore Cryotronics Inc., A new cryogenic diode thermometer (2003) 1620–1627. https://www.lakeshore.com/docs/default-source/product-downloads/literature/dt-670-overview.pdf?sfvrsn=51861d91{_}4

[237] Lake Shore Cryotronics Inc., User's Manual Model 336 Temperature Controller, Tech. Rep. April 2011 (2011). https://www.lakeshore.com/docs/default-source/product-downloads/336_manual0ebc9b06cbbb456491c65cf1337983e4.pdf?sfvrsn=2e8633a3_1

[238] ICEOxford, Thermal Properties of Low-Temperature Varnish ICEles03, personal correspondence (2020).

[239] L. D. Landau, E. M. Lifshitz, J. B. Sykes, W. H. Reid, E. H. Dill, Theory of Elasticity: Vol. 7 of Course of Theoretical Physics, Vol. 13, 1960. https://physicstoday.scitation.org/doi/10.1063/1.3057037

[240] AMTEK, ZYGO GPI XP Specifications (2000). https://www.lambdaphoto.co.uk/pdfs/GPI_XP_specs.pdf

[241] C. Zener, General Theory of Thermoelastic Internal Friction, Physical Review Letters 53 (1) (1938) 90–99. https://journals.aps.org/pr/pdf/10.1103/PhysRev.53.90

[242] C. Zener, Anelasticity of metals, Il Nuovo Cimento Series 10 7 (2 Supplement) (1958) 544–568. https://link.springer.com/article/10.1007$%$2FBF02751494

[243] R. Lifshitz, M. Roukes, Thermoelastic damping in micro and nanomechanical systems, Physical Review B - Condensed Matter and Materials Physics 61 (8) (2000) 5600–5609. https://journals.aps.org/prb/abstract/10.1103/PhysRevB.61.5600

[244] D. G. Blair, The Detection of Gravitational Waves, 1991. https://www.cambridge.org/core/books/detection-of-gravitational-waves/4840FA275A002D3F5D271512DDB8C5CC

[245] Y. Sun, H. Tohmyoh, Thermoelastic damping of the axisymmetric vibration of circular plate resonators, Journal of Sound and Vibration 319 (2009) 392–405. https://www.sciencedirect.com/science/article/pii/S0022460X08005543

[246] S. Erbay, E. S. Şuhubi, Longitudinal wave propagation in a generalized thermoelastic cylinder, Journal of Thermal Stresses 9 (3) (1986) 279–295.

https://www.tandfonline.com/doi/abs/10.1080/01495738608961904?journalCode=uths20

[247] Y. S. Touloukian, E. H. Buyco, Thermophysical Properties of Matter: Specific Heat of Metallic Elements and Alloys, Vol. 4, 1971. https://www.osti.gov/biblio/5439707-thermophysical-properties-matter-tprc-data-series-volume-specific-heat-metallic-elements-alloys-reannouncement-data-book

[248] M. C. Masso Herrera, Properties of bonded silicon for future generations of gravitational wave observatories, Ph.D. thesis, University of Glasgow (2019). http://theses.gla.ac.uk/41177/

[249] M. Granata, A. Amato, L. Balzarini, M. Canepa, J. Degallaix, D. Forest, V. Dolique, L. Mereni, C. Michel, L. Pinard, B. Sassolas, J. Teillon, G. Cagnoli, Amorphous optical coatings of present gravitational-wave interferometers, Classical and Quantum Gravity 37 (9) (2020) 095004. https://doi.org/10.1088/1361-6382/ab77e9

[250] M. Granata, L. Balzarini, J. Degallaix, V. Dolique, R. Flaminio, D. Forest, D. Hofman, C. Michel, R. Pedurand, L. Pinard, et al., Internal friction and young's modulus measurements of sio2 and ta2o5 films done with an ultra-high q silicon-wafer suspension, Versita 60 (2016) 1–7. http://hal.in2p3.fr/in2p3-01357312/document

[251] Montana Instruments, Technical Drawing Montana Instruments S50. https://www.montanainstruments.com/library/files/4106-510-DIM001.pdf

[252] R. Nawrodt, C. Schwarz, S. Kroker, I. W. Martin, R. Bassiri, F. Brückner, L. Cunningham, G. D. Hammond, D. Heinert, J. Hough, T. Käsebier, E. B. Kley, R. Neubert, S. Reid, S. Rowan, P. Seidel, A. Tünnermann, Investigation of mechanical losses of thin silicon flexures at low temperatures, Classical and Quantum Gravity 30 (11) (2013) 115008. http://arxiv.org/abs/1003.2893v1

[253] S. Reid, G. Cagnoli, D. R. Crooks, J. Hough, P. Murray, S. Rowan, M. M. Fejer, R. Route, S. Zappe, Mechanical dissipation in silicon flexures, Physics Letters, Section A: General, Atomic and Solid State Physics 351 (4-5) (2006) 205–211. https://arxiv.org/abs/gr-qc/0504134

[254] G. G. Stoney, C. A. Parsons, The tension of metallic films deposited by electrolysis, Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character 82 (553) (1909) 172–175. https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1909.0021

[255] L. Kuo, H. Pan, S. Wang, S. Chao, M. Principe, I. Pinto, J. Neilson, R. De Salvo, Annealing effect on the nano-meter scale tit.aniay silica multi-layers for mirror coatings of the laser interferometer gravitational waves detector, in: 2019 PhotonIcs Electromagnetics Research Symposium - Spring (PIERS-Spring), 2019, pp. 2437–2442. https://ieeexplore.ieee.org/abstract/document/9017449

[256] J. Steinlechner, I. W. Martin, How can amorphous silicon improve current gravitational-wave detectors?, Physical Review D 103 (2021) 042001. https://link.aps.org/doi/10.1103/PhysRevD.103.042001

[257] X. Liu, B. E. White, R. O. Pohl, E. Iwanizcko, K. M. Jones, A. H. Mahan, B. N. Nelson, R. S. Crandall, S. Veprek, Amorphous solid without low energy excitations, Physical Review Letters 78 (23) (1997) 4418–4421. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.78.4418

[258] S. Nayebossadri, D. Smith, J. Speight, D. Book, Amorphous Zr-based thin films fabricated by magnetron sputtering for potential application in hydrogen purification, Journal of Alloys and Compounds 645 (S1) (2015) S56–S60. http://dx.doi.org/10.1016/j.jallcom.2015.01.230

[259] H. Müller, S. W. Chiow, Q. Long, S. Herrmann, S. Chu, Atom interferometry with up to 24-photon-momentum-transfer beam splitters, Physical Review Letters 100 (18) (2008) 3–6. http://arxiv.org/abs/0712.1990v2

[260] T. Hong, H. Yang, E. K. Gustafson, R. X. Adhikari, Y. Chen, Brownian thermal noise in multilayer coated mirrors, Physical Review Letters D 87 (2013) 082001. https://link.aps.org/doi/10.1103/PhysRevD.87.082001

[261] H. Mathur, K. Jones-Smith, A. Lowenstein, An analysis of the LIGO discovery based on Introductory Physics (2016) 1–18. http://arxiv.org/abs/1609.09349

[262] The LIGO Scientific Collaboration, The Virgo Collaboration, GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, Physical Review Letters 119 (16) (2017) 30–33. http://arxiv.org/abs/1710.05836v2

[263] S. D. Penn, P. H. Sneddon, H. Armandula, J. C. Betzwieser, G. Cagnoli, et al., Mechanical loss in tantala/silica dielectric mirror coatings, Classical and Quantum Gravity 20 (13) (2003) 2917–2928. https://doi.org/10.1088/0264-9381/20/13/334

[264] COMSOL, COMSOL Multiphysics User's Guide, The Heat Transfer Branch (2012) 709–745. https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/COMSOL_ReferenceManual.pdf

[265] Solis Beijing Co. Ltd, Fused silica blanks (2020). http://www.optoelect.com/2018/Quartz-Glass_0411/11.html

[266] S. Chen, M. Zhu, D. Li, H. He, Y. Zhao, J. Shao, Z. Fan, Effects of electric field distribution and pulse duration on the ultra-short pulse laser damage resistance of laser coatings, Vol. 7842, 2010, p. 78420D. http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.867225

[267] E. Bertschinger, E. F. Taylor, Gravitational Waves, Physics 456 (1) (2010) 1–26. http://arxiv.org/abs/gr-qc/9506086v1

[268] D. A. Coulter, R. J. Foley, C. D. Kilpatrick, M. R. Drout, A. L. Piro, et al., Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source, Science 358 (6370) (2017) 1556–1558.

[269] H. de Waal, Influence of Proton Exchange on Internal Friction in Alkali Silicate Glasses, Journal of the American Cermaic Society 52 (3). https://ceramics.onlinelibrary.wiley.com/doi/10.1111/j.1151-2916.1969.tb11206.x

[270] The LIGO Scientific Collaboration, Instrument Science White Paper, LIGO Technical Note (225). https://dcc.ligo.org/LIGO-T2000407/public

[271] S. Reinsch, R. Müller, J. Deubener, H. Behrens, Internal friction of hydrated soda-lime-silicate glasses, The Journal of Chemical Physics 139 (17) (2013) 174506. http://aip.scitation.org/doi/10.1063/1.4828740

[272] L. Yu, Internal thermal noise in the LIGO test masses: A direct approach, Physical Review D - Particles, Fields, Gravitation and Cosmology 57 (2) (1998) 1–4. https://journals-aps-org.libproxy.mit.edu/prd/pdf/10.1103/PhysRevD.57.659$%$0Ahttps://journals.aps.org/prd/pdf/10.1103/PhysRevD.57.659

[273] G. M. Harry, H. Armandula, E. Black, D. R. Crooks, G. Cagnoli, et al., Optical coatings for gravitational wave detection, in: J. D. T. Kruschwitz, J. B. Oliver (Eds.), Advances in Thin Film Coatings for Optical Applications, Vol. 5527, International Society for Optics and Photonics, SPIE, 2004, pp. 33 – 40. https://doi.org/10.1117/12.555780

[274] S. J. Waldman, The Advanced LIGO Gravitational Wave Detector (2011). https://arxiv.org/abs/1103.2728

[275] G. Parisi, F. Zamponi, The ideal glass transition of hard spheres, Journal of Chemical Physics 123 (14) (2005) 144501. http://arxiv.org/abs/cond-mat/0506445v1

[276] N. Andersson, K. D. Kokkotas, The r-mode instability in rotating neutron stars, International Journal of Modern Physics D 10 (04) (2001) 381–441. http://dx.doi.org/10.1142/S0218271801001062

[277] S. Gras, K. Kuns, P. Fritschel, Point absorbers in aLIGO coating, no. December, 2019, pp. 1–16. https://dcc.ligo.org/LIGO-G1902305

[278] D. A. Buckley, I. Andreoni, S. Barway, J. Cooke, S. M. Crawford, et al., A comparison between SALT/SAAO observations and kilonova models for AT 2017gfo: The first electromagnetic counterpart of a gravitational wave transient - GW170817, Monthly Notices of the Royal Astronomical Society: Letters 474 (1) (2018) L71–L75. http://arxiv.org/abs/1710.05855v1

[279] R. Flaminio, J. Franc, C. Michel, N. Morgado, L. Pinard, B. Sassolas, A study of coating mechanical and optical losses in view of reducing mirror thermal noise in gravitational wave detectors, Classical and Quantum Gravity 27 (8) (2010) 084030. https://iopscience.iop.org/article/10.1088/0264-9381/27/8/084030/meta

[280] S. Hild, S. Chelkowski, A. Freise, J. Franc, N. Morgado, R. Flaminio, R. DeSalvo, A xylophone configuration for a third-generation gravitational wave detector, Classical and Quantum Gravity 27 (1) (2009) 015003. http://dx.doi.org/10.1088/0264-9381/27/1/015003

[281] E. Gustafson, P. Fritschel, Advanced LIGO H1 Optical Layout (2016). https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=8135

[282] S. Nayebossadri, D. Smith, J. Speight, D. Book, Amorphous Zr-based thin films fabricated by magnetron sputtering for potential application in hydrogen purification, Journal of Alloys and Compounds 645 (S1) (2015) S56–S60. http://dx.doi.org/10.1016/j.jallcom.2015.01.230

[283] H. Mathur, K. Brown, A. Lowenstein, An analysis of the LIGO discovery based on introductory physics, American Journal of Physics 85 (9) (2017) 676–682. http://dx.doi.org/10.1119/1.4985727

[284] R. Bork, R. Abbott, D. Barker, J. Heefner, An overview of the LIGO control and data acquisition systems, 2001. https://accelconf.web.cern.ch/ica01/papers/tubi001.pdf

[285] T. W. Tu, S. Y. Lee, Analytical Solution of Heat Conduction for Hollow Cylinders with Time-Dependent Boundary Condition and Time-Dependent Heat Transfer Coefficient, Journal of Applied Mathematics 2015. https://www.hindawi.com/journals/jam/2015/203404/

[286] ANSYS Inc, ANSYS Fluent Theory Guide (November 2013). http://www.afs.enea.it/project/neptunius/docs/fluent/html/th/main_pre.htm

[287] A. Heptonstall, G. Cagnoli, J. Hough, S. Rowan, Characterisation of mechanical loss in synthetic fused silica ribbons, Physics Letters, Section A: General, Atomic and Solid State Physics 354 (5-6) (2006) 353–359. https://www.sciencedirect.com/science/article/pii/S0375960106001630

[288] C. Bourgeois, E. Steinsland, N. Blanc, N. F. de Rooij, Design of resonators for the determination of the temperature coefficients of elastic constants of monocrystalline silicon (1997) 791–799. https://ieeexplore.ieee.org/document/639192?reload=true$&$arnumber=639192&contentType=Conference$%$20Publications

[289] C. Bourgeois, E. Steinstand, N. Blanc, d. R. N.F, Design Of Resonators For the Determination Of the Temperature Coefficients Of Elastic Constants Of Monocrystalline Silicon, IEEE International Frequency Control Symposium (59) (1997) 791–799. http://www.scopus.com/inward/record.url?eid=2-s2.0-79959555301&partnerID=40&md5=59399bbf1cf7eb5ae1cb01a9753e0c4d

[290] C. J. Myers, R. C. Allen, Development of an analytical mirror model addressing the problem of thermoelastic deformation, Applied Optics 24 (13) (1985) 1933. https://www.osapublishing.org/ao/abstract.cfm?uri=ao-24-13-1933

[291] M. Fletcher, S. Tait, J. Steinlechner, I. W. Martin, A. S. Bell, J. Hough, S. Rowan, R. Schnabel, Effect of stress and temperature on the optical properties of silicon nitride membranes at 1550 nm, Frontiers in Materials 5. http://eprints.gla.ac.uk/154641/

[292] D. R. Queen, X. Liu, J. Karel, T. H. Metcalf, F. Hellman, Excess specific heat in evaporated amorphous silicon, Physical Review Letters 110 (13) (2013) 135901. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.110.135901

[293] M. Abernathy, A. Amato, A. Ananyeva, S. Angelova, B. Baloukas, R. Bassiri, et al., Exploration of co-sputtered ta_2o_5-zro_2 thin films for gravitational-wave detectors. https://arxiv.org/abs/2103.14140

[294] Y.-B. Yi, Finite Element Analysis of Thermoelastic Damping in Contour-Mode Vibrations of Micro- and Nanoscale Ring, Disk, and Elliptical Plate Resonators, Journal of Vibration and Acoustics 132 (4) (2010) 1–7. https://asmedigitalcollection.asme.org/vibrationacoustics/article-abstract/132/4/041015/442758/Finite-Element-Analysis-of-Thermoelastic-Damping?redirectedFrom=fulltext

[295] S. Stöttinger, G. Hinze, G. Diezemann, I. Oesterling, K. Müllen, T. Basché, Impact of local compressive stress on the optical transitions of single organic dye molecules, Nature Nanotechnology 9 (3) (2014) 182–186. https://www.nature.com/articles/nnano.2013.303

[296] S. Reinsch, R. Müller, J. Deubener, H. Behrens, Internal friction of hydrated soda-lime-silicate glasses, Journal of Chemical Physics 139 (17) (2013) 1–7. https://aip.scitation.org/doi/10.1063/1.4828740

[297] L. Yu, Internal thermal noise in the LIGO test masses: A direct approach, Physical Review D - Particles, Fields, Gravitation and Cosmology 57 (2) (1998) 1–4. https://journals-aps-org.libproxy.mit.edu/prd/pdf/10.1103/PhysRevD.57.659$%$0Ahttps://journals.aps.org/prd/pdf/10.1103/PhysRevD.57.659

[298] F. Nürnberg, M. Stamminger, B. Kuehn, M. Altwein, R. Takke, IR grade fused silica for high-power laser applications. https://www.researchgate.net/publication/267024486_IR_grade_fused_silica_for_high-power_laser_applications

[299] K. Takimoto, A. Fukuta, Y. Yamamoto, N. Yoshida, T. Itoh, S. Nonomura, Linear thermal expansion coefficients of amorphous and microcrystalline silicon films, Journal of Non-Crystalline Solids 299-302 (2002) 314–317. https://www.sciencedirect.com/science/article/pii/S0022309302009304

[300] P. B. Karlmann, K. J. Klein, P. G. Halverson, R. D. Peters, M. B. Levine, D. Van Buren, M. J. Dudik, Linear thermal expansion measurements of single crystal silicon for validation of interferometer based cryogenic dilatometer, AIP Conference Proceedings 824 I (2000) (2006) 35–42. https://aip.scitation.org/doi/abs/10.1063/1.2192331?journalCode=apc

[301] D. R. M. Crooks, Mechanical Loss and its Significance in the Test Mass Mirrors of Gravitational Wave Detectors, Ph.D. thesis, University of Glasgow (2002). http://theses.gla.ac.uk/http://theses.gla.ac.uk/2893/

[302] J. Tauc, Optical properties and electronic structure of amorphous ge and si, Materials Research Bulletin 3 (1) (1968) 37 – 46. http://www.sciencedirect.com/science/article/pii/0025540868900238

[303] T. Quinn, C. Speake, R. Davis, W. Tew, Stress-dependent damping in cube torsion and flexure suspensions at stresses up to 1.1 gpa, Physics Letters A 197 (3) (1995) 197 – 208. http://www.sciencedirect.com/science/article/pii/037596019400921B

[304] LIGO Collaboration, Virgo Collaboration, Search for gravitational-wave bursts associated with gamma-ray bursts using data from LIGO science run 5 and virgo science run 1, The Astrophysical Journal 715. https://arxiv.org/abs/0908.3824

[305] A. Alexandrovski, R. K. Route, M. M. Fejer, Absorption Studies in Sapphire, 2001. https://dcc.ligo.org/DocDB/0033/G010152/000/G010152-00.pdf

[306] K. Craig, Studies of the mechanical dissipation of thin films for mirrors in interferometric gravitational wave detectors, Ph.D. thesis, University of Glasgow (2015). https://theses.gla.ac.uk/6582/

[307] S. Hild, S. Chelkowski, A. Freise, J. Franc, N. Morgado, R. Flaminio, R. DeSalvo, A xylophone configuration for a third-generation gravitational wave detector, Classical and Quantum Gravity 27 (1) (2009) 015003. http://dx.doi.org/10.1088/0264-9381/27/1/015003

[308] K. Danzmann, E. al, LISA ESA L3 Proposal, Esa. https://www.elisascience.org/files/publications/LISA{_}L3{_}20170120.pdf

[309] K. Craig, J. Steinlechner, P. G. Murray, A. S. Bell, R. Birney, et al., Mirror coating solution for the cryogenic einstein telescope, Physical Review Letters 122 (2019) 231102. https://link.aps.org/doi/10.1103/PhysRevLett.122.231102

[310] W. Street, PEEK (PolyEtherEtherKetone) Specifications (since 1998) (2009) 1–4. http://www.boedeker.com/peek_p.htm

[311] L. P. Singer, L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, Physical Review D 93 (2) (2016) 1–24. http://arxiv.org/abs/1508.03634v6

[312] D. Smith, P. Baumeister, Refractive index of some oxide and fluoride coating materials, Applied Optics 18 (1) (1979) 111. https://www.osapublishing.org/ao/abstract.cfm?URI=ao-18-1-111

[313] S. Fairhurst, Source localization with an advanced gravitational wave detector network, Classical and Quantum Gravity 28 (10) (2011) 105021. http://arxiv.org/abs/1010.6192v2

[314] J. P. Masse, H. Szymanowski, O. Zabeida, A. Amassian, J. E. Klemberg-Sapieha, L. Martinu, Stability and effect of annealing on the optical properties of plasma-deposited Ta_2O_5 and Nb_2O_5films, Thin Solid Films 515 (4) (2006) 1674–1682. https://www.sciencedirect.com/science/article/pii/S0040609006006882

[315] A. J. Levan, J. D. Lyman, N. R. Tanvir, J. Hjorth, I. Mandel, E. R. Stanway, D. Steeghs, A. S. Fruchter, E. Troja, S. L. Schrøder, , others., The environment of the binary neutron star merger gw170817, The Astrophysical Journal 848 (2) (2017) L28. http://dx.doi.org/10.3847/2041-8213/aa905f

[316] D. G. Blair, J. Ferreirinho, Thermoelastic effect in niobium at the superconducting transition, Physical Review Letters 49 (1982) 375–378. https://link.aps.org/doi/10.1103/PhysRevLett.49.375

[317] A. Einstein, Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen, Annalen der Physik 322 (8) 549–560. https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19053220806

[318] Montclaire State University, Sounds of Space Time (2016). http://www.soundsofspacetime.org/detection.html

[319] R. C. Douglas, Aspects of hydroxide catalysis bonding of sapphire and silicon for use in future gravitational wave detectors, Ph.D. thesis, University of Glasgow (2017). http://theses.gla.ac.uk/7993/

[320] D. Becker, D. Collinson, M. Chiesa, R. Kappl, N. Khan, S. S. Eaton, G. Smith, P. Riedi, Electron Paramagnetic Resonance, Vol. 20 of SPR - Electron Paramagnetic Resonance, The Royal Society of Chemistry, 2006. http://dx.doi.org/10.1039/9781847557568

[321] B. S. Berry, W. C. Pritchet, Vibrating Reed Internal Friction Apparatus for Films and Foils., IBM Journal of Research and Development 19 (4) (1975) 334–343. https://dl.acm.org/doi/abs/10.1147/rd.194.0334

[322] IOLITEC, Technical data sheet, personal correspondence (2005). https://iolitec.de/en

[323] W. Yam, S. Gras, M. Evans, Multimaterial coatings with reduced thermal noise, Physical Review D 91 (2015) 042002. https://link.aps.org/doi/10.1103/PhysRevD.91.042002

[324] R. Flaminio, J. Franc, C. Michel, N. Morgado, L. Pinard, B. Sassolas, A study of coating mechanical and optical losses in view of reducing mirror thermal noise in gravitational wave detectors, Classical and Quantum Gravity 27 (8) (2010) 084030. https://doi.org/10.1088/0264-9381/27/8/084030

[325] A. Bell, A. Markosyan, Progress towards delivering silicon testmasses for cryogenic detectors, 2017, pp. 1–12. https://events.ego-gw.it/indico/getFile.py/access?contribId=6&sessionId=0&resId=1&materialId=slides&confId=52

[326] S. Gras, M. Evans, Direct measurement of coating thermal noise in optical resonators, Phys. Rev. D 98 (2018) 122001. https://link.aps.org/doi/10.1103/PhysRevD.98.122001

[327] N. Demos, S. Gras, M. Evans, MIT Coating Thermal Noise Group Update (2020). https://dcc-lho.ligo.org/LIGO-G2001592

[328] K. Srinivasan, Coating Strain Induced Distortion in LIGO Optics Abstract : Introduction :, LIGO Technical Journal (1997) 1–27. https://dcc.ligo.org/public/0028/T970176/000/T970176-00.pdf

[329] R. Flaminio, J. Franc, C. Michel, N. Morgado, L. Pinard, B. Sassolas, A study of coating mechanical and optical losses in view of reducing mirror thermal noise in gravitational wave detectors, Classical and Quantum Gravity 27 (8). https://iopscience.iop.org/article/10.1088/0264-9381/27/8/084030/pdf